Page Proof Instructions and Queries

Journal Title: ADB
Article Number: 576392

Greetings, and thank you for publishing with SAGE. We have prepared this page proof for your review. Please review the entire proof for accuracy and provide necessary corrections by digitally marking this PDF using Adobe Reader.

Click "Comment" in the upper right corner of Adobe Reader to access the mark-up tools as follows:

For textual edits, please use the "Annotations" tools. Please refrain from using the two tools crossed out below, as data loss can occur when using these tools.

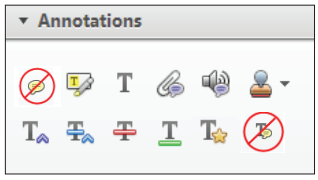

For formatting requests, questions, or other complicated changes, please insert a comment using "Drawing Markups."

Detailed annotation guidelines can be viewed at: http://www.sagepub.com/repository/binaries/pdfs/AnnotationGuidelines.pdf
Adobe Reader can be downloaded (free) at: http://www.adobe.com/products/reader.html.

Sl. No.	Query
1	No queries

Adaptive Behavior

Adaptive Behavior
1–18
© The Author(s) 2015
Reprints and permissions:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/1059712315576392
adb.sagepub.com
SAGE

Evolutionary robotics techniques used to model information and control of visually guided braking

Didem Kadıhasanoğlu[1], Randall D Beer[2,3] and Geoffrey P Bingham[2,4]

Abstract

This paper utilizes evolutionary robotics techniques as a hypothesis generator to explore optical variables and control strategies that could be used to solve a driving-like braking task. Given such a task, humans exhibit two different braking behaviors: continuously regulated braking and impulsive braking. Based on an oft-used experimental task in human perception/action research, a series of evolutionary robotics simulations were developed to explore the space of possible braking strategies by examining how braking strategies change as the optical information is manipulated. Our results can be summarized as follows: (1) behaviors similar to human behavior were observed only when the constraints were selected correctly; (2) the optical variables τ and proportional rate yielded significantly better braking performance; (3) two different classes of impulsive braking behaviors were observed, including one not reported in previous studies: discrete impulsive braking and oscillatory impulsive braking; (4) the optical variable τ is used to initiate and terminate braking; (5) the evolved model agents use a proportional rate control strategy to regulate braking continuously. We argue that combining psychological experiments and evolutionary robotics simulations is a promising research methodology that is useful for testing existing hypotheses and generating new ones.

Keywords

Evolutionary robotics, visually guided braking, proportional rate control, constant tau-dot strategy

1 Introduction

Inspired by Darwinian evolution, evolutionary robotics is a relatively new technique for designing autonomous robots and their control systems using population-based artificial evolution (Holland, 1975; Nolfi & Floreano, 2000). Recently, there is a growing interest in applying the evolutionary robotics techniques to model human cognition and behavior. These techniques have been used to develop models of associative learning (Phattanasri, Chiel & Beer, 2007; Izquierdo, Harvey & Beer, 2008), selective attention (Ward & Ward, 2008), categorical perception (Beer, 2003) and relational categorization (Williams, Beer & Gasser, 2008), agency detection (Iizuka & Di Paolo, 2007), and the famous "A-not-B" error (Piaget, 1954) exhibited by 7–12-month-old infants (Wood & Di Paolo, 2007). The main motivations for using the evolutionary robotics techniques in the study of human cognition and behavior are: (1) to enhance our theoretical thinking and knowledge; (2) to provide existence proofs and minimal conditions of a cognitive phenomenon; and (3) to test existing hypotheses about a cognitive phenomenon and to form novel ones.

What makes the evolutionary robotics techniques an attractive modeling tool for cognitive scientists and psychologists is the methodological advantages that they offer. In evolutionary robotics, better solutions to a problem are evolved from an initial population of candidate solutions. Generally, the problem to be solved is to evolve a robot that exhibits a behavior of interest and the solutions are the parameters of an artificial neural network controlling the robot to generate the specified behavior. The initial population of candidate solutions is created randomly to explore the entire search space that contains all possible solutions. Also,

[1]Department of Psychology, TOBB University of Economics and Technology, Ankara, Turkey
[2]Cognitive Science Program, Indiana University, Bloomington, IN, USA
[3]School of Informatics and Computing, Indiana University, Bloomington, IN, USA
[4]Department of Psychological and Brain Sciences, Indiana University, Bloomington, IN, USA

Corresponding author:
Didem Kadıhasanoğlu, Department of Psychology, TOBB University of Economics and Technology, Sogutozu Caddesi No: 43, Sogutozu, Ankara 06560, Turkey
Email: dkadi@etu.edu.tr

different runs of the evolutionary algorithm can produce different solutions, most of which cannot be discovered by thought alone. Usually, these solutions are different from the expectations of the experimenter. As a result, evolutionary robotics techniques enable us to test the existing hypotheses about an observed behavior as well as to form novel hypotheses and ideas that can be tested in further empirical studies. The analysis of the evolved models can also provide insights into the mechanisms underlying the observed behavior (e.g., Izquierdo & Beer's (2013) work on the klinotaxis behavior of the nematode worm *Caenorhabditis elegans*).

Continuous-time recurrent neural networks (CTRNN) are one of the most commonly used controllers in evolutionary robotics due to their dynamical properties. It has been proven that CTRNNs are universal approximators of smooth dynamics, which means that given a dynamical system, there exists a CTRNN that can approximate its behavior for any finite interval of time (Funahashi & Nakamura, 1993; Kimura & Nakano, 1998). It is also well known that even small CTRNNs can exhibit a wide variety of complex dynamics (Beer, 1995), which makes them well suited: (1) for evolving adaptive, dynamically complex and non-reactive behaviors; and consequently, (2) for exploring a wide range of possible solutions to a given problem.

Considering the sensory and motor capabilities of evolved model agents, ecological psychology in general, and visual control of locomotion in particular, provides an excellent research area in which the evolutionary robotics techniques can be used. Visual control of braking, which is an essential part of locomotion, arises in many situations such as while driving or walking in a crowded street. It can be described as approaching an object to make contact without collision. The work presented in this paper investigates the visual control of braking from an ecological psychology approach to perception and action developed by James J. Gibson (1979). The central idea of the ecological approach is that behavior is controlled by the task specific information that is available in *the optic flow* and it emerges as control laws relate information variables to action parameters (Warren, 2006).

In this paper, we focused on visually guided braking in the context of a driving-like braking task, i.e., the task of slowing down from a high speed to stop safely before reaching an obstacle in the path of motion. Given a driving-like braking task, humans exhibit two different classes of braking behaviors: (1) continuously regulated braking and (2) impulsive braking in which deceleration is used discretely. For continuously regulated braking, different control strategies that are based on different information variables have been proposed in the literature and different studies have provided evidence for different strategies. No studies have yet investigated the information variable(s) used in impulsive braking and how braking is controlled on the basis of this information.

