The Dependence of Braking Strategies on Optical Variables in an Evolved Model of Visually-Guided Braking

Didem Kadihasanoglu[1], Randall D. Beer[1,2], and Geoffrey P. Bingham[1,3]

[1] Cognitive Science Program
[2] School of Informatics and Computing
[3] Department of Psychological and Brain Sciences
Indiana University, Bloomington, IN 47406 USA
{dikadiha,rdbeer,gbingham}@indiana.edu

Abstract. This paper presents results from two sets of experiments which investigate how strategies used by embodied dynamical agents in a simple braking task are affected by the perceptual information that the agents receive. Agents are evolved in a simple 2D environment containing one stationary object. The task of the agents is to stop as close as possible to the object without hitting it. The results of these experiments demonstrate that most of the evolved agents use an impulsive braking strategy, in which deceleration is not controlled continuously. Potential causes of this impulsive braking strategy and possible future directions are discussed.

Key words: Evolutionary robotics, visually-guided braking, image size, image expansion rate, tau, tau-dot

1 Introduction

There is a growing interest in applying the evolutionary approach to model experimental paradigms from psychology. For example, inspired by the psychological experiments such as double-TV-monitor experiments and perceptual crossing, Iizuka and Di Paolo [1] investigated how embodied agents establish live interactions and discriminate this type of interaction from the identical recorded motions. Rohde and Di Paolo [2] implemented an evolutionary robotics simulation to guide the analysis of empirical data on adaptation to sensory delays. In another work, Wood and Di Paolo [3] applied evolutionary robotics techniques to model the famous "A-not-B" error paradigm. Considering the sensory and motor capabilities of evolved model agents, ecological psychology in general, and control of locomotion in particular, provides an excellent research area in which evolutionary robotics techniques can be used.

Visual control of locomotion, which requires coordination between perception and action, is essential for any mobile agent, whether it is a human, animal or a robot, to move around, explore and interact with the world. One approach to the control of locomotion is based on internal representations such as world models and plans [4]. An alternative approach, which is developed by Gibson [5–7] is based on the idea that adaptive behavior is controlled by the perceptual information that is available to the observer. When an observer moves in an environment, a pattern of motion is produced at the eye of the observer called *the optic flow* [5, 7]. Optic flow provides information about the 3D layout of the environment, objects in the environment, and observer's self-motion through the environment, and can be used to control locomotion.

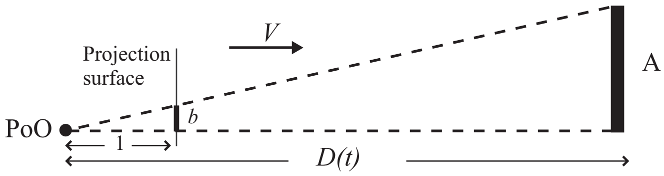

Fig. 1. A schematic view of an observer approaching an object of size A. PoO denotes the point of observation, D is the distance from the object, V is the observer's velocity, and b is the image size at time t. The retina of the observer (i.e. the projection surface) is approximated as being 1 unit distance away from the PoO.

In this paper, we explore how perceptual information received by an agent affects the behavior of the agent in the context of a simple braking task based on the experimental paradigm used in Yilmaz and Warren [8]. During a direct approach to an object, the image of the object on the observer's retina expands (see, Fig. 1). Gibson [6] argued that the rate of optical expansion could be used to control braking. There are a number of ways in which optical expansion rate could be used. The first one is to keep expansion rate at a constant positive value, the magnitude of which depends on the point when braking is initiated [8]. Another strategy is based on the tau (τ) variable. Lee [9] demonstrated that the optical variable tau, which is the ratio of object's image size (b) to the image expansion rate ($\dot{b}$), specifies the time-to-contact (TTC) with the object as long as the current velocity is held constant:

