

# An Evolutionary Robotics Model of Visually-Guided Braking: Testing Optical Variables

Didem Kadihasanoglu<sup>1</sup>, Randall D. Beer<sup>2,3</sup> and Geoffrey P. Bingham<sup>2,4</sup>

<sup>1</sup>Department of Psychology, TOBB University of Economics and Technology, Ankara, Turkey

<sup>2</sup>Cognitive Science Program, Indiana University, Bloomington, IN 47406, USA

<sup>3</sup>School of Informatics and Computing, Indiana University, Bloomington, IN 47406, USA

<sup>4</sup>Department of Psychological and Brain Sciences, Indiana University, Bloomington, IN 47405, USA

dkadi@etu.edu.tr

## Abstract

This paper presents results from a series of evolutionary robotics simulations that were designed to investigate the informational basis of visually-guided braking. Evolutionary robotics techniques were used to develop models of visually-guided braking behavior in humans to aid in resolving existing questions in the literature. Based on a well-used experimental paradigm from psychology, model agents were evolved to solve a driving-like braking task in a simple 2D environment involving one object. Agents had five sensors to detect image size of the object, image expansion rate, tau, tau-dot and proportional rate, respectively. These optical variables were those tested in experimental investigations of visually guided-braking in humans. The aim of the present work was to investigate which of these optical variables were used by the evolved agents to solve the braking task when all variables were available to control braking. Our results indicated that the agent with the highest performance used exclusively proportional rate to control braking. The agent with the lowest performance was found to be using primarily tau-dot together with image size and image expansion rate.

## Introduction

The task of braking to avoid collisions, such as approaching an object so as to grasp it or slowing down to stop before reaching an obstacle in the path of motion, is a fundamental task that we perform in our everyday life. In the context of a driving-like braking task, the aim of the present work is to investigate how braking behavior is performed by modeling using evolutionary robotics techniques. Given a driving-like braking task, humans exhibit two different braking behaviors: impulsive braking and continuously regulated braking. In impulsive braking, the brake is used discretely, most of the time in an on-or-off fashion. In continuously regulated braking, deceleration is increased or decreased continuously and the brake is never released once it is applied. In this paper, we focused on continuously regulated braking and investigated its informational basis using evolutionary robotics techniques as a modeling methodology. The theoretical framework that underlies our work is the information-based control to perception/action developed by James J. Gibson (Gibson, 1979). The central idea of information-based control is that behavior is controlled by the

task specific information that is available in the optic flow and this task-specific information is used in a control strategy to achieve the desired behavior.

When an observer moves in an environment a pattern of optical motion is created at the eye of the observer, which is called “optic flow” (Gibson, 1979). The optic flow is a rich source of information about the three-dimensional structure of the environment, the observer’s motion through and relative to the environment, as well as the motion of the other objects in the environment. Since the optic flow is created by the relative motion between an observer and environment, perception is an active process in which the observer moves around in the world to generate and pick up the information specifying the world properties with respect to herself. To achieve a given goal behavior, the observer should move so as to create and maintain a certain pattern in the optic flow. These patterns are called optical variables. Optical variables provide task-specific information and are used in control strategies that guide behavior.

There are a number of optical variables proposed in the literature that could be used to control braking continuously and different control strategies relying on different optical variables have been proposed. The image size of the object on the observer’s retina and the image expansion rate are the first two optical variables that could be used to control braking. However, given that the same image size can be created by different sized objects at different distances and the same image expansion rate can be created by approaching different sized objects with different speeds, a control strategy relying on image size or image expansion rate might not be an efficient control strategy that results in successful braking performance in every situation.

A third optical variable that could be used to control braking is tau, which is mathematically defined as the image size divided by image expansion rate (Lee, 1976). Tau specifies the time-to-contact with the object as long as the current approach velocity is held constant. Bingham (1995) suggested that braking could be controlled by using a strategy called the constant tau strategy. To stop safely at the object, this strategy requires one to move so as to keep tau constant at a certain value, the magnitude of which depends on the initial conditions and the maximum brake capacity. This strategy results in continuously regulated braking, in which the

approach velocity decreases linearly. No studies have yet found evidence supporting a pure constant tau strategy. However, our previous modeling work (Kadihasanoglu, Beer and Bingham, 2015) provided evidence that supports the use of a form of the constant tau strategy, in which discrete tau values are used to determine the initiation and termination of deceleration.