This paper utilized the evolutionary robotics techniques as a hypothesis generator to explore the space of possible control strategies that could be used to solve a driving-like braking task. We designed a series of evolutionary robotics simulations to aid in resolving the existing questions as to the control strategies underlying continuously regulated braking and to investigate the information variable(s) and the control strategy(ies) that could be used in impulsive braking. Model agents with CTRNN controllers were evolved to solve a braking task in a simple 2D environment containing one stationary object. We manipulated the visual information that the agents received from the environment and investigated how the evolved braking strategies changed depending on the visual information available to the agents.

We would like to reiterate that our goal was not to design a controller for the braking task but to use the evolution of CTRNNs to explore the different strategies that an agent receiving a particular type of optical information can use to solve the braking task. Our choice of CTRNNs as controllers was based on their dynamical properties. As they are capable of producing dynamically complex and non-reactive behaviors, they allow us to simulate the dynamics of perception and action (Warren, 2006) and to explore a larger space of possible braking strategies. The significance of our simulations lies in two main aspects: (1) the simulation set-up was based on an experimental task from perception/action research and (2) the different information variables that the agents received from the environment were those hypothesized to be used and tested in experimental investigations of visually guided braking in humans.

We performed in-depth behavioral analyses of the evolved agents to uncover the information variables and the control strategies that could be effectively used to perform the braking task. The results we obtained showed that some information variables did not work or only worked in limited conditions. Some of the evolved agents used strategies that are proposed in the literature, some agents relied on novel strategies that are not reported in previous studies. Some of these strategies resulted in successful braking performance only in the conditions in which they were evolved without generalizing, but others did generalize.

The rest of this paper is organized as follows. As the underlying framework is Gibson's information-based approach to perception and action, the next section summarizes Gibson's main ideas and the principles of his theory. Then, we discuss the information variables and control strategies that have been proposed in the literature, together with the corresponding evidence. In the third section, the general simulation set-up is described. The fourth section is devoted to in-depth

discussion of the results obtained from the analyses of the agents, together with the implications for visual control of braking in humans. Finally, we conclude with a general discussion of the results and an outline for future work.

2 An ecological approach to visually guided braking

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). Optic flow provides 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. It is a rich source of information that is both spatially and temporally continuous.

Observers do not passively receive the information from the optic flow. Perception is an active process in which an observer moves around in the world to generate and pick up the information specifying the world properties with respect to the observer. This formulation has the following consequences. (1) Perception and action are inseparable, continuous, and cyclic. We act in order to perceive and what we perceive in turn guides our actions. (2) Perception does not occur inside the observer's head. Instead, it occurs as the observer moves in the environment, indicating that the observer and his/her environment form an inseparable pair. Thus, at the core of Gibson's theory is the rejection of any kind of dualism between an animal and its environment, between perception and action.

The optic flow is created by the relative motion between an observer and his/her environment and there is a geometry underlying the optic flow, which reflects the structure of the surrounds. This means that the observer has a spatio-temporal relation to the surrounds. Then, the idea is that to achieve a given goal behavior, the observer should move to create and maintain a certain pattern in the optic flow. These patterns are called *optical variables*.

In the case of a driving-like braking task investigated in this paper, there are a number of optical variables proposed in the literature that could be used to solve the task and different control strategies that are based on different optical variables have been suggested. Consider the case of approaching a stationary object shown in Figure 1. As the observer approaches the object, the optical structure available at the observer's eye changes. More specifically, the image size, b, of the object expands. The image size (b) and the image expansion rate ($\dot{b}$) are the first two optical variables that are available for the control of braking. However, the nature of these optical variables undermines their usefulness. The same image size can be created by different

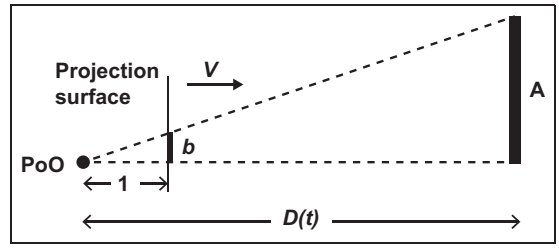

Figure 1. A schematic view of an observer approaching a stationary object of size *A* at a distance *D*, with a velocity *V*. PoO denotes the point of observation. *b* is the image size of the object on the observer's retina at time *t*. The observer's retina (i.e., the projection surface) is approximated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

sized objects at different distances. Similarly, the same image expansion rate can be produced by approaching different sized objects at different speeds. Given the complex nature of the environment we live in, relying solely on b or $\dot{b}$ might not be an efficient braking strategy.

The third optical variable is τ (tau), which can be described as the ratio of image size to image expansion rate. During a direct approach with constant velocity, τ specifies the time-to-contact (T_c) with the object (Lee, 1976):

$$\tau = \frac{b}{\dot{b}} = -\frac{D}{V} = T_c \tag{1}$$

In other words, τ specifies the time remaining before the observer collides with the object as long as the current velocity is held constant.

Bingham (1995) suggested that braking could be controlled by τ using a strategy called *the constant τ strategy*. There are two versions of this strategy. The strong version suggests that to stop at an object without colliding, an observer should move as to keep τ constant at a certain value, the magnitude of which depends on the initial conditions and the observer's braking capabilities. Even though this is an efficient braking strategy that results in a linear decrease in velocity and shorter stopping times, no studies have found evidence supporting the strategy. The weak version of the constant τ strategy suggests that τ is used to determine the initiation and termination of deceleration. It can be described as "never let τ go below a certain critical value." Unlike its strong version, which results in continuously regulated braking, the weak version results in impulsive braking, in which brake is used discretely. No studies have yet tested the use of the weak version of the constant τ strategy.

Yilmaz and Warren (1995) list two other strategies, in which deceleration is not controlled continuously but the brake is used in an impulsive fashion. In *the slam-on-the-brake strategy*, an observer approaches the

object with a constant velocity and then applies maximum deceleration later in the approach to stop at the target. *The bang-bang strategy* corresponds to applying a large deceleration at the beginning of the approach and then slowly drifting to the object, using one or more deceleration spikes later to stop.

Lee (1976) also showed that the first time derivative of τ ($\dot{\tau}$ or tau-dot) can be used to control deceleration during braking, as it provides information about whether the current deceleration is adequate to stop at the object without collision. A $\dot{\tau}$ value of -1.0 corresponds to a constant velocity approach. If $\dot{\tau} > -0.5$, the current deceleration is too high and if the deceleration is maintained, the observer will stop short of the object. If $\dot{\tau} < -0.5$, the current deceleration is too low and if it is maintained, it will result in crash. A $\dot{\tau}$ value of -0.5 provides a special case, which brings the observer to a stop right at the object with a constant deceleration. So, $\dot{\tau}$ is the fourth optical variable that can be used to control braking and *the constant $\dot{\tau}$ strategy* suggests that to stop successfully at a target, an observer should move to maintain $\dot{\tau}$ constant at a value around -0.5 (Lee, 1976).

