To be presented at the International Conference on Robotics and Automation, 2002, Washington, D.C.

Detecting Surface Features During Locomotion Using Optic Flow

M. ANTHONY LEWIS

Iguana Robotics, Inc.
P.O. Box 628
Mahomet, IL 61853
tlewis@iguana-robotics.com

Abstract- *We test the hypothesis that: (1) Optic flow can be used to detect significant environmental features during locomotion in a biped, even given significant up and down movement and jarring of the robot during locomotion. (2) Reliable detection is only possible if a* prediction *of the expected optic flow field is made at each instance. This prediction should be driven by the phase of the robot's gait as well as other information about the state of the robot. (3) This prediction can be accomplished in a distributed, biologically plausible framework.*

Our results using a walking biped mechanism strongly support this hypothesis and indicate that optic flow is a viable strategy and that the prediction of optic flow is a critical component in this behavior.

We believe this is the first time that this approach has been used and demonstrated successfully on a walking machine.

1.0 Introduction

For walking robots to realize their full potential, distal environment sensing must be tightly integrated with the walking cycle. Distal sensing is crucial to allow anticipatory gait adjustment to accommodate varying terrain. Close coupling of the visual and locomotor cycle can lead to rapid, adaptive adjustment of the robot. Biological systems have a tightly integrated action perception cycle, unlike that seen in most legged robots.

In the 1950's, Gibson pointed out the importance of the flow of the 'optic array' (i.e. *optic flow),* a visual cue arising from relative movement of the environment, in the control of human locomotion [1]. Relative motion of the environment can reveal environment structure. His pioneering ideas have influenced generations of experimentalists in the brain and behavioral sciences.

In robotics, the route to using optic flow has come principally from studies of insects. Insect motion processing pathways are relatively well understood. There have been a number of robotic implementations that use optic flow as its primary input for visual navigation. For a flavor of that work, see [2-9].

Vision is a particularly elegant, multifunctional distal sense. It is not surprising that biological systems have relied so heavily on vision to guide locomotion.

However, there has been little or no work using optic flow to control *legged* locomotion although there has been interest in other visual techniques for the control of walking machines (visual servoing, color, stereopsis, terrain reconstruction using laser ranging). Given the strong belief in the biological community that optic flow is a key component of visuomotor coordination (e.g. [10-12]), and that special purpose chips are becoming available which can rapidly compute optic flow components, we decided to investigate how optic flow might be used to control a legged robot.

In previous work [13], we have shown how a robot can learn to step over obstacles using a reinforcement paradigm. The control architecture featured stereo input, and uses sensory expectation (i.e. building up an expectation of the surface in front of the robot), detection of surface anomalies (i.e. an obstacle) and the association of those novelties with motor behavior (control of foot placement and triggering when to step over the obstacle). A key observation of that work was that what constitutes an 'obstacle' is dependent on the capabilities of the robot. The robot learned to perceive only those 'obstacles' that could potentially destabilize it.

Here we focus on replacing the front-end of this algorithm with a new visual cue: optic flow. It is not clear whether optic flow is stable enough to allow detection of fine surface features. Our experiments here are designed to evaluate if this is possible.

The main hypothesis which we will test is that: (1) Optic flow can be used to detect potentially destabilizing environmental features during locomotion in a biped, even given significant up and down movement and jarring of the robot during locomotion. (2) Reliable detection is only possible if a prediction of the expected optic flow field is made at each instance. This prediction should be driven by the phase of the robot's gait as well as other information about the state of the robot. (3) This prediction can be accomplished in a distributed, biologically plausible framework.

The architecture of the model is biologically 'plausible.' The architecture used is motivated by the architecture of the vertebrate nervous system. Space limitation prevents a complete discourse on this topic. The biological details will be presented in a forthcoming article. An overview is given in [13].

2.0 Materials and Methods

2.1 The Robot Mechanism

A 20-cm tall tethered biped is used in the following experiments (See Fig. 1). The tether allows forward/backward and up/down translation of the body. The hip's rotation is held fixed. Two miniature cameras give the robot a view from its feet to about 40 cm in front of the robot. The robot itself uses 4 hobby type servos to actuate the limbs (Futaba 3002 for hips and 5203 for knees). The servos are controlled by a custom board (ServoX24 Board, Digital Designs and Systems, Cambridge, MA). Joint level commands are streamed over an RS-232 port from a Win 2000 Laptop (550Mhz Dell). Visual computation is performed on a Linux workstation (1.7 GHz Dell).