A fourth optical variable that could be used to control braking is tau-dot. Tau-dot is the first time derivative of tau and provides information as to whether the current deceleration is adequate to stop safely at the object (Lee, 1976). A tau-dot value of  $-1.0$  corresponds to a constant velocity approach. If  $\text{tau-dot} > -0.5$ , the current deceleration is too high and it should be decreased to stop at the object. If  $\text{tau-dot} < -0.5$ , then the current deceleration is too low and it should be increased in order not to crash to the object. A tau-dot value of  $-0.5$  brings the observer to a stop right at the object with a constant deceleration. The constant tau-dot strategy proposed by Lee (1976) suggests that to stop safely at an object, an observer should move so as to keep tau-dot constant at a value around  $-0.5$ .

Yilmaz and Warren (1995) tested the use of the constant tau-dot strategy with a simulated braking task. On a computer screen, participants viewed a simulated approach to a set of road signs in a 3D environment and were asked to stop as close as possible to the signs using a spring-loaded mouse as a brake. Their results showed that tau-dot was not held constant at  $-0.5$  while approaching to the signs. However, they argued that participants relied on tau-dot to control braking such that tau-dot =  $-0.5$  was used as a reference value to regulate the required change in deceleration.

A fifth optical variable was proposed by Anderson and Bingham (2010) for the control of braking. This optical variable is called proportional rate (*PR*), and mathematically it is defined as the ratio of tau and tau-dot. Anderson and Bingham (2010) suggested that to stop successfully at an object, an observer should move so as to keep *PR* constant at a certain value. In other words, the observer should try to maintain a constant proportion between tau and the rate of change of tau. This control strategy is called proportional rate control. They also provide evidence that proportional rate control is used to guide reaching (Anderson & Bingham, 2010) as well as locomoting-to-reach behaviors (Anderson & Bingham, 2011).

When compared to other control strategies, an advantage of the proportional rate control is that it offers a range of *PR* values that could be used to control braking. Choosing a different *PR* value within this range does not result in crashing; it determines whether approach is accomplished faster or slower. The *PR* value can fluctuate over the range without ill effect. Therefore, unlike the constant tau-dot strategy, which requires tau-dot to be maintained around a single value of  $-0.5$ , the proportional rate control is more flexible and more stable. It is robust to perturbations.

To sum up, there are different optical variables and control strategies proposed in the literature that could be used to control braking and different studies have provided evidence for the use of different optical variables and control strategies. Therefore, the primary aim of the present study is to test the use of these optical variables by developing models of braking behavior using evolutionary robotics simulations. The models

were developed to aid in resolving the existing questions in the literature on visually guided braking by investigating which optical variables give rise better braking performance. In a simple 2D environment with one stationary object, we evolved model agents with neural network controllers to solve a braking task. The agents had five sensors to detect, the image size, image expansion rate, tau, tau-dot and *PR*, respectively. The simulation set-up was based on a well-used experimental paradigm developed by Yilmaz and Warren (1995). We selected the agents with the best and worst braking performance to investigate the optical variables that they used to solve the braking task.

We chose the evolutionary robotics techniques as the modeling methodology since they allowed us to simulate the whole perception/action cycle. The motion of the agents in the environment created the information they received and this information in turn was used to guide their behavior. Furthermore, in evolutionary robotics, a robot's behavior is a result of a self-organization process and emerges gradually through the interactions between the robot and its environment. Robots are evaluated according to their ability to perform the desired behavior as a whole. As a result, the experimenter's assumptions about how to achieve the desired behavior are not required. Therefore, by minimizing a priori assumptions that are built into the model, these techniques also allowed us to develop models in a relatively prejudice-free fashion (Harvey, et al. 2005). Moreover, different runs of the evolutionary algorithm can produce different solutions. Usually, these solutions are different from the expectations of the experimenter and cannot be discovered by thought alone. Consequently, evolutionary robotics techniques can also be used as hypothesis generators.

The present paper extends our previous modeling work in which model agents were each evolved to solve the same braking task with access to only one of the optical variables (Kadihasanoglu, Beer & Bingham, 2015). In our previous work, we isolated optical variables to test the effectiveness of each optical variable alone. We, then, investigated what kind of control strategies best enabled agents to solve the braking task. However, in the natural environment we live in, all these optical variables are available to control braking. Therefore, to increase the model's ecological validity, in the present work, we allowed each agent access to all five optical variables. We then investigated which optical variables that the evolved agents used to solve the braking task, when all optical variables were available for controlling braking.