Yilmaz and Warren (1995) tested the use of the constant $\dot{\tau}$ strategy in an experiment, which involved a simulated braking task. In this task, participants viewed a simulated approach to a set of road signs on a computer screen and were asked to stop as close as possible to the signs without colliding with them. They regulated their deceleration using a brake composed of a spring-loaded mouse. For each trial, the slope of the regression line fitted to the time series of τ was taken to be the mean $\dot{\tau}$ value for that trial.

The results showed that during approach, $\dot{\tau}$ was not held constant at -0.5 but $\dot{\tau}$ trajectories oscillated around -0.5 with a mean $\dot{\tau}$ value around -0.51. Yilmaz and Warren (1995) argued that once braking is initiated, $\dot{\tau} = -0.5$ was used as a reference value in a closed-loop control to regulate deceleration. More specifically, $\dot{\tau}$ was used to control the required change in deceleration and the adjustments of deceleration were proportional to the difference between the current value of $\dot{\tau}$ and $\dot{\tau} = -0.5$.

Even though Yilmaz and Warren's (1995) results suggest that the participants used $\dot{\tau}$ to control braking, there are some issues that need to be considered. First of all, some of their participants did not exhibit continuous regulation of the brake and used an impulsive braking strategy. Impulsive braking is also exhibited in driving especially by novice drivers. However, no studies have yet investigated the optical variable(s) and the control strategy(ies) that could be used in impulsive braking. Another issue with any $\dot{\tau}$-based strategy is that $\dot{\tau}$ does not specify when to initiate braking. $\dot{\tau}$ remains at -1.0 unless braking is initiated. This implies that if $\dot{\tau}$ is used to control braking, humans must be relying on another optical variable to initiate braking. The question of how and when people initiate braking still remains open.

Anderson and Bingham (2010) proposed an alternative optical variable and a new control strategy that could be used to control braking during approach. They suggested that to stop successfully at a target, an observer should move to maintain a constant proportion between the rate of change of τ (i.e., $\dot{\tau}$) and τ itself. This is called *proportional rate control*. They also found evidence supporting that proportional rate control is used in visually guided reaching (Anderson & Bingham, 2010). Thus, proportional rate ($\dot{\tau}/\tau$) is the fifth optical variable that is available for the control of braking.

Proportional rate control offers a number of advantages over any $\dot{\tau}$-based strategy. Firstly, proportional rate control allows a range of proportional rate (PR) values that will result in successful braking. Choosing different PR values will determine whether braking occurs faster or slower. This means proportional rate control does not result in crashing when a slightly different PR value is chosen, which makes it more resistant to perturbations, and therefore more stable, than a $\dot{\tau}$-based strategy, which would require $\dot{\tau}$ to be maintained around a single value of -0.5. Secondly, unlike $\dot{\tau}$, PR evolves over time and can be used to initiate braking without requiring any other optical variable.

In a subsequent study, Anderson and Bingham (2011) investigated the visual information and the control strategy used to guide locomoting-to-reach behavior (such as walking to open a door). They asked participants to jog to a target location and bring their nose to a stop at the target. Linear regression analysis on the time series of τ revealed that the mean $\dot{\tau}$ values were around -0.5, suggesting the use of $\dot{\tau}$ to control braking.

Anderson and Bingham (2011) also performed a split-half analysis on the τ and PR trajectories. Using the median time sample in the trajectories, they split the τ and PR trajectories in two halves. For the PR trajectories, the mean PR values were calculated for each half. For the τ trajectories, a regression line was fitted to each half to calculate the mean $\dot{\tau}$ values. The mean $\dot{\tau}$ values in the first and second halves were then compared to examine whether $\dot{\tau}$ values changed as participants approached the target. The analysis revealed a significant difference between the mean $\dot{\tau}$ values in two halves, indicating that $\dot{\tau}$ was not held constant. Moreover, Anderson and Bingham (2010, 2011) found that the slope of the first half was always greater than that in the second half, which cannot be predicted by the argument that $\dot{\tau}$ is used as a reference value to control deceleration. Conversely, when the mean PR values computed for the first and second halves were compared, no significant difference between the mean PR values was found. The PR was maintained constant at around -0.2, indicating that participants controlled their approach using a proportional rate control. These

results introduce the possibility that proportional rate control can also be used to control deceleration during a driving-like braking task used by Yilmaz and Warren (1995).

In light of the issues discussed above, we designed a series of evolutionary robotics simulations. Our aims were threefold: (1) to investigate the optical variable(s) and the control strategy(ies) underlying impulsive braking; (2) to aid in resolving the existing questions in the literature on visually guided braking by investigating which control strategy underlies continuously regulated braking; and (3) to explore the space of possible braking strategies by examining how braking strategies change as the available information is manipulated.

We chose to use the evolutionary robotics techniques as a modeling methodology because in addition to the methodological advantages discussed above, evolving these simple brain–body–environment systems offers further advantages. Firstly, the underlying assumption of this effort is that the behavior is the property of the entire brain–body–environment system and cannot be attributed to any subsystem in isolation (Beer, 2014). These ideas are in line with the ecological approach to perception and action, which emphasizes the inseparability of animal and environment and perception and action. Secondly, the evolutionary robotics techniques allow us to simulate the whole perception–action cycle. The movements of the agents in the environment create the information that they receive and the information they receive in turn guides their behavior. Finally, these evolved agents are simple enough that we can analyze them in detail to discover the strategies used to solve the braking task.

3 The methods

3.1 The agent–environment set-up

A simulated braking task has been used in experiments investigating the control strategies underlying visually guided braking. Participants viewing a 3D display find themselves speeding along a linear path over a simulated ground surface toward a set of signs. They are asked to use a brake to regulate their deceleration to stop in front of the signs. The brake is a hand held joystick or mouse, attached to the computer. This task and set-up was used in our experiments involving human subjects in Perception/Action Laboratory at Indiana University, Bloomington. Accordingly, the evolutionary robotics simulations reported in this paper were based on this well used experimental task.

The central idea of the ecological approach is that behavior is controlled by task specific information that is available in the optic flow and it emerges as control laws relate information variables to action parameters. Therefore, we employed five sensors, each detecting an optical variable proposed in the literature, and then, used evolutionary search to let the agents find their own control laws. As the information in the optic flow is created by the relative motion between an agent and its environment, we placed these sensors on a simple body that is capable of both moving along one dimension and braking.

We evolved model agents that were placed in a simple 2D environment with one stationary line object to solve a simple braking task. The simulation set-up can be seen in Figures 2a and 2b. The agent has a circular body with a diameter of 30 and its "retina" is assumed to be 1 unit distance away from its center. The first sensor receives an input proportional to the image size (b) of the object on the retina. The second sensor receives an input proportional to the image expansion rate ($\dot{b}$). The third, fourth and the fifth sensors detect τ, $\dot{\tau}$ and 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, i.e., its heading is fixed and it can only decelerate.