2.1.1 Locomotor Generator

The locomotor pattern is generated using a Central Pattern Generator (CPG) network. The network also generates a distributed representation (i.e. a Population Code) of the phase of the left and right halves of the CPG network. In general, this phase information is used to coordinate visual processing with the step cycle. In the particular case presented here, this phase information and joint information is used to predict the expected optic flow at 18 horizontal band regions in the incoming image. Balance is not a concern here as the tether provides a great deal of stability to the biped. The walking speed of the robot was about 2 cm/sec.

2.2 Computing Optic Flow Events

Normal flow is computed using an algorithm previously described. [14]. The components of flow are computed as:

$$\frac{dx}{dt} \propto \int \frac{E_t(\nabla E \cdot [1 \quad 0]^T)}{\|\nabla E\|^2}$$

$$\frac{dy}{dt} \propto \int \frac{E_t(\nabla E \cdot [0 \quad 1]^T)}{\|\nabla E\|^2}$$

where dx/dt is the yaw velocity and dy/dt is the pitch velocity of a region of the image, E_t is the temporal derivative of a pixel in the image, and ∇E is the spatial gradient. The integral indicates averaging over a small patch of the image.

We have shown that this algorithm will give a good estimate of the true optic flow if the statistics of spatial gradient orientations are balanced in their statistical distribution (e.g. with the assumption of a flat distribution for edge orientation [14]). The results should also hold here where the ratio of horizontal and vertical edges is about 1:1.

Information is selectively filtered based on an optic flow "event' representation. For an optic flow estimate to be forwarded to further stages of processing, we check that E_t and ∇E are of sufficient magnitude to give a reliable estimate of optic flow. Thus, if sparse edges are available, few 'events' will be sent to latter stages of processing. The lack of an event does not imply that there is or is not an environmental surface present. The lack of an event only means that there is insufficient evidence to

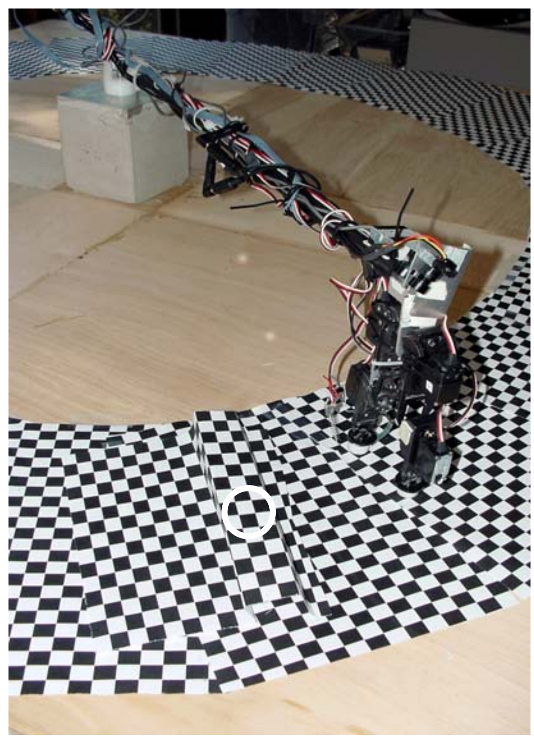

Figure 1. Walking biped mechanism and camouflaged obstacle. Obstacle has the same texture pattern as track. A white circle is placed on the center of the obstacle to help the reader localize it.