The rest of the paper is organized as follows. In the next section, the general simulation set-up, the neural network model and the evolutionary algorithm used to evolve the agents are described. After that, we present the results obtained from the analyses of the evolved agents. Finally, we conclude with a general discussion of the results and an outline for future work.

## Methods

The simulation set-up is essentially the same as that used in Kadihasanoglu, Beer and Bingham (2015). It is based on a simulated braking task that has been well used in experiments investigating visually guided braking. An agent, which has a circular body with a diameter of 30, is placed in a 2D

environment with one stationary line object (Figure 1(a)). The agent's retina is assumed to be 1 unit distance away from its center. The retina provides the source information for the agent's five sensors. Each of the five sensors receives an input proportional to image size of the object on its retina ( $b$ ), image expansion rate ( $\dot{b}$ ), tau ( $\tau$ ), tau-dot ( $\dot{\tau}$ ) and proportional rate ( $PR$ ), 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 along a linear path toward the object and is only able to decelerate.

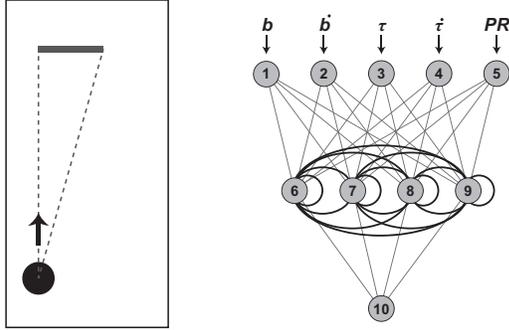


Figure 1: The basic agent–environment set-up used in the evolutionary robotics simulations. (a) The black circle represents the agent and the line represents the stationary object. (b) The architecture of the network that controls the agent's behavior.

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

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

where  $N$  is the number of nodes in the CTRNN,  $s$  is the state of each neuron,  $T$  is the time constant,  $w_{ji}$  is the strength of the connection from the  $j^{\text{th}}$  neuron to the  $i^{\text{th}}$  neuron,  $g$  is a gain term,  $\theta$  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  $o_i = \sigma(s_i + \theta_i)$ . The five sensors are fully connected to a layer of four interneurons (Figure 1(b)). The interneurons are fully interconnected and project fully to a motor neuron, controlling the deceleration of the agent. The agent's deceleration is calculated using the following formula:

$$-\dot{v} = k o_m$$

where  $o_m$  is the output of the motor neuron,  $k$  is a scaling constant that determines the maximum deceleration.  $k$  was set to 3.0. Thus, mathematically, the model is a dynamical system consisting of 11 differential equations. Ten of the equations describe how the state of each node in the CTRNN changes over time. The last equation describes the motion of

the agent in the environment. For both simplicity and speed, simulations were integrated using the Euler method with an integration step size of 0.1.

We used a population-based, real-valued hill-climbing algorithm to evolve the connection weights, time constants, and biases of all neurons in the CTRNN. Only sensor gains were evolved. Interneurons and the motor neuron had a gain of 1.0. The ranges for the neural network parameters were as follows: connection weights  $\in [-16, 16]$ , biases  $\in [-16, 16]$ , time constants  $\in [1, 10]$ , sensor gains  $\in [1, 5]$ . These ranges were selected to increase the probabilities of getting a dynamically rich parameter space and encountering these interesting dynamics (Beer, 2006). The population size was 150 and maximum generation number was set to 5000. New generations were created by applying random Gaussian mutations to the parents, with a mutation variance of 0.45. The fitness scaling multiple was set to 1.03.

The agent's performance was determined based on its behavior in a series of evaluation trials. The agent has seven different initial distances from the object (120.0, 135.0, 150.0, 165.0, 180.0, 205.0 and 210.0) and six initial velocities (10.0, 11.0, 12.0, 13.0, 14.0 and 15.0), resulting in initial time-to-contact values varying between 8.0 and 21.0. The object had four different sizes (45.0, 55.0, 65.0 and 75.0). As a result, there were  $7 \times 6 \times 4 = 168$  evaluation trials. Each trial proceeded as follows. First, the agent's neural states were initialized to zero and the agent was placed at one of the seven distances. The agent's center was always aligned with the object's left end and the position of the object's right end changed depending on the object's size. Then, the agent's velocity was set to one of the six velocities. The agent moved along a linear path with a constant velocity unless the brake was applied. A trial ended when at least one of the following was true: (1) the agent's velocity was 0.0, (2) the agent collided with the object, or (3) maximum trial duration, which was set to 500 time steps, was reached. The agent's overall fitness was determined by averaging its fitness values attained in evaluation trials.