3.2 The neural model and the evolutionary algorithm

The behavior of the agents is controlled by a CTRNN with the following state equation:

$$T_i s_i = -s_i + \sum_{j=1}^{N} w_{ji}\sigma\big(g_j(s_j + \theta_j)\big) + I_i \quad i = 1, \cdots, N \tag{2}$$

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 jth neuron to the ith neuron, g is a gain term, θ is a bias term, $\sigma(x) = \frac{1}{(1+e^{-x})}$ 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 sensors, have a gain of 1.0. The agent's five sensors are fully connected to four fully interconnected interneurons, which are in turn fully connected to one motor neuron controlling the deceleration of the agent. The network architecture is given in Figure 2c. 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 was set to 3.0 in all of the simulations reported below.

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 population-based, real-valued hill-climbing algorithm with fitness-proportionate selection. New generations were created by applying random Gaussian mutations to the selected parents. The mutation variance was set to 0.45. The fitness scaling multiple was 1.03. Simulations were

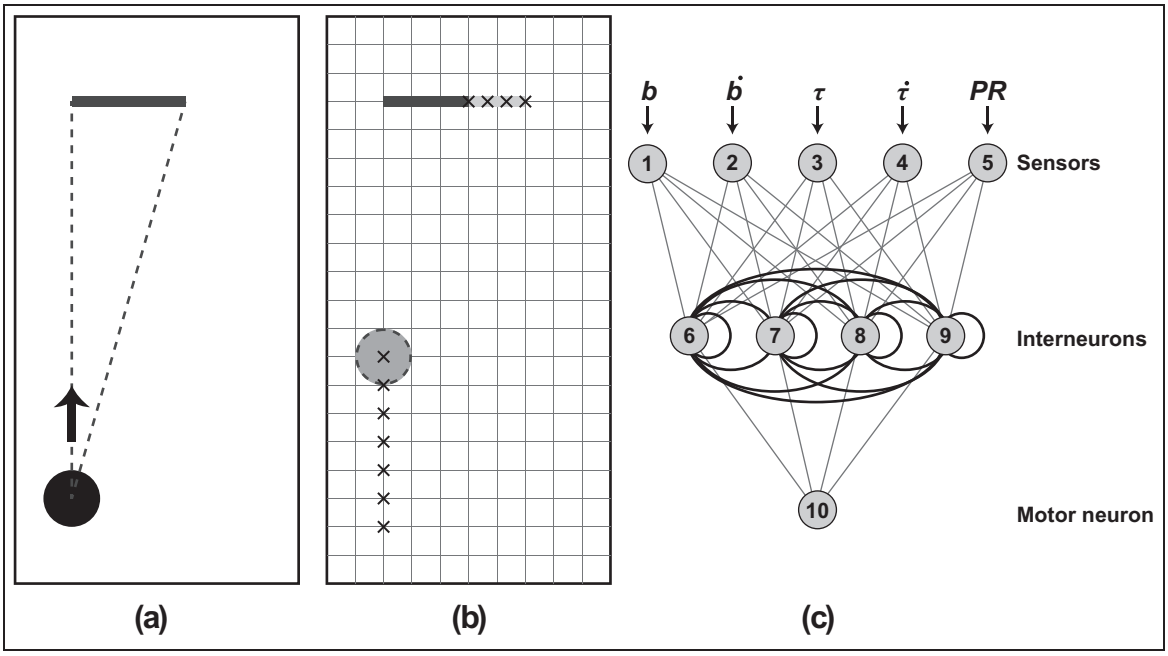

Figure 2. The basic agent–environment set-up used in the evolutionary robotics simulations. (a) The black circle represents the agent and the line represents the stationary object. The agent has five sensors to detect b, $\dot{b}$, τ, $\dot{\tau}$, and *PR*, respectively. (b) An illustration of four object sizes (45.0, 55.0, 65.0, and 75.0) and seven initial distances from the object (120.0, 135.0, 150.0, 165.0, 180.0, 205.0, and 210.0). The left end of the object is fixed in the environment. Black crosses denote the object's right end depending on its size; grey crosses denote the position of the center of the agent's body at each initial distance. Each square represents a 15 unit × 15 unit area. (c) The architecture of the network that controls the agent's behavior.

integrated using the Euler method with an integration step size of 0.1. For all evolutionary searches reported below, the population size was 150 and the maximum generation number was 5000.

3.3 Fitness evaluation

An agent's performance was determined based on its behavior in a number of evaluation trials. The object had four different sizes (45.0, 55.0, 65.0, and 75.0). The position of the object's left end was fixed but the position of its right end changed depending on the object's size. The horizontal position of the agent was also fixed but the vertical distance between the agent and the object varied. The agent had 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). As a result, initial time-to-contacts with the object varied between 8.0 and 21.0. The object sizes and the initial distances were selected so that the same image size could be created by each of the four objects at one of the initial distances.

Each possible combination of the object size, agent's initial distance, and velocity was presented as a trial, resulting in $4 \times 7 \times 6 = 168$ evaluation trials. At the beginning of each trial, the agent's neural states were initialized to zero. Then, the agent was placed in one of the seven locations and its velocity was initialized to one of the six velocities. The agent moved with a constant velocity unless braking was initiated. A trial ended when at least one of the following conditions was met: (1) the maximum trial duration was reached, (2) the agent's velocity was 0.0 or (3) the agent collided with the object, i.e., the vertical distance between the object and the center of the agent was less than or equal to the agent's radius. The overall fitness of the agent was determined by averaging its fitness over the 168 evaluation trials.

The fitness measure to evaluate the performance of the agents was determined based on the instructions given to human participants in our experiments. We instructed participants to stop as close as possible to the targets without hitting them. This instruction was covered with a distance and a velocity term in the fitness function. To prevent the use of the "bang-bang" and "slam-on-the-brake" strategies, participants were also instructed to use the brake as smoothly and continuously as possible, avoiding sudden changes in deceleration. This was translated into the fitness function as a jerk term. Therefore, the fitness function had three components to minimize: (1) the distance between the agent and the object without colliding with the object, (2) the velocity of the agent and (3) the total jerk for each trial. Then, the performance measure to be maximized was:

$$\frac{\sum_{i=1}^{NumTrials}\left[((1 - d_i/dMax_i) + (1 - v_i/vMax_i))/2 - W_j jerk_i\right]}{NumTrials} \quad (4)$$

where $NumTrials$ is the total number of trials, d_i is the vertical distance between the agent and the object at the end of the ith trial and $dMax_i$ is the initial vertical distance in that trial, v_i is the agent's velocity at the end of the ith trial, and $vMax_i$ is the agent's initial velocity. W_j is the weight of the jerk component.