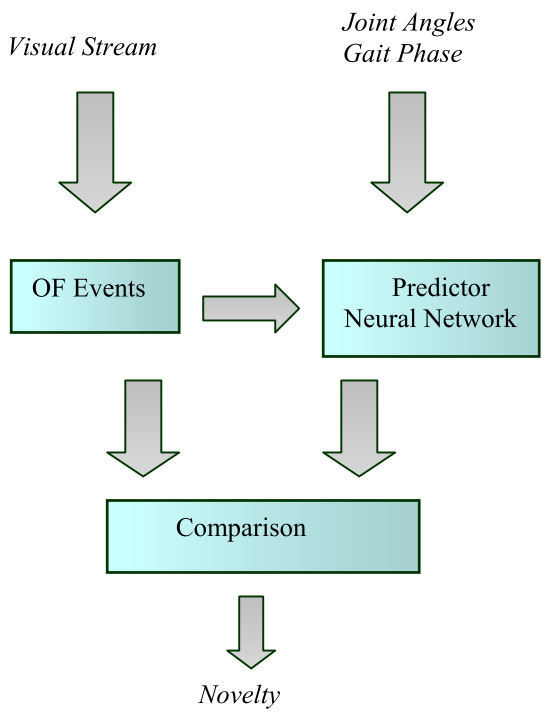

Figure 2. Prediction of Optic Flow Processing. Based on Gait phase and joint angle, a Neural network predicts possible OF events in the image. Simultaneously, a visual stream is processed and optic flow events are generated from real data. The Neural Predictor updates its estimates. A comparison module computes the difference between the actual and expected optic flow events. The comparator adjusts its sensitivity based on statistical assumptions about the frequency of 'novel stimuli' in the environment. If the rate of novel stimuli reaches a predefined threshold. the robot stops, having detected an environmental anomaly. In future implementation, this anomaly detection may be used to adjust foot placement and trigger the robot to step over the obstacles. See a description of how this was done using another visual cue, stereopsis, in [13].

decide one way or another. This idea of detecting 'events' is basic to communication between "Neuromorphic" chips (i.e. chips based on biological process, some of which can compute 1- and 2-d flow fields at high rates). The interested reader is referred to [15, 16] for more detail.

2.2.1 Optic Flow During Locomotion

The central problem of using Optic Flow to guide a *walking* robot during locomotion is that the up and down (and sideways) movement of the robot generates an Optic Flow field that varies significantly during a gait cycle, even when stepping over flat surfaces. This robot captures the key features of this variable movement (except sideways movement due to the tether constraint).

Because the movement is periodic, we hypothesize that at different points in the step cycle, there is a different likelihood of a given optic flow vector occurring at a point in the image. If we can *predict* the most likely heading and its magnitude, we can then detect anomalous conditions in the flow field. We describe our approach to this below.

2.2.2 Generating Optic Flow Expectations

In previous work [13] we describe a walking robot that learns to step over obstacles. A key component of this work was the generation of 'sensory expectations.' Here, we rely on the same computation to generate expectancies of optic flow. Referring to Figure 2, camera information is processed as described in section 2.1.2. Optic flow events are then compared to the predicted optic flow values. The prediction is constantly updated and improved by observing the difference between the actual optic flow values and predicted values. The adaptive sensitivity seeks to produce a certain probability of detecting events. This is based on assumptions about the statistics of anomalies in the environment. In the work here, we assume 5% as the expected rate of anomalies per optic flow region per step.

As the focus of this article is on the optic flow, the details of the predictor are not given here.

2.2.3 Criteria for 'stopping' the Robot

Based on the 5% threshold for each firing, and given that in our model about 18 cells have the potential for detecting an obstacle, we set the threshold for obstacle detection at 140% of the expected spontaneous firing (in practice, this threshold can be found optimally for a given application). The detection is based on a running average of 'novelty pixels' with a time constant of about 1 step cycle.

2.2.4 Obstacles

Figure 3 (A) shows an obstacle from the robot's view. As can bee seen, the obstacle is well camouflaged. In figure 3 (B) we see the same obstacle from the side.

Because of the camouflage, the robot must use optic flow or a similar cue to detect the obstacle.

The obstacles were constructed from paper printed with a special pattern as shown in the figure. The paper was then folded appropriately to form obstacles of the needed heights.