The fitness function used to evaluate the performance of the agents had three components to minimize: (1) the distance between the agent and the object, (2) the velocity of the agent, and (3) the total jerk for a trial. Impulsive braking implicates sudden changes in deceleration resulting in high jerk. Total jerk was included in the fitness function to prevent agents from using an impulsive braking strategy. Then, the fitness measure to be maximized was:

$$\frac{\sum_{i=1}^{NumTrials} [((1 - d_i / dMax_i) + (1 - v_i / vMax_i)) / 2 - W_j \text{jerk}_i]}{NumTrials}$$

where  $NumTrials$  is the total number of trials,  $d_i$  is the longitudinal distance between the agent and the object and  $v_i$  is the agent's velocity at the end of  $i^{\text{th}}$  trial,  $dMax_i$  and  $vMax_i$  are the agent's initial distance and the velocity in  $i^{\text{th}}$  trial.  $W_j$  is the weight of the jerk term. For the weight of the jerk term, we tested different values within the range of 500 and 1200. For values less than 1000, it was still possible to observe impulsive braking behavior. Values greater than 1000 reduced

the fitness values of the evolved agents. Therefore, we chose the value 1000 as the weight of the jerk term.

We calculated the total jerk for a trial using the following formula (Flash & Hogan, 1985):

$$jerk_i = \frac{1}{2} \int_0^{t_i} \left( \frac{d^3y}{dt^3} \right)^2 dt$$

where  $y$  is the time-varying distance between the agent and the object,  $jerk_i$  is the time integral of the square of the jerk in  $i^{th}$  trial and  $t_i$  is the trial duration. When the agent’s velocity is less than or equal to 10 % of its initial velocity, jerk did not contribute to the integral to allow the agent to apply the brake and stop in front of the object. Even though the agents were evolved to minimize jerk, there was always some jerk in the system. Therefore, it was not possible for an agent to attain a 100% fitness value.

To successfully perform the braking task, first, the agent must not crash into the object, and second, it should stop very close to the object. As is the case with crashes, stopping far away from the object indicates that the visual information is used poorly, and therefore, the task is not achieved. Since it might be possible for an agent to attain a relatively high fitness value by stopping relatively farther away from the object, assessing the performance of the agents only by looking at their fitness value might be inadequate. Therefore, as in Kadihasanoglu, Beer and Bingham (2015), we defined four classes of trials to better evaluate the agents’ performance: (1) crash trials, (2) premature stops, (3) failure-to-reach trials and (4) successful trials. Crash trials are the trials in which the agent’s final velocity ( $v_{final}$ ) is greater than zero when the final distance between the agent and the object ( $d_{final}$ ) is zero. Premature stops correspond to the trials in which the agent stopped far away from the object. Any trial in which  $d_{final}$  is greater than 15.0 (the radius of the agent’s body) is a premature stop. Failure-to-reach trials are the trials, which ended because the maximum trail duration is reached. In failure-to-reach trials,  $v_{final}$  and  $d_{final}$  are both greater than zero. Crash trials, premature stops, and failure-to-reach trials constitute unsuccessful trials. Successful trials correspond to the trials in which  $d_{final}$  is less than 15.0 when the agent stopped. It is important to note that we also use the same trial classification when analyzing human performance (Kadihasanoglu, 2012).

## Results

Ten evolutionary searches were performed with different random seeds, resulting in 10 best agents. During evolution, all agents received all five optical variables. In other words, all optical variables were available to be used to control braking. The fitness values of the best-evolved agents, together with their successful trial rates, can be seen in Table 1. Numbers in the parentheses denote the percentages of crash trials, failure-to-reach trials and premature stops, respectively. As can be seen from the table, Agent 5 has the lowest fitness value (82.66%) and the lowest successful trial rate (28.6%). Agent 7 has the highest fitness value (86.78%). However, the highest successful trial rate belongs to Agent 9 (72.0%). When compared to Agent 7, Agent 9 stops closer to the object in

more trials. It has less premature stops, and therefore, has a higher successful trial rate. However, its final velocities in crash trials (mean = 3.74, max = 6.88) are higher than those of Agent 7 (mean = 2.24, max = 3.02), causing its fitness value (85.89%) to be less than that of Agent 7.