$$\tau = \frac{b}{\dot{b}} = \frac{D}{V} = TTC \; . \tag{1}$$

Lee [9] also showed that the time derivative of τ ($\dot{\tau}$, or tau-dot) could be used to control deceleration during braking. If $\dot{\tau} < -0.5$, the current deceleration is too low and if it is maintained, it will result in a crash. If $\dot{\tau} > -0.5$, the current deceleration is too high and one will stop away from the object. If $\dot{\tau} = -0.5$, then the current deceleration will bring the observer to a stop right at the object. One can control braking by adjusting deceleration so as to keep $\dot{\tau}$ around -0.5, which is known as the "constant $\dot{\tau}$" strategy. A third hypothesis claims that braking could be controlled by computing the required deceleration from spatial variables such as distance to the object, the velocity of the observer, object size together with the optical variables [8]. Yilmaz and Warren [8] list two other strategies, in which deceleration is not continuously controlled but the brake is used in an impulsive fashion. The first strategy is the "slam on the brake" strategy in which an observer approaches the object with a constant velocity and then applies maximum deceleration. The second strategy is to apply a large deceleration at the beginning of the approach and then slowly drift to the object, using one or more deceleration spikes later to stop. Yilmaz and Warren call this strategy the "bang-bang" strategy.

The work presented in this paper explores how braking strategies of the evolved model agents are influenced by the perceptual information that is available to the agents. The simulations are based on the experiment carried out by Yilmaz and Warren [8], in which participants viewed computer displays simulating an approach to a stationary road sign. The task of the participants was to stop as close as possible to the road sign. The deceleration was regulated via a spring-loaded mouse. The only difference between our simulations and original experiment is that in our simulations the initial TTCs are longer. The reason to keep the TTC values longer is to investigate when braking is initiated.

2 Methods

In a series of experiments, model agents that are placed in a simple 2D environment with one stationary line object are evolved (see, Fig. 2(a)). The length of the object was 60. The agent has a circular body with a diameter of 30, and four sensors. The first sensor receives an input proportional to the image size (b) of the object, which is calculated using the geometry illustrated in Fig. 1. The second sensor receives an input proportional to the image expansion rate ($\dot{b}$). In the experiments reported here, we simulated translation through a rigid environment with no rotations, such as eye and head rotations. In this case, $\dot{b}$ can be calculated using Formula 1. The third and the fourth sensors detect the optical variable τ and $\dot{\tau}$, respectively. The task of the agent is to stop as close as possible to the object without hitting it. The agent can only move forward, i.e., its heading is fixed and it can only decelerate. So, it is a second-order, Newtonian system. The braking force is controlled by the motor neuron.

The behavior of each agent is controlled by a continuous-time recurrent neural network (CTRNN) with the following state equation:

$$T_i \dot{s}_i = -s_i + \sum_{j=1}^{N} w_{ji} \sigma(g_j(s_j + \theta_j)) + I_i \qquad i = 1, ..., N \ . \tag{2}$$

where N is the number of the CTRNN nodes, s is the state of each neuron, T_i is the time constant, w_{ji} is the strength of the connection from the j^{th} neuron to the i^{th} neuron, g is a gain, θ is a bias term, $\sigma(x) = 1/(1 + e^{-x})$ is the standard logistic activation function and I is the external input. The output of a neuron is $O_i = \sigma(s_i + \theta_i)$. All neurons, except for the sensory neurons, had a gain of 1.0. The agent's four sensors are fully connected to four fully interconnected interneurons which are in turn fully connected to one motor neuron controlling the vertical motion of the agent (see, Fig. 2(b)). The agent's deceleration is calculated using the following formula:

$$-\dot{V} = k \times O_m \ . \tag{3}$$

where O_m is the output of the motor neuron and k is a scaling constant which is set to be 3.0 in all of the experiments that will be reported in the next section.