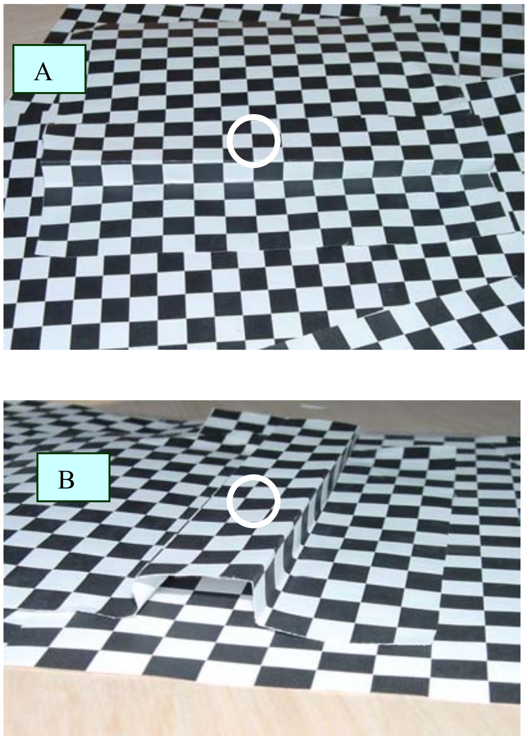

Figure 3. (A) Obstacle from robot view is almost 'invisible.' We placed a white circle on the obstacle to help the reader find it. (B) The same obstacle, side view.

2.3 Experimental Procedure

2.3.1 Statistics of the optic flow vector

We first wanted to determine what the statistical distribution of the optic flow vector was during locomotion. We computed the optic flow field as described above. We then computed the angle of the flow vectors ($\theta = a\tan 2(dy/dt, dx/dt)$ and the magnitude ($mag = \sqrt{(dx/dt)^2 + (dy/dt)^2}$). We also collected data on the frequency at which each histogram value was chosen.

We then computed a magnitude histogram and a frequency histogram of this data. This information gives us some quantitative idea of the statistical variation of the flow field during normal locomotion.

2.3.2 Does the Use of Prediction Improve Obstacle Detection Performance?

To test the hypothesis that prediction of optic flow is necessary for effective detection of the obstacles, we performed the following experiment.

We tested the case of Prediction versus no-prediction in detecting obstacles. The no-prediction cases are the control group. We used 5 different height obstacles: 1, 2, 3, 4, and 5 cm. All obstacles were 3 cm wide and about 15 cm across.

For each obstacle, and within each group, a total of 10 trials were performed that resulted in either a collision with the obstacle or the robot stopping short of the obstacle. For 5 obstacles, and for the prediction and the control cases, we performed a total of 100 experiments (5x2x10).

We also recorded the number of times the robot stopped before encountering an obstacle. This number was very low in both groups.

The strategy used to detect obstacles needed to work in both cases. We conditioned the data in the following way. Animals normally fixate on aspects of the environment using eye movement. In our robot, the eyes are fixed to the body. To emulate fixation, we used a center patch in the image as a reference point. We then subtracted the velocity of this reference point from all of the other optic flow velocities in the image. In this way we produced an approximation to a stabled, fixed gaze in the environment.

3.0 Results

3.1 Statistics of Optic Flow

Figure 4 (A) shows the distribution of optic flow vector magnitudes versus direction (0 to 360 degrees). As can be seen, there is a significant probability of an optic flow vector taking on *any* possible direction. In Figure 4 (B) we see however that the number of times vectors between 80 and 100 degree (in the direction of movement of the robot) are chosen is far higher than in other directions. However, there is a better than 50% possibility that the robot has an optic flow vector not pointed in the major direction.

It seems unlikely that a simple scheme such as detecting optic flow above a certain threshold would work. As seen in Figure 4(B), sometimes the optic flow field reverses itself during locomotion on this particular robot. A more sophisticated procedure, as implemented here, is warranted.

It seems likely that the probability distribution is related to the phase of gait. This hypothesis is tested in the next experiment. If the probability of choosing a particular direction is a function of the phase of gait, then we should be able to significantly improve our performance in the task of detecting fine surface features.