The total jerk for each trial was calculated using the following formula (Flash & Hogan, 1985):

$$jerk_i = \frac{1}{2}\int_0^{t_i}\left(\frac{d^3y}{dt^3}\right)^2 dt \qquad (5)$$

where $jerk_i$ is the time integral of the square of the magnitude of the total jerk in the ith trial and t_i is the corresponding trial duration. At the end of a trial, the agent has to apply braking to stop. Therefore, when the agent's velocity was less then or equal to 10% of its initial velocity, jerk did not contribute to the integral. It is important to note that, as there is always some jerk, it is not possible to evolve agents with perfect fitness values using this fitness measure unless the weight of the jerk component is zero. For all of the simulations reported below the maximum trial duration was set to 500 time steps.

We evolved two sets of agents. Rapid changes in deceleration, which is a hallmark of impulsive braking, result in high jerk. This implies that having a jerk term in the fitness function might prevent agents from adopting an impulsive braking strategy. As a result, in the first set of simulations, W_j, the weight of the jerk term in the fitness function, was set to zero. In the second set of simulations, W_j was set to 1000.0, which was determined empirically.

For each set of simulations, five groups of agents, each receiving a different type of visual information, were evolved. The agents in the first group received only b. The input to the remaining four sensors was set to zero. Similarly, the second, third, fourth and the fifth groups received only $\dot{b}$, τ, $\dot{\tau}$ and PR, respectively. The aim was to isolate the optical variables and then examine the braking strategies that they give rise to. From now on, agents in different groups will be referred to by the information they received such as τ agents. Likewise, the first and the second sets of simulations will be referred as Experiment 1 and Experiment 2, respectively. For each group in both experiments, 25 different evolutionary searches were performed with different random seeds, resulting in 25 best-evolved agents for each group.

3.4 Performance evaluation

We performed in-depth behavioral analyses of the evolved agents to uncover the optical variables and the control strategies that could effectively be used to perform the braking task. In addition to the optical variables, the fitness measure has also an impact on the agents' behavior. In both experiments, a multi-criteria fitness function was used to evolve the agents. As a result, it is possible that an agent might attain a relatively high fitness value by performing very well with respect to one criterion, but performing on average with respect to another. As both crashes and stopping too short of the object indicate poor use of visual information, it is important to examine the agents' performance in more detail, with respect to each of the distance and velocity criteria of the fitness function.

For this purpose, we defined four classes of trials: (1) crash trials: the trials in which the final velocity of the agent (v_{final}) was greater than zero when the final distance between the agent and the object (d_{final}) was zero; (2) premature stops: the trials in which the agents stopped too early, too far away from the object. Any trial in which $d_{final} > 15.0$ (i.e., the agent's radius) was considered a premature stop; (3) failure-to-reach trials: the trials that ended not because of a crash or a premature stop, but because the maximum trial duration of 500 time steps was reached. Failure-to-reach trials together with crash trials and premature stops constitute the unsuccessful trials; (4) successful trials: the trials in which the agents stopped safely, very close to the object. Any trial in which $0 \leqslant d_{final} \leqslant 15.0$ when $v_{final} = 0$ was considered a successful trial. It is important to note that we do not expect either our human participants or the evolved agents to have 100% successful trial rate. Given the nature of the task, crashes and premature stops do occur and are acceptable as long as their numbers are small. Our analysis focuses on the performance and the behavior of the agents in successful trials.

4 The results

4.1 Testing the constraints for a given behavioral capability

What can these evolved brain–body–environment systems tell us about the visual control of braking in humans? They can allow us to explore the constraints on the braking behavior. To illustrate, we first evolved a set of agents using a fitness function, which minimized only the velocity of the agent and the distance between the agent and the object. In these simulations, the trial duration was not limited. The results showed that all of the evolved agents reduced their velocities to almost zero values right at the beginning of the trials and then slowly drifted to the object, resulting in extremely long trial durations. Besides, some of these evolutionary searches could not be completed within feasible time because of the near-zero velocities of the agents and they had to be stopped before reaching the maximum number of generations. This type of behavior was never observed in humans.

Table 1. The numbers of successful and unsuccessful trials for the best agents in Experiment 1 and Experiment 2.

	Unsuccessful trials			Successful trials	Successful trial rate
	Crash trials	Premature stops	Failure-to-reach trials		
Experiment 1					
The b agent	9	4	4	151	89.9%
The $\dot{b}$ agent	8	0	0	160	95.2%
The τ agent	20	0	0	148	88.1%
The $\dot{\tau}$ agent	4	0	36	128	76.2%
The *PR* agent	0	0	0	168	100.0%
Experiment 2					
The b agent	67	78	0	23	16.7%
The $\dot{b}$ agent	47	70	2	49	29.2%
The τ agent	28	16	0	124	73.8%
The $\dot{\tau}$ agent	52	0	8	108	64.3%
The *PR* agent	20	32	0	116	69.0%

To prevent agents exhibiting this behavior, trial duration was included in the fitness function, which forced all agents to adopt the slam-on-the-brake strategy. This suggests that time is a key constraint that gives rise to the slam-on-the-brake strategy. We, then, excluded trial duration from the fitness function but time was included indirectly in the fitness measure by limiting the trial duration. This change in the fitness function successfully eliminated the slam-on-the-brake strategy. However, all agents were found to be using an impulsive braking strategy, details of which will be discussed in more detail in the following section. Continuously regulated braking was observed only when trial duration was limited and jerk was also included in the fitness function. This finding suggests that jerk is also an important constraint that affects braking behavior. It is especially important in continuously regulated braking.

4.2 Testing the existing hypotheses, forming new ones

4.2.1 Comparison of the optical variables. What impact do different optical variables have on braking performance? Did all of the agents perform the task comparably well? In other words, did all optical variables result in successful performance? Or are there any optical variables that yielded significantly better performance? To answer these questions, we first examined how well the different groups of agents in each experiment performed the task.

The fitness values of the best-evolved agents in Experiment 1 and Experiment 2 can be seen in Figures 3a and 3b, respectively. The gray circles represent the mean fitness values of the 25 best-evolved agents, with error bars indicating ±1 standard deviation. The black squares correspond to the best-evolved agent that attained the highest fitness value out of 25 agents.