Agent No	Agents’ Performance	
	Fitness value	Successful trial rate
9	85.89%	72.0% (13.7%, 0%, 14.3%)
4	85.02%	66.1% (22.6%, 0%, 11.3%)
1	85.57%	64.3% (21.4%, 0%, 14.3%)
2	83.98%	63.7% (24.4%, 0%, 11.9%)
7	86.78%	61.9% (11.9%, 0%, 26.2%)
6	84.95%	59.5% (23.8%, 0%, 16.7%)
3	85.23%	59.5% (16.7%, 0%, 23.8%)
8	85.99%	59.5% (16.7%, 0%, 23.8%)
10	83.61%	56.6% (28.0%, 7.7%, 7.7%)
5	82.66%	28.6% (36.3%, 0%, 35.1%)

Note. Numbers in the parenthesis denote the percentages of crash trials, failure-to-reach trials and premature stops, respectively.

Table 1: Fitness values of the 10 best-evolved agents, together with their successful trial rates.

We selected the agents with the highest and lowest successful trial rates, i.e., Agent 9 and Agent 5, to investigate the optical variable(s) that they used to solve the braking task. To achieve this, we systematically removed one or more optical variables from the agent-environment system and examined how the agents’ performance changed. In other words, the agents were placed in the environment, input to one or more sensors was set to zero and the agent performed the same 168 evaluation trials in the absence of the optical variable(s) to receive a fitness value. Each possible 1-, 2-, 3- and 4-combinations of optical variables were removed and the fitness values attained in each case were then compared.

Figure 2 shows the fitness values attained by Agent 9 when only one optical variable was removed from the agent-environment system. The green line in the figure represents the agent’s original fitness value of 85.89%, the red line represents the fitness value attained by the agent (66.90%) when all optical variables were removed, i.e. when the agent performed the braking task without visual information. A fitness value of 66.90% might seem to be a relatively high fitness value, and therefore, it could be argued that Agent 9 performed the task relatively well even in the absence of any visual information. However, this is not the case. Consider an agent that stops too soon and far away from the object in all trials. Since the fitness is a continuous measure and the final velocity of this agent is zero in all trials, this agent will receive 50.0% fitness value from the velocity component of the fitness function. However, the task is not just to stop, but to stop as close as possible to the object. Therefore, unless the agent stops very close to the objects (i.e.,  $d_{final} < 15.0$ ) in majority of the trials, the agent cannot be considered as

successful in solving the braking task. The same thing is also true for an agent that crashes to the object in all trials. In this case, the final distance between the agent and the object would be zero in all trials. As a result, the agent would receive 50.0% fitness value from the distance component of the fitness function even though it did not perform the braking task successfully.

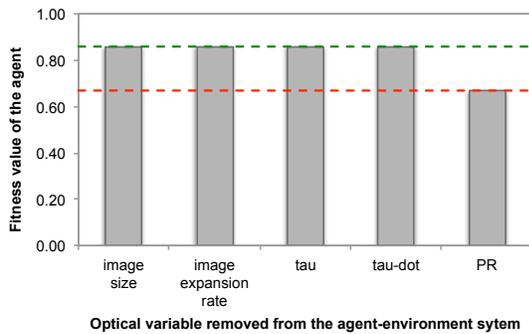


Figure 2: The fitness values attained by Agent 9 when only one optical variable is removed from the agent-environment system.

As can be seen from Figure 2, removing image size, image expansion rate, tau or tau-dot did not affect the performance of the agent at all. However, when *PR* was removed from the agent-environment system, the fitness value of the agent dropped to the no-visual-information level. In other words, the effect of removing *PR* was the same as that of removing all optical variables.

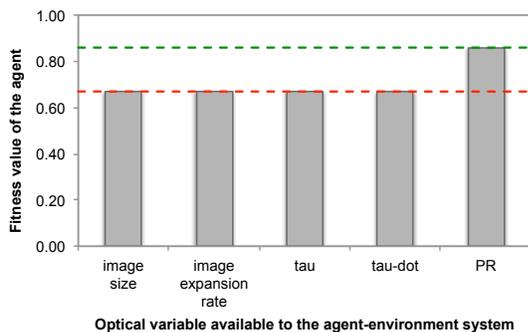


Figure 3: The fitness values attained by Agent 9 when the agent received only one optical variable.