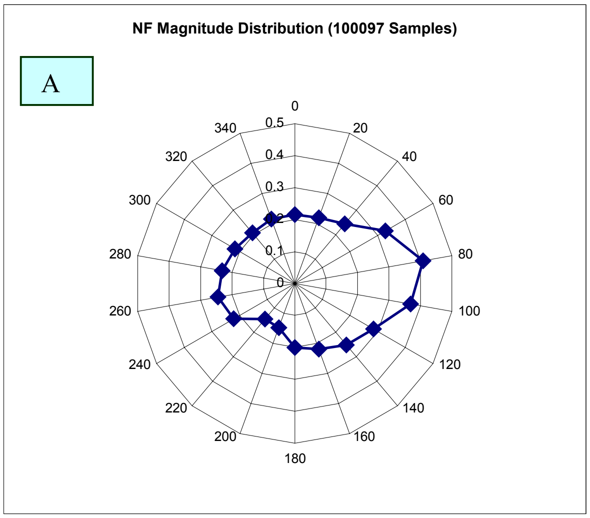

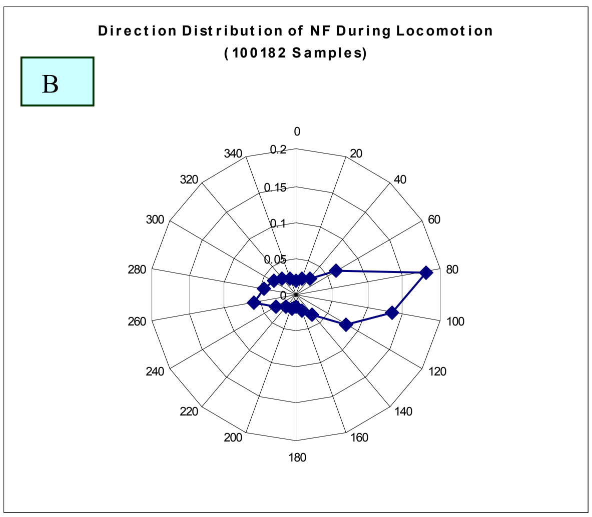

Figure 4. Statistical distribution of Optic Flow vectors during normal locomotion. (A) the average magnitude of optic flow versus vector angle (B) The probability of choosing any one vector at a given time.

3.2 *Use of prediction*

Figure 5 shows the results of the experiment in the use of prediction versus control in the task of detecting anomalous surface features. The lower line is the control. As can be seen, the robot can detect obstacles reliably when the obstacle is about 4 CM tall in the control case. However, in the case where prediction is used, the results are much better. The obstacle is detected reliably even when the obstacle is 1 cm tall.

These results indicate that the use of prediction significantly improves the detection of fine obstacles.

We also computed how many times the biped stopped when no obstacle was present. In the case of the use of prediction, it stopped 7 times out of a total of 107 trials.

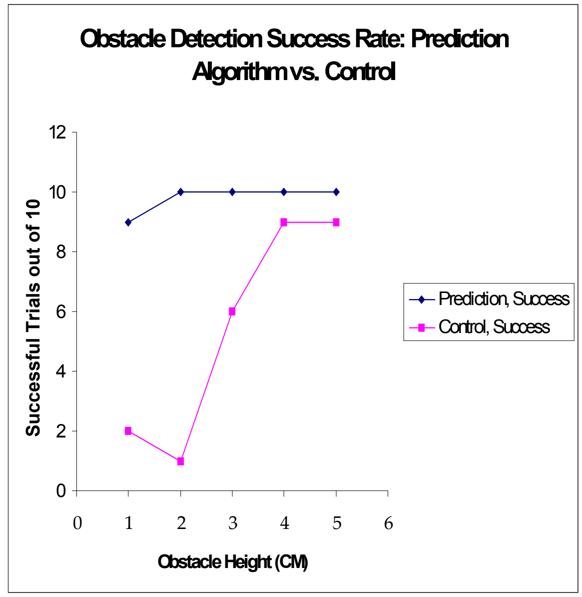

Figure 5.Obstacle detection success rate.

In the control group, it stopped 2 times in 102 trials. Future experimentation will allow us to determine if this effect is significant, and, if so, the cause of this slight difference.

4.0 Conclusions

Our results lend compelling evidence to the idea that (1) it is possible to predict the optic flow based on the phase of gait and kinematic configuration of the robot. (2) This prediction is necessary (at least in this robot) for the reliable detection of surface features.

In previous work, we showed how a robot could learn to adjust its foot placement and timing to step over obstacles using stereopsis as a cue[13]. It is likely that optic flow can serve a similar function.