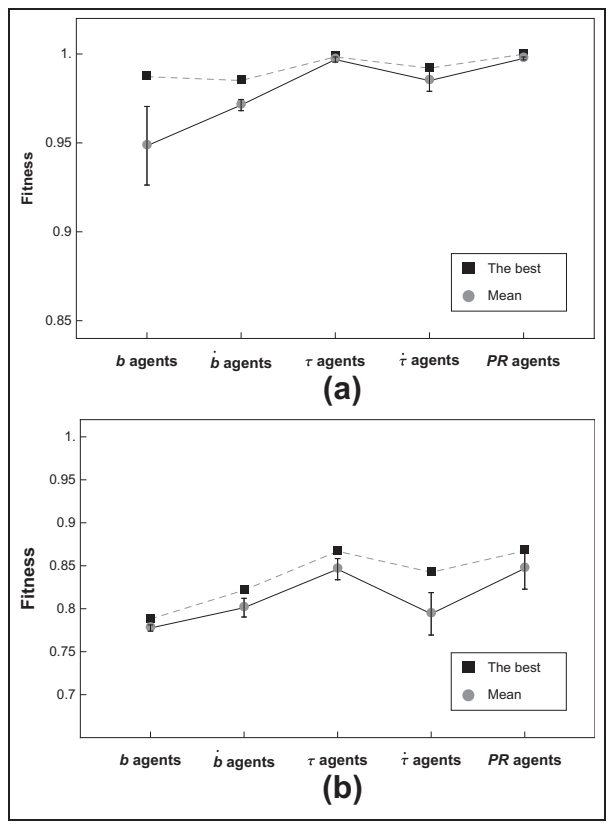

Figure 3. Fitness values of the best-evolved agents (a) in Experiment 1, in which the weight of the jerk term (W_j) was 0.0 and (b) in Experiment 2, with W_j=1000.0. The gray circles represent the mean fitness values of the 25 best-evolved agents, with error bars indicating ± 1 standard deviation. The black squares correspond to the best-evolved agent that achieved the highest fitness value out of 25 agents.

From now on, the best of the best-evolved 25 agents in each group will be referred as "the best" agent for that group.

The numbers of successful and unsuccessful trials of the best agents are given in Table 1. In Experiment 1,

all of the best agents performed the task with a fitness value >98% and a successful trial rate >75%. The most successful agent was the best *PR* agent (99.99%) with no unsuccessful trials and it stopped right at the object (see Table 2 in Appendix A for the final stopping distances). The next successful agent was the best τ agent (99.84%), which was followed by the best $\dot{\tau}$ agent (99.20%), the best b agent (98.72%), and the best $\dot{b}$ agent (98.51%), respectively.

To confirm these results, a Kruskal–Wallis[1] test on the fitness values of all 125 best-evolved agents (5 groups × 25 best agents in each group) was conducted to compare the effect of the optical variables on the fitness values. The analysis indicated a significant effect of the optical variable, $\chi^2(4) = 104.12$, $p<.001$, with an effect size of 0.840. *Post hoc* comparisons using Mann–Whitney tests with Bonferroni correction revealed that the fitness values of the $\dot{b}$ agents were significantly higher than those of the b agents ($p<.001$). The fitness values of the $\dot{\tau}$ agents were significantly higher than those of $\dot{b}$ agents ($p<.001$). The only non-significant comparison was between the τ agents and those of *PR* agents ($p = 1.0$), indicating the τ agents and the *PR* agents performed the task equally well. Both τ and *PR* agents have fitness values that were significantly higher than those of the $\dot{\tau}$ agents (both at $p<.001$ significance level).

In Experiment 2, the best b and the $\dot{b}$ agents could not perform the task at all. Both had a very low successful trial rate (Table 1). As a result, these agents will be excluded from further analysis. The best performances were again achieved by the best *PR* agent (86.73%) and the best τ agent (86.69%). These agents were followed by the best $\dot{\tau}$ agent (84.24%). The Kruskal–Wallis test on the fitness values of the 125 best agents evolved in Experiment 2 again revealed a significant effect of the optical variable, $\chi^2(4) = 85.27$, $p<.001$, with an effect size of 0.688. As in Experiment 1, *post hoc* comparisons using Mann–Whitney tests with Bonferroni correction indicated that the fitness values of the *PR* agents were not significantly different than those of the τ agents ($p = 1.0$). Both groups performed significantly better than the $\dot{\tau}$ agents (both at $p<.001$ significance level). Even though the best $\dot{\tau}$ agent achieved a performance close to the best *PR* and τ agents, pairwise comparisons revealed that overall the $\dot{\tau}$ agents did not perform the task significantly better than the b agents ($p = 1.0$) and the $\dot{b}$ agents ($p = .95$). Therefore, the $\dot{\tau}$ agents are also considered unsuccessful and will be excluded from further analysis.

Why were the τ and *PR* agents successful in performing the task, but not the $\dot{\tau}$ agents? As previously stated, unlike $\dot{\tau}$ that stays constant at -1.0 unless braking is initiated, both τ and *PR* evolve over time. As a result, the τ and *PR* agents experienced change in the visual information from the beginning of the trials. However, the $\dot{\tau}$ agents had to initiate braking not only to stop but also to experience change in the visual information, which probably made the task harder for the $\dot{\tau}$ agents. The poorer performance of the $\dot{\tau}$ agents in Experiment 2 can be taken as evidence indicating that, when isolated from other variables, $\dot{\tau}$ alone is not sufficient to control braking successfully. What this means is that with $\dot{\tau}$, there is no information that can be used to indicate when to start braking. In contrast, the evolving τ and *PR* trajectories can be so used.

To sum up, in Experiment 1, all best agents performed the task successfully whereas only the τ and *PR* agents were successful in both experiments. The τ and *PR* agents were also the best performing agents in both experiments and the analysis revealed that they performed the task equally well. These results suggest that when the optical variables are isolated, the optical variables τ and *PR* result in better braking performance than b, $\dot{b}$, and $\dot{\tau}$. Now that we have identified the agents that can successfully solve the braking task, the next step in our analysis is to investigate these agents' behavior in more detail.

4.2.2 Impulsive vs. continuously regulated braking. How does the behavior of the successful agents change depending on the optical variable and on the fitness measure? Do the agents exhibit impulsive braking or do they regulate deceleration continuously? To answer these questions, we examined the velocity and deceleration trajectories of the successful agents, together with the resulting times series of the optical variables.

The aim of the Experiment 1 was to explore the informational basis of impulsive braking and all five best agents evolved in Experiment 1 were found to be using an impulsive braking strategy. The best b agent used the slam-on-the-brake and the bang-bang strategies. The best $\dot{b}$ agent was found to be using a pure bang-bang strategy. As expected, these strategies resulted in highly variable stopping distances (see Table 2 in Appendix A). The analysis of the τ trajectories of the best τ agent in Experiment 1 suggested that this agent used the weak version of the constant τ strategy (Figure 4). In other words, the agent used τ to determine the initiation and termination of deceleration (Figure 5). The mean τ value at which the first braking is initiated is 6.68 ($SD = 0.19$). The mean τ value, at which the brake was released, is 13.79 ($SD = 0.40$) and the mean τ at which the final braking is initiated is 7.59 ($SD = 0.04$).