The connection weights ($w_{ji} \in [-16, 16]$), biases ($\theta \in [-16, 16]$), time constants ($T \in [1, 10]$) and the gains ($g \in [1, 5]$) were evolved using a real-valued hill

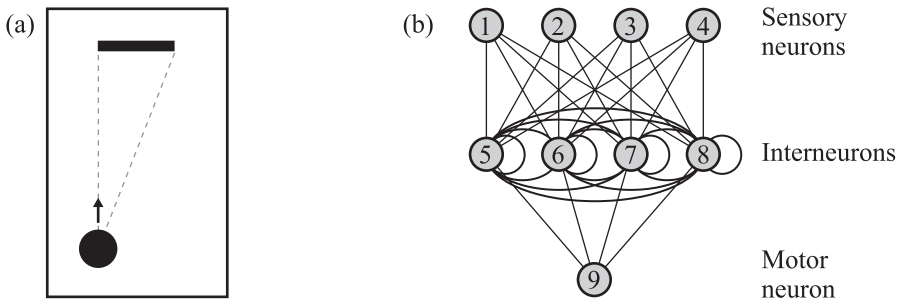

Fig. 2. (a) Basic experimental set-up (b) The CTRNN architecture.

climbing algorithm with fitness-proportionate selection. New generations were created by applying random Gaussian mutations to the selected parents. The mutation variance was 0.45. The fitness scaling multiple was 1.03. Simulations were integrated using the Euler method with an integration step size of 0.1.

An agent's performance is determined based on its behavior in a number of evaluation trials. The object's position in the environment is fixed across all trials. The horizontal position of the agent is also the same across trials but the vertical distance between the agent and the object varied. The agent has 7 different initial distances from the object (120, 135, 150, 165, 180, 205 and 210) and 6 initial velocities (10.0, 11.0, 12.0, 13.0, 14.0 and 15.0). As a result, initial TTC with the object varies between 8.0 and 21.0. Each possible combination of the agent's initial distances and velocities was presented as a trial, resulting in $7 \times 6 = 42$ evaluation trials. At the beginning of each trial, the agent's neural states are initialized to zero. Then, the agent is placed in one of the 7 locations and its velocity is initialized to one of the 6 velocities. A trial ends when the velocity of the agent is 0.0 or when the agent touches the object, i.e., when the vertical distance between the center of the agent and the object is less than or equal to the radius of the agent. The overall fitness of the agent was determined by averaging the fitness of the agent over 42 evaluation trials.

Two different fitness functions were used in two different sets of experiments. The first fitness function was based on the velocity of the agent and the vertical Euclidean distance between the agent and the object at the end of a trial. It minimizes the agent's velocity and the distance between the agent and the object. Then, the performance measure to be maximized was:

$$\frac{\sum_{i=1}^{NumTrials} (1 - d_i/dMax_i)(1 - v_i/vMax_i)}{NumTrials} . \tag{4}$$

where $NumTrials$ is the total number of trials, d_i is the vertical distance between the agent and the object at the end of i^{th} trial, $dMax_i$ is the initial vertical distance of the agent from the object, v_i is the agent's velocity at the end of the i^{th} trial and $vMax_i$ is the agent's initial velocity. The second fitness function also minimizes the trial duration in addition to the velocity and the distance. In this case, the performance measure to be maximized was:

$$\frac{\sum_{i=1}^{NumTrials}((1 - t_i/t_{max}) + (1 - d_i/dMax_i)(1 - v_i/vMax_i)/2)}{NumTrials} . \quad (5)$$

where t_i is the duration of the i^{th} trial and t_{max} is the maximum trial duration. Since it always takes some time for agents to end a trial, it is not possible to evolve agents with perfect fitness values using this second fitness measure.

3 Results

We conducted two sets of experiments. The aim of the first set of experiments is to investigate the effect of the perceptual information on the evolved braking strategies. The second set of experiments investigates the effect of the fitness measure on the evolved braking strategies. In the first set of experiments, we manipulated the perceptual information available to the agents. The fitness values of the agents were calculated using the fitness measure given in Formula 4. There were four groups of agents, each receiving a different type of information. The first group only received image size as the information. The input to the remaining three sensors was set to be zero. Similarly, the second, third and fourth groups received only the image expansion rate, tau and tau-dot as the information, respectively. From now on, agents in different groups will be referred by the information they receive such as image size agent or tau agent.