The results presented here are surprising. Optic flow signals can be relatively noisy. The detection of small amounts of changes in the optic flow field, especially with a walking robot seems a difficult problem. Further work needs to be done to fully characterize the performance of the performance of this methodology. For example, what is the smallest object that can be reliably detected? Can the size of the obstacle be deduced from the optic flow signal? Finally, in our system optic flow and stereopsis give redundant information about the environment. Under which condition is a visual cue dominant in the task of avoiding obstacle and determining the transverability of the path immediately in front of the robot? What is the relationship between walking speed and visuomotor performance?

Acknowledgments

The author wishes to thank Yoichiro Endo of the Georgia Institute of Technology for many interesting discussion related to this work, and for the Linux vision setup and the vision processing code. The author thanks Dr. Richard Reeve and others at the 2001 Telluride Neuromorphic Engineering workshop for useful comments related to this work. The author also thanks Francesco Tenore (Johns Hopkins University) for a careful reading of the manuscript and valuable comments.

The author gratefully acknowledges support of ONR Grant No. N00014-99-0984.

5.0 References

[1] J. J. Gibson, "Visually Controlled Locomotion and Visual Orientation in Animals," *British Journal of Psychology*, vol. 49, pp. 182-194, 1958.

[2] M. V. Srinivasan and S. Venkatesh, *From Living Eyes to Seeing Machines*. UK: Oxford University Press, 1997.

[3] N. Franceschini, J. M. Pichon, and C. Blanes, "From Insect vision to Robot Vision," *Phil. Trans. R. Soc. Lond B*, pp. 283-294, 1992.

[4] R. R. Harrison, "An Analog VLSI Motion Sensor Based on the Fly Visual System," Ph.D. Thesis, *CNS*. California Institute of Technology, 2000.

[5] K. Weber, S. Venkatesh, and M. V. Srinivasan, "Insect inspired behaviours for the autonomous control of mobile robots," presented at Proceedings of the 13th International Conference on Pattern Recognition, 1996.

[6] P. J. Sobey, "Active navigation with a monocular robot," *Biol Cybern*, vol. 71, pp. 443-440, 1996.

[7] M. A. Lewis and M. E. Nelson, "Look Before You Leap: Peering Behavior for Depth Perception," presented at Simulation of Adaptive Behavior, Zurich, 1998.

[8] M. A. Lewis, "Visual Navigation in a Robot using Zig-Zag Behavior," presented at Advances in Neural Information Processing Systems (NIPS*10)., 1998.

[9] D. Coombs, M. Herman, T. Hong, and M. Nashman, "Real-time obstacle avoidance using central flow divergence and peripheral flow," *IEEE Trans. on Robotics and Automation*, vol. 14, pp. 49-59, 1998.

[10] D. N. Lee, "The Optic Flow Field: the Foundation of Vision," *Phil. Trans. R. Soc. Lond. B*, vol. 290, pp. 169-179, 1980.

[11] D. N. Lee and J. R. Lishman, "Visual proprioceptive control of stance," *Journal of Human Movement Studies*, vol. 1, pp. 87-95, 1975.

[12] W. H. Warren, B. A. Kay, W. D. Zosh, A. P. Duchon, and S. Sahuc, "Optic Flow Is Used to Control Human Walking," *Nature Neuroscience*, vol. 4, pp. 213-216, 2001.

[13] M. A. Lewis and L. S. Simó, "Certain Principles of Biomorphic Robots," *Autonomous Robots*, vol. 11, pp. 221-226, 2001. (*A preprint is available at: http://www-iguana-robotics.com)*

[14] M. A. Lewis, "Egomotion Computation Using Template Matching and Inner Product of Optic Flow Filter," presented at Conference on Information Sciences and Systems, Baltimore, MD, 1999.

[15] J. Lazzaro and J. Wawrzynek, "A Multi-Sender Asynchrounous Extension to the AER Protocol," presented at Proceedings of 1995 Conference on Advanced Research in VLSI, 1955.

[16] K. A. Boahen, "Communicating Neuronal Ensembles between Neuromorphic Chips," in *Neuromorphic Systems Engineering*, T. S. Lande, Ed. Boston, MA: Kluwer Academic Publishers, 1998, pp. 229-261.