All three agents examined so far exhibited discrete impulsive braking, in which brake is used in an on-or-off fashion. However, what separates the best τagent from the best b and $\dot{b}$ agents is that it initiated and terminated braking on the basis of the optical variable τ. This braking strategy resulted in an 88.10% successful trial rate and the agent stopped very close to the target with a very small variation in the final stopping distances (see Table 2 in Appendix A). This suggests that

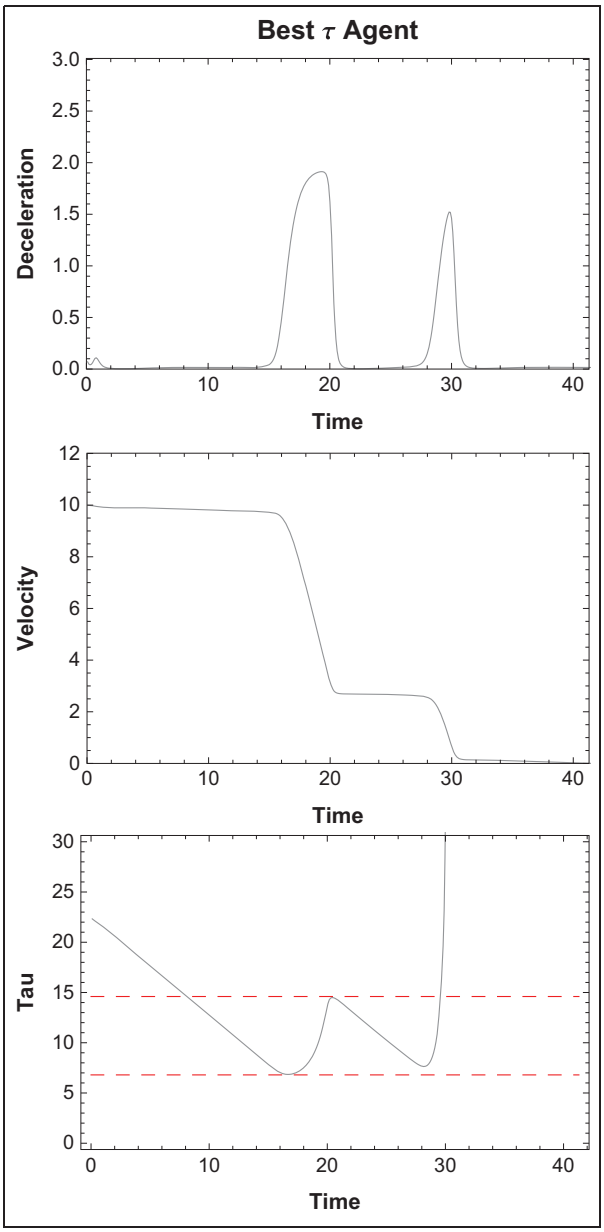

Figure 4. Sample velocity, deceleration, and τ trajectories of the best τ agent in a successful trial in Experiment 1. Dashed lines represent the τ values corresponding to the initiation and termination of deceleration.

even though the weak version of the constant τ strategy is impulsive, it is a very efficient impulsive braking strategy that brought the agent to a stop very close to the target in the majority of the trials.

The deceleration trajectories of the best $\dot{\tau}$ and PR agents in Experiment 1 revealed a second type of impulsive braking, in which brake is applied in a periodic fashion, creating oscillations in deceleration. Oscillations in deceleration, in turn, cause oscillations in τ, $\dot{\tau}$ and PR (Figure 6). This oscillatory impulsive braking behavior has not been reported in the literature on visually guided braking before. Even though the best $\dot{\tau}$ and PR agents used the brake in a periodic

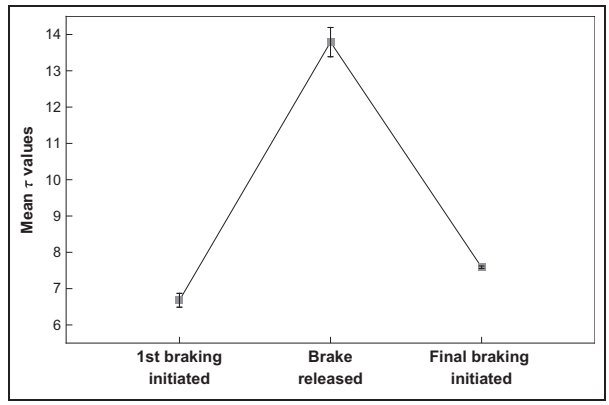

Figure 5. The mean values of the best τ agent in Experiment 1, at which braking is initiated and terminated. Error bars represent ± 1 standard deviation.

fashion, the oscillations in deceleration created oscillations with increasing amplitude in the $\dot{\tau}$ and PR trajectories of the best $\dot{\tau}$ agent. However, the oscillatory deceleration of the best PR agent produced oscillations with decreasing amplitudes in $\dot{\tau}$ and PR trajectories. Unlike the oscillations with increasing amplitude, the decreasing amplitude of the oscillations in the $\dot{\tau}$ and PR trajectories of the best PR agent kept those variables bounded within a region. This might be the reason why the best PR agent performed the task better than the best $\dot{\tau}$ agent.

To recapitulate, there are two main findings in Experiment 1: (1) two different classes of impulsive braking behaviors were observed, including one not reported in previous studies: discrete impulsive braking, in which the brake is used in an on/off fashion and oscillatory impulsive braking, in which the brake is used in a periodic fashion; and (2) the weak version of the constant τ strategy is an efficient impulsive braking strategy, in which the optical variable τ is used to initiate and terminate braking.

The aim of Experiment 2 was investigate optical variables and the control strategies used in continuously regulated braking. The only successful agents in Experiment 2 were the best τ. and PR agents. Both agents regulated their deceleration continuously in all trials, indicating that including a minimum jerk criterion in the fitness measure, indeed, resulted in continuously regulated braking (Figure 7). Two competing strategies for the continuously regulated braking are the constant $\dot{\tau}$ strategy (Lee, 1976; Yilmaz & Warren, 1995) and the proportional rate control (Anderson & Bingham, 2010, 2011). The next step in our analysis is to investigate whether the best τ and PR agents evolved in Experiment 2 used one of these strategies.