Figure 3 shows the fitness values of Agent 9 when it received only one optical variable. Input to the remaining four sensors was set to zero. As can be seen from the figure, receiving only image size, image expansion rate, tau or tau-dot had the same effect on the agent's fitness value as receiving no visual information at all. Conversely, when the agent received only *PR*, the fitness value it had was the same

as the fitness value it attained when it received all five optical variables. In other words, receiving only *PR* had the same effect on the agent's performance as receiving all optical variables. Taken all together, these results suggest that Agent 9, i.e. the agent with the highest successful trial rate, relied exclusively on *PR* to control braking even if all five optical variables were available.

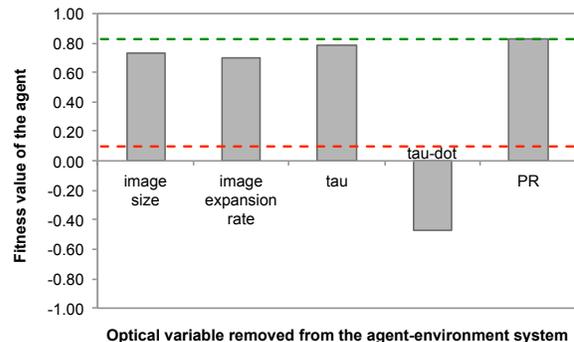


Figure 4: The fitness values attained by Agent 5 when only one optical variable is removed from the agent-environment system.

The fitness values attained by Agent 5 (the agent with the lowest successful trial rate) when one optical variable was removed from the agent-environment system are given in Figure 4. The green line in the figure again represents the agent's original fitness value (82.66%) and the red line represents the fitness value attained when all optical variables were removed (9.81%). As can be seen from the figure, removing *PR* did not affect the fitness value of the agent at all. Similarly, removing tau had a minimal effect on the fitness, suggesting that optical variables tau and *PR* were not used by the agent. When image size or image expansion rate was removed, there was a relatively small but noticeable decrease in the fitness value, suggesting that these optical variables might play some role in controlling braking. However, when tau-dot was removed, the performance of the agent was severely affected and its fitness value became negative. In other words, receiving image size, image expansion rate, tau and *PR* together in the absence of tau-dot was worse than receiving no information at all. This suggests that tau-dot is critical for the agent to perform the braking task successfully. Why did the agent's fitness value become negative when tau-dot was removed? When tau-dot is not available, the agent increases its deceleration rapidly, stops very soon and far away from the object. This behavior has two consequences: (1) the agent receives little fitness from the distance component of the fitness function and (2) rapid increase in deceleration results in very high total jerk. Since the total jerk is a subtractive term in the fitness function, very high total jerk gives rise to a negative fitness value.

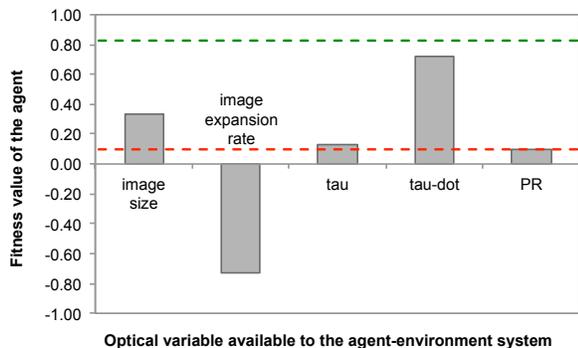


Figure 5: The fitness values attained by Agent 5 when the agent received only one optical variable.

Figure 5 shows the fitness values of Agent 5 when only one optical variable was available for controlling braking. As can be seen from the figure, receiving only tau or *PR* did not increase the fitness value of the agent from the no-visual-information level, indicating that the agent did not rely on tau and *PR* when decelerating. The fitness value of the agent when it received only image size was 33.34%, which was slightly above the no-visual-information level, suggesting that image size was used by the agent. Receiving only image expansion rate was worse than receiving no information at all and resulted in a negative fitness value. This indicates that image expansion rate by itself is not sufficient for Agent 5 to brake successfully but it plays a critical role on the observed behavior of the agent. If it had not played a role, then the performance of the agent would have been around the no-visual-information level. Finally, when the agent received only tau-dot as visual information, it attained a fitness value of 71.63%, which was quite close to its original fitness value. Taken all together, these results suggest that Agent 5, i.e. the agent with the lowest successful trial rate, used primarily, but not exclusively, tau-dot to control braking. It also relied on image expansion rate and image size. The negative fitness values observed when only tau-dot was removed from the agent-environment system and when the agent received only image expansion rate suggests that there is an interaction between these two variables in producing the observed behavior of the agent.