Preliminary results indicated that if the trial duration was not limited, most of the evolved agents exhibited the "bang-bang" strategy regardless of the visual information they were receiving. In other words, the agents decreased their velocities to near zero values right at the beginning of the trials and then slowly drifted to the object, giving rise to very long trial durations. In order to prevent agents from using the "bang-bang" strategy, the maximum trial duration was set to be 500 time steps in all of the experiments. For each group, 10 evolutionary runs were performed with different random seeds. For all evolutionary runs, the population size was 150 and the maximum generation number was set to 5000. Agents that can successfully solve the task were evolved in all four groups.

For the image size group, 9 out of 10 evolutionary runs produced agents that had a fitness value over 90% on the 42 evaluation trials. The best evolved image size agent across 10 runs attained a fitness value of 99.34%. It is important to note that the fitness measure does not explicitly punish the agents for non-zero velocities. This means that it is possible for the agents to have very small but non-zero velocities at the end of the trials. The velocity profile of the best evolved image size agent across 42 evaluation trials is given in Fig. 3(a). As can be seen from the figure, at the beginning of each trial the agent applies maximum deceleration and then moves with a constant velocity, the magnitude of which depends only on the agent's initial velocity. After that, it initiates its final braking. Although there are slight variations with the decreasing fitness values, these velocity profiles were essentially the same across agents. The velocity of the best agent is zero at the end of each evaluation trial and it stops very close to the target. The average final distance between the agent and the object across 42 evaluation trials was 0.94. The performance of the agent was also tested on 4641 generalization trials in which the agent's initial distance from the object was varied between 120 and 210, with an increment of 1 and the agent's initial velocity was changed from 10.0 to 15.0, with an increment of 0.1. The average performance of the best evolved image size agent was 98.76%. In 819 of the generalization trials, the agent touched the object with an average velocity of 0.42. All of these trials correspond to the trials in which the agent's initial velocity varied between 11.1 and 11.9. The average final distance of the agent from the object across generalization trials was 0.87.

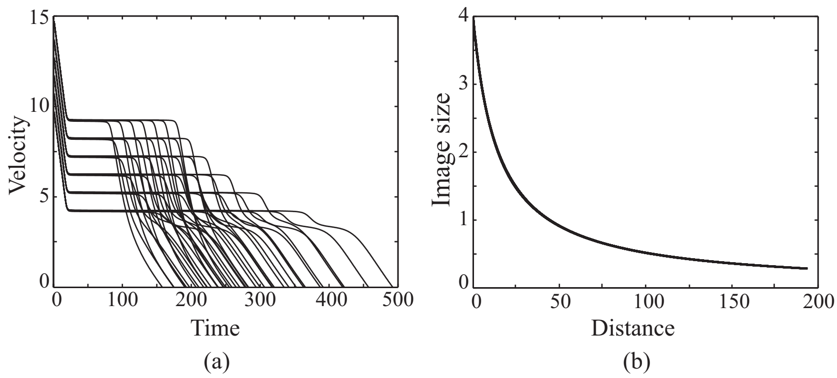

Fig. 3. (a) Velocity profiles of the best evolved image size agent across 42 evaluation trials. (b) Change in the image size as a function of distance across 42 evaluation trials.

Fig. 3(b) shows how the image size changes as the distance between the agent and the object changes across 42 evaluation trials. As can be seen from the figure, the shape of the curve is the same regardless of the initial conditions. Even though the agent did not explicitly receive image expansion rate or tau as the information, we also examined how these variables changed as the agent approaches to the object. The examination revealed that the image expansion rate was never held constant and the agent initiated its final braking when tau reached a certain value, the magnitude of which changes with the agent's initial velocity only. As the agent's initial velocity increased from 10.0 to 15.0, the tau value at which the braking was initiated decreased from 9.1 to 4.3.

For the image expansion rate group, all of the 10 evolutionary runs produced agents that had a fitness value over 96% on the 42 evaluation trials. The fitness value of the best evolved image expansion rate agent was 97.16%. The agent's velocity was always zero at the end of the evaluation trials but it stopped farther from the object compared to the image size agent. The average final distance was 4.09. The velocity profiles of the best image expansion rate agent are given in Fig. 4(a). The agent uses "slam on the brake" strategy and adjusts its braking in an impulsive fashion. Similar to the image size agent, the image expansion rate agent, too, applies maximum deceleration right at the beginning of the trials and then continues to move with a constant velocity. However, unlike the image size agent, it also applies maximum deceleration at the end of the trials. The image expansion rate was never held constant and the final braking was initiated when tau reaches a certain value, which varied between 6.6 and 3.2. The behavior of the rest of the agents was also very similar. The average performance of the best image expansion rate agent on the generalization trials was 97.01%. At the end of each generalization trial, the agent's velocity was zero. The average final distance between the agent and the object was 4.31, which is greater than the average final distance for the image size agent.