4.2.3 Constant $\dot{\tau}$ strategy vs. proportional rate control. CTRNNs are proved to be universal approximators of smooth dynamics (Funahashi & Nakamura,

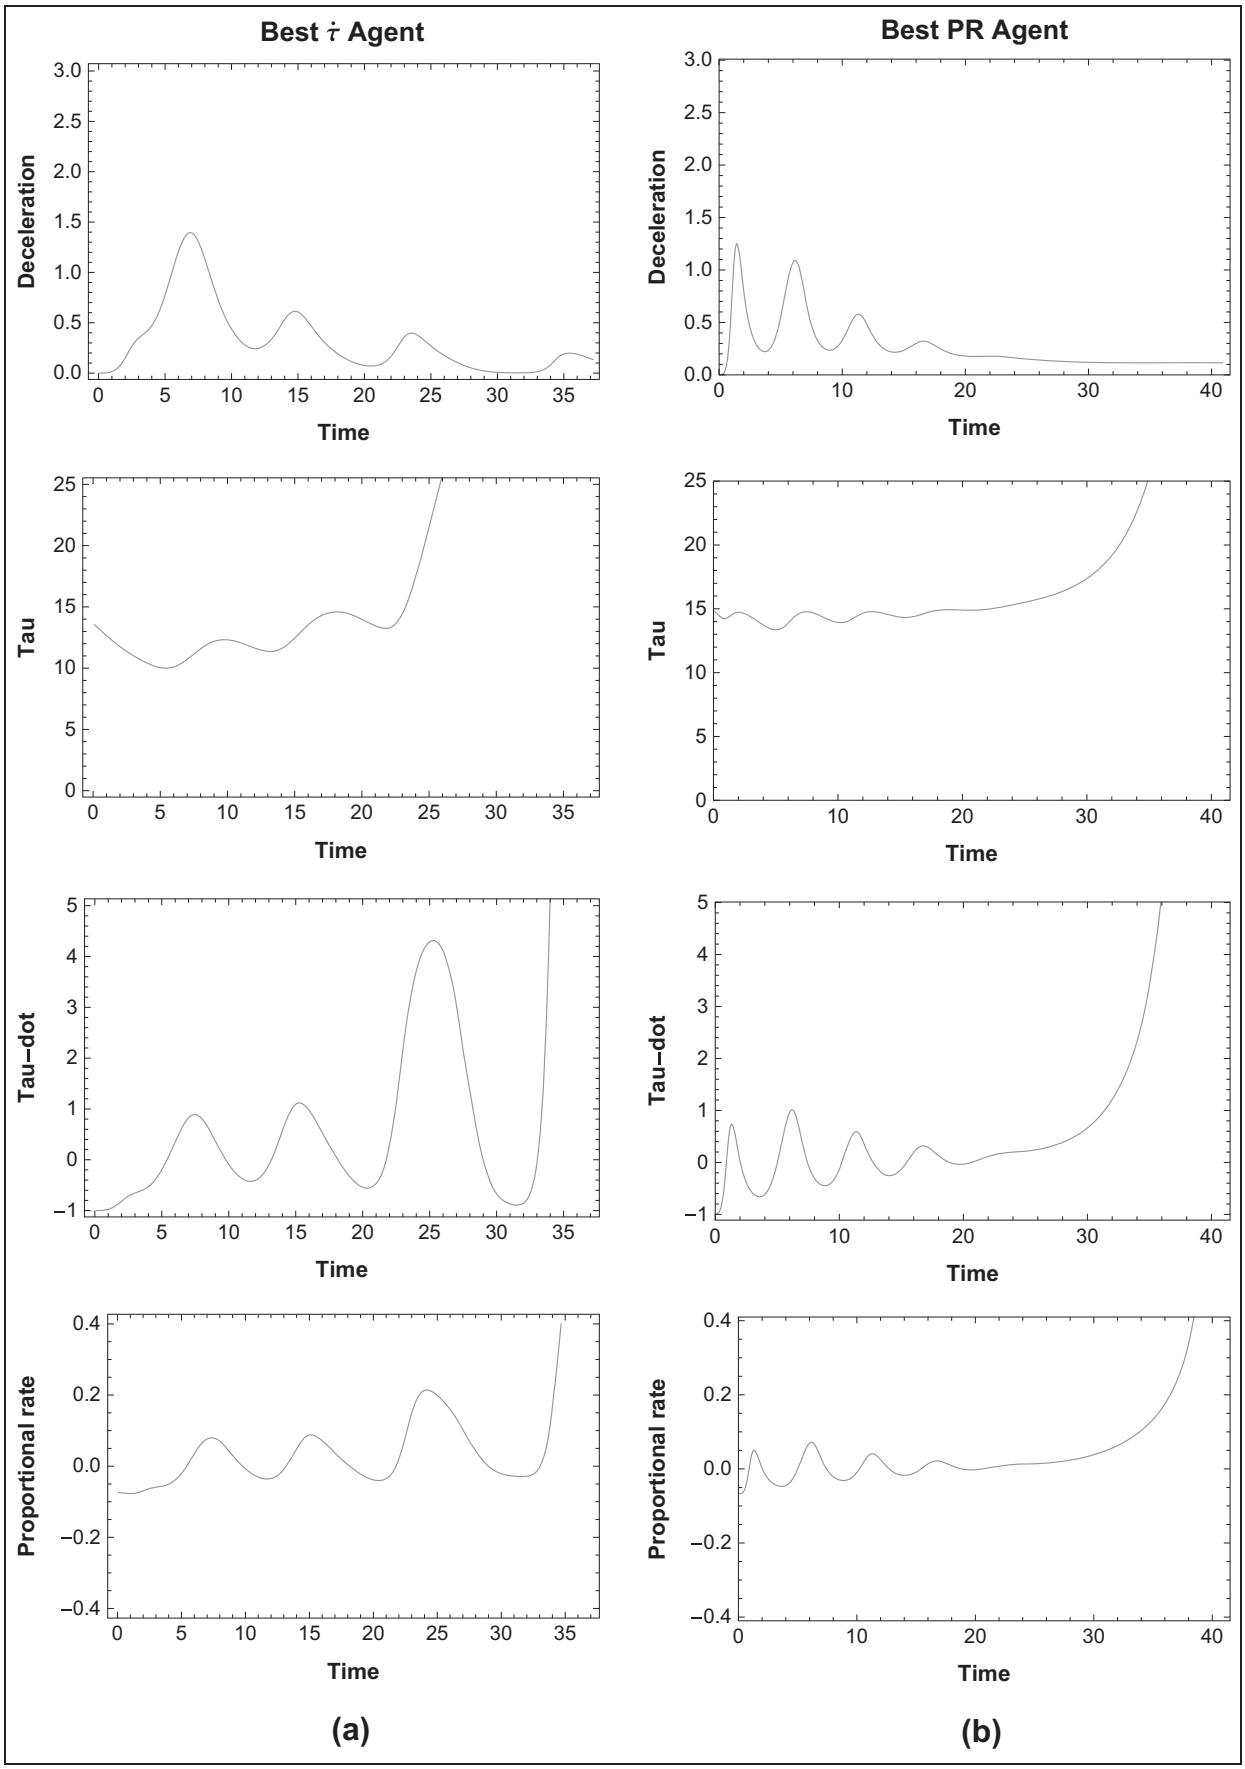

Figure 6. Sample deceleration trajectories of the (a) best $\dot{\tau}$ agent and (b) best *PR* agent in a successful trial in Experiment 1, together with the corresponding τ, $\dot{\tau}$, and *PR* trajectories.