## Discussion

In this paper, we presented results from a set of evolutionary robotics simulations that were developed to investigate the informational basis of visually-guided braking in the context of a driving-like braking task. In a simple 2D environment with one stationary object, we evolved model agents with CTRNN controllers to slow down from high speed to stop as close as possible to the object without hitting it. Simulations were designed based on a widely used experimental paradigm from perception/action research and the different optical variables used as input to the CTRNN were those argued to be used and tested in experimental investigations of visually guided braking in humans. Mathematically, the model is a

dynamical system consisting of 11 differential equations. Ten equations describe the behavior of the CTRNN nodes and one equation defines the agent's behavior in the environment.

The present paper extends our previous work (Kadihasanoglu, Beer & Bingham, 2015) in which model agents were evolved to solve the same braking task with only one optical variable. By isolating each optical variable, our aim was to test the effectiveness of the optical variables and to investigate what kind of control strategies agents used to solve the braking task when they only had access to one of the optical variables. Our results indicated that the optical variables tau and *PR* are the most efficient optical variables and they each individually are sufficient to control braking continuously. Tau-dot alone did not result in successful braking performance.

Even though braking behavior is predominantly guided by visual information, other sensory information such as the information coming from the vestibular organ might also affect the braking behavior. However, the aim of the present work was not to develop a full model of the braking behavior but to test the effectiveness of different optical variables proposed in the literature in solving the braking task. Moreover, the models were based on a well-used experimental paradigm, which involves a simulated braking task providing only visual information. During evolution, we let the agents detect all five optical variables simultaneously. We then chose the agents with the highest and the lowest successful trial rates and investigated the optical variables used by these agents to solve the braking task.

To achieve this, we systematically removed the optical variables from the agent-environment system and compared the changes in the fitness value of the agents in the absence of these optical variables. The idea is as follows. If an optical variable is actually used by the agent to control braking, then removing that variable, i.e. setting the input to the corresponding sensor to zero, should decrease the agent's performance. Conversely, if removing an optical variable does not affect the agent's performance, then the agent does not use that particular optical variable to control braking even if it was available during evolution. It is possible that setting the input to a sensor to zero might cause too much disruption for an agent. A potentially less disruptive method for removing an optical variable might be taking the average value of a sensor over a normal run and then clamping it to that mean value instead of zero.

Consistent with our previous findings, the agent with the highest successful trial rate was found to be using *PR* exclusively to regulate its deceleration when approaching to the object, even if all five optical variables were present. An advantage of *PR* is that there is a range of *PR* values that could be used to control braking. Unlike tau-dot, which has to be maintained around a single value of  $-0.5$ , *PR* values can fluctuate over this range and slight changes in *PR* do not result in crashing. Therefore, *PR* provides more flexibility and more stability. The agent with the lowest successful trial rate was found to be relying primarily on tau-dot to control braking. This agent also used image expansion rate and image size together with tau-dot. In line with our previous work, relying on tau-dot did not result in successful braking performance, even when tau-dot was used together with image size and image expansion rate. This result is to be expected because

tau-dot fails to provide information about when to initiate braking to ensure successful performance.

What would happen if we removed the optical variables from the onset of the evolution? For example, what would be the performance of the agents if they did not receive *PR* at all? Would they fail to perform the task? Or, would they adapt via evolution and have a comparable performance to the agents receiving different combination of variables (or even all optical variables)? Both cases are possible. The results of our previous work (Kadihasanoglu, Beer & Bingham, 2015) indicated that the optical variables tau and *PR* are alone sufficient to control braking continuously. This means that an agent receiving tau as the visual information can, in principle, perform the task successfully. However, receiving an optical variable does not guarantee that the agent will use that optical variable. To illustrate, the ten agents evolved in this study received all five optical variables during evolution. Some agents controlled braking successfully; some of them performed the braking task poorly. The best-performing agent was found to be using exclusively *PR*. The agent with the lowest successful trial rate was using primarily tau-dot even if *PR* was available. In addition, CTRNNs are proven to be universal approximators of smooth dynamics (Funahashi & Nakamura, 1993). They can differentiate, integrate, and combine multiple inputs. So, an agent receiving only image size can, in theory, derive all remaining four variables and use them to control braking. However, whether a four-node CTRNN used in this study can, in practice, do all that approximation is an empirical question. Our previous results suggest that it probably can't.