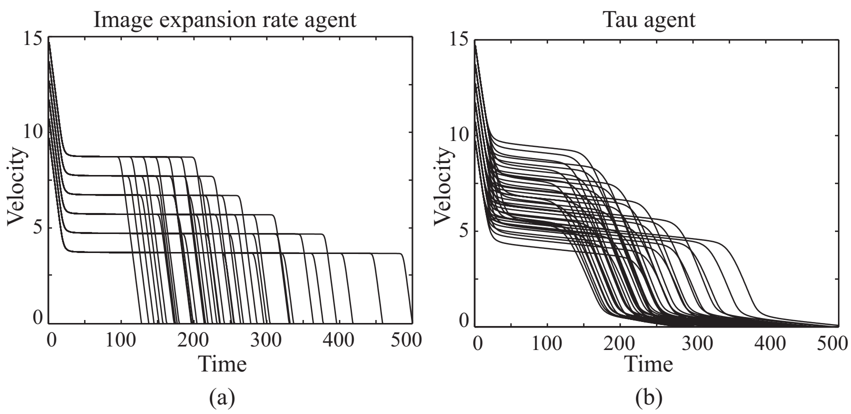

Fig. 4. Velocity profiles of the best evolved image expansion rate agent (a) and tau agent (b) across 42 evaluation trials.

For the tau group, the best evolved agent in each evolutionary run had a fitness value of at least 97.7%. The best evolved agent across 10 runs achieved a fitness value of 99.33%. This agent touched the object in 5 of the evaluation trials with an average velocity of 0.11. The average final distance between the agent and the object across all evaluation trials was 0.20. Similar to the image size agent, the tau agent also stopped very close to the object and risked touching it in some of the trials. However, its velocity profiles, which are illustrated in Fig. 4(b), are different. The agent seems to use a mixture of the "bang-bang" and the "slam on the brake" strategies. The tau agent also applies maximum deceleration at the beginning of the trials. But this time, the magnitude of the reduced velocity is not only dependent on the agent's initial velocity but also on the agent's initial distance. It varies between approximately 5.0 and 10.0. Then, the tau agent continues to decrease its velocity. After initiating its final braking, it slowly drifts to the object. The final braking was initiated when tau reached a certain value, which varied between 7.5 and 5.6. The agent's average performance on the generalization trials was 99.82%. The agent touched the object in 482 of the generalization trials with an average velocity of 0.05. The agent's average distance from the object at the end of generalization trials was 0.22. The examination of the velocity profiles of the rest of the agents revealed two more behaviors. The first behavior is similar to the behavior of the image expansion rate agent. The second behavior is more like the "slam on the brake" strategy, in which the agent moves with its initial velocity for a period of time and then rapidly increases its deceleration to the maximum value.

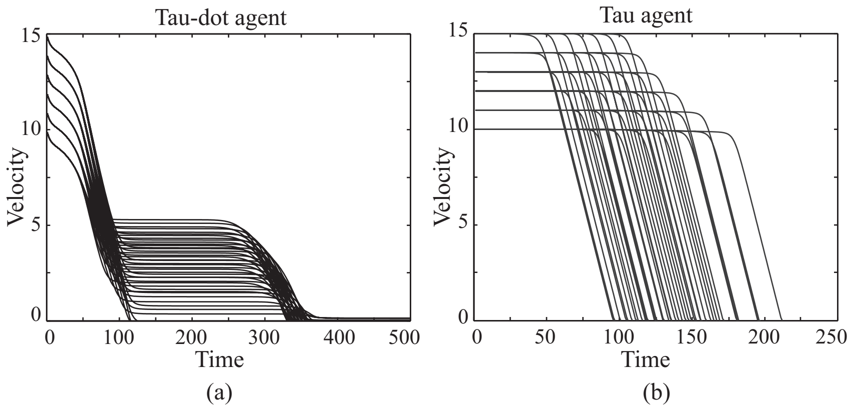

Fig. 5. (a) Velocity profile of the best evolved tau-dot agent across 42 evaluation trial. (b) Velocity profile of the best tau-agent that is evolved when the trial duration is explicitly included in the fitness measure.