Investigating how visually-guided braking is achieved requires investigating both the optical variables and the control strategies used to perform it. We found that the most successful agent used *PR* to control braking. However, this is only the first step in the analysis. The next step is to investigate the control strategy used by the agent to solve the braking task. In other words, the next step is to investigate how *PR* is used to control braking. After discovering the control strategy used by the agent, the underlying dynamics of the agent-environment system can be analyzed using tools from dynamical systems theory to understand the mechanism that produces the observed behavior of the agent.

We also plan to evolve more agents to get a group of agents with a successful trial rate greater than 70%. The aim is to investigate whether it is possible to observe a tendency in these well-performing agents to rely on certain optical variables even if all the variables are available. If, for example, the majority of these well-performing agents are found to be using *PR*, this provides more support to the hypothesis that *PR* is the optical variable that is used to guide braking behavior. Analyses of these well-performing agents might also reveal combinations of optical variables that could effectively be used to control braking.

Another way to get agents with higher successful trial rates might be modifying the fitness function so that it emphasizes the importance of not crashing. This could be done by increasing the weight of the velocity term in the fitness function. However, unsuccessful trials also involve premature stops, in which the agent stopped far away from the object. Accentuating the importance of not crashing might result in more premature stops, and therefore, might not increase the

successful trial rate significantly. An alternative way might be to include the successful trial rate in the fitness function.

The focus of the present work was to test the use of the optical variables that are proposed and tested in experimental investigations of visually guided braking in humans. This simulation set-up can also be used to investigate the constraints on human vision by evolving agents in the context of limitations and failure modes. To illustrate, one assumption made in the work presented in this paper is the absence of biological noise in extracting the optical variables from the optic flow. Introducing noise to the sensors might lead to changes in the observed behavior of the agents. Another interesting question to pursue is what happens when an optical variable is removed during evolution. How does the behavior of the agents change when they are faced with such a failure mode? Providing answers to these questions can provide valuable insights into the control of braking in humans.

## References

- Anderson, J., & Bingham, G. P. (2010). A solution to the online guidance problem for targeted reaches: Proportional rate control using relative disparity. *Experimental Brain Research*, 205(3), 291–306.
- Anderson, J., & Bingham, G. P. (2011). Locomoting-to-reach: Information variables and control strategies for nested actions. *Experimental Brain Research*, 214(4), 631–644.
- Beer, R. D. (2006). Parameter space structure of continuous-time recurrent neural networks. *Neural Computation*, 18(12), 3009–3051.
- Bingham, G. P. (1995). The role of perception in timing: Feed-back control in motor programming and task dynamics. In E. Covey, H. Hawkins, T. McMullen, & R. Port (Eds.) *Neural representation of temporal patterns* (pp. 129–157). New York: Plenum Press.
- Flash, T., & Hogan, N. (1985). The coordination of arm movements: an experimentally confirmed mathematical model. *Journal of Neuroscience*, 5, 1688–1703.
- Funahashi, K. & Nakamura, Y. (1993). Approximation of dynamical systems by continuous time recurrent neural networks. *Neural Networks*, 6, 801–806.
- Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston, MA: Houghton Mifflin.
- Harvey, I., Di Paolo, E., Wood, R., Quinn, M., & Tuci, E. (2005). Evolutionary robotics: A new scientific tool for studying cognition. *Artificial life*, 11(1-2), 79-98.
- Kadihasanoglu, D. (2012). *An evolutionary robotics approach to visually-guided braking: Data and theory*, Ph.D. Thesis, Cognitive Science Program, Indiana University, Bloomington.
- Kadihasanoglu, D., Beer, R.D., & Bingham, G.P. (2015). Evolutionary robotics techniques used to model information and control of visually-guided braking. *Adaptive Behavior*, 23(3), 125–142.
- Lee, D. N. (1976). A theory of visual control of braking based on information about time- to-collision. *Perception*, 5, 437–459.
- Yilmaz, E. H., & Warren, W. H. (1995). Visual control of braking: A test of the tau-dot hypothesis. *Journal of Experimental Psychology: Human Perception and Performance*, 21, 996–1014.