Finally, for the tau-dot group, in 9 of the evolutionary runs, the best evolved agent achieved a fitness value 90% or higher. The best evolved tau-dot agent across 10 evolutionary runs had a fitness value of 97.06%. The agent touched the object in 4 of the evaluation trials with an average velocity of 0.12. The average final distance between the agent and the object was 4.2 but the variation between the final distances among trials was greater. Fig. 5(a) shows the velocity profiles of the best evolved tau-dot agent, which were very similar across agents. At the beginning of each trial, the agent decelerates at a decreasing rate. Regardless of the initial velocity or the distance, the agent applies the same deceleration. Then, the agent increases the deceleration to a maximum value, which varies with the agent's initial velocity and the distance. In three of the evaluation trials, this initial braking brings the agent to a stop very close to the object. In the remaining trials, the agent uses a kind of "bang-bang" strategy. It moves with a constant velocity, then initiates the final braking and then slowly approaches to the object. The tau value at which the final braking was initiated now varies with the agent's initial distance and the velocity but the agent never let the tau values go below 7.3. The agent's average performance on the generalization trials was 97.65%. It touched the object in 476 of the generalization trials with an average velocity of 0.09. Its final distance from the object varied between 0.01 and 18.16, with a mean of 3.36.

One common feature of the evolved agents is that all of them applied maximum deceleration at the beginning of each trial regardless of the information that they received. It is possible that at the beginning of an evolutionary search, it is easier for agents to increase their fitness by decreasing their velocity. As a result, rather than being governed by the visual information, the initial braking behavior might be due to the fitness measure. In order to test this prediction, we ran a second series of experiments using the fitness measure given in Formula 5. However, since this fitness measure puts an explicit time pressure on agents, it might encourage agents to adopt "slam on the brake" strategy. 10 evolutionary runs were performed with different random seeds, each having a population size of 150 and the maximum generation number of 5000. The agents received tau as the information. In all of the runs agents whose fitness values are at least 90% were evolved. Maximum trial duration was set to be 1000 time steps. The best evolved agent had a fitness value of 92.07% across evaluation trials and 92.09% across generalization trials. At the end of each evaluation trial, the agent's velocity was zero and the agent touched the object in 9 of the generalization trials with an average velocity of 0.46. The average final distance between the agent and the object was 2.30 across all generalization trials. The velocity profile of this agent can be seen in Fig. 5(b). As predicted, putting an explicit time pressure on the agents eliminated the maximum deceleration at the beginning of the trials. However, it also encouraged the agent to adopt the "slam on the brake" strategy. The tau values that the agent initiated the braking varied between 5.6 and 7.5. This behavior was the same across agents.

4 Discussion

In this paper, we presented results from a series of experiments in which successful simulated agents are evolved to solve a simple braking task. In the first set of experiments, there were four different groups of agents, each receiving different perceptual information: image size, image expansion rate, tau and tau-dot. In each group, the agents that can successfully solve the task were evolved. All of the best evolved agents used an impulsive braking strategy in which the deceleration was not controlled continuously. One common feature of the velocity profiles of all of the best agents was the maximum deceleration that was applied at the beginning of each trial. This is probably due to the fitness measure since it makes it easier for the agents to increase their fitness by reducing their velocity. This behavior is eliminated in the second set of experiments, when the trial duration is explicitly included in the fitness measure. However, putting a time pressure on agents encourage them to use a pure "slam on the brake" strategy, in which they approached the object by keeping their initial velocities constant and then applied maximum deceleration at the end of the trials. Yilmaz and Warren [8] indicate that the "slam on the brake" strategy cannot be an efficient strategy to control braking in actual driving because of its inertial consequences. However, in our simulations the agents do not suffer from the side effects of applying a large amount of deceleration in a short time period. One way to prevent agents from adopting this strategy could be punishing the agents for high jerk. It is also possible that the use of a sigmoid activation function for the motor neuron caused the agents to apply either full braking or no braking at all, therefore, preventing them from continuously adjusting their velocity. One possible solution could be evolving the gain of the motor neuron which was set to be 1.0 in the experiments reported in this paper.

Another common feature of the image size, tau and tau-dot agents is that they all stopped very close to the object and sometimes risked touching the object with a velocity close to zero. The image expansion rate agent was safer. Its velocity was always zero at the end of the trials. However, it stopped farther away from the object compared to the other agents. The fitness measures that we used in the experiments give equal weights to the distance and velocity components. Stopping very close to the object and touching it with a low velocity result in similar fitness values. As a result, rather than being an indication of the inefficiency of the perceptual information, the non-zero velocities at the end of the trials might simply be an artifact of the fitness measures. They might be prevented by increasing the weight of the velocity component in the fitness measure or by punishing the agent for non-zero velocities.

Our main goal for evolving these model agents is to use them as a tool for studying human control of locomotion. Yilmaz and Warren [8] provided evidence in favor of the constant tau-dot strategy. For now, none of the evolved agents seem to apply the constant tau-dot strategy or the constant image expansion rate strategy. Currently, we are investigating under what conditions those strategies evolve. We are also investigating the braking strategies used by humans in ongoing experiments involving human subjects. Another point is that although the velocity profiles of the agents change with the changing perceptual information, the perceptual information does not significantly alter the strategies adopted by the agents. However, it is possible that the effect of the changing perceptual information might be suppressed by issues related to the fitness measures that are mentioned above. Once those issues are solved we might be able to see the effect of changing perceptual information. Then, the next step is to allow agents to receive various combinations of visual information and to investigate under what conditions one type of information is preferred as opposed to others.

Acknowledgments. This work was supported in part by NSF grant IIS-0916409 to RDB.

References

1. Iizuka, H., Di Paolo, E. A.: Minimal agency detection of embodied agents. In: Almeida e Costa, F., Rocha, L.M., Costa, E., Harvey, I., Coutinho, A. (eds.) Proceedings of the 9th European Conference on Artificial Life. LNCS (LNAI), 485–494. Springer, Heidelberg (2007)
2. Rohde, M., Di Paolo, E.: Adaptation to sensory delays: An evolutionary robotics model of an empirical study. In: Almeida e Costa, F., Rocha, L., Costa, E., Harvey, I., Coutinho, A. (eds.) Proceedings of the 9th European Conference on Artificial Life. LNCS (LNAI), 193–202. Springer, Heidelberg (2007)
3. Wood, R., Di Paolo, E.: New models for old questions: Evolutionary robotics and the 'A Not B' error. In: Almeida e Costa, F., Rocha, L., Costa, E., Harvey, I., Coutinho, A. (eds.) Proceedings of the 9th European Conference on Artificial Life. LNCS (LNAI), 1141-1150. Springer, Heidelberg (2007)
4. Loomis, J. M., Beall, A. C.: Model-based control of perception/action. In: Vaina, L. M., Beardsley, S. A., Rushton, S. K. (eds.) Optic flow and beyond, 421–441. Kluwer, Dordrecht, the Netherlands (2004)
5. Gibson, J. J.: The Perception of the Visual World. Houghton Mifflin, Boston (1950)
6. Gibson, J. J.: Visually controlled locomotion and visual orientation in animals. British journal of psychology 49(3), 182–194 (1958)
7. Gibson, J. J.: The ecological approach to visual perception. Houghton Mifflin, Boston (1979)
8. Yilmaz, E. H., Warren, W. H., Jr.: Visual control of braking: A test of the hypothesis. Journal of Experimental Psychology: Human Perception and Performance 21, 996–1014 (1995)
9. Lee, D. N.: A theory of visual control of braking based on information about time-to-collision. Perception 5, 437–459 (1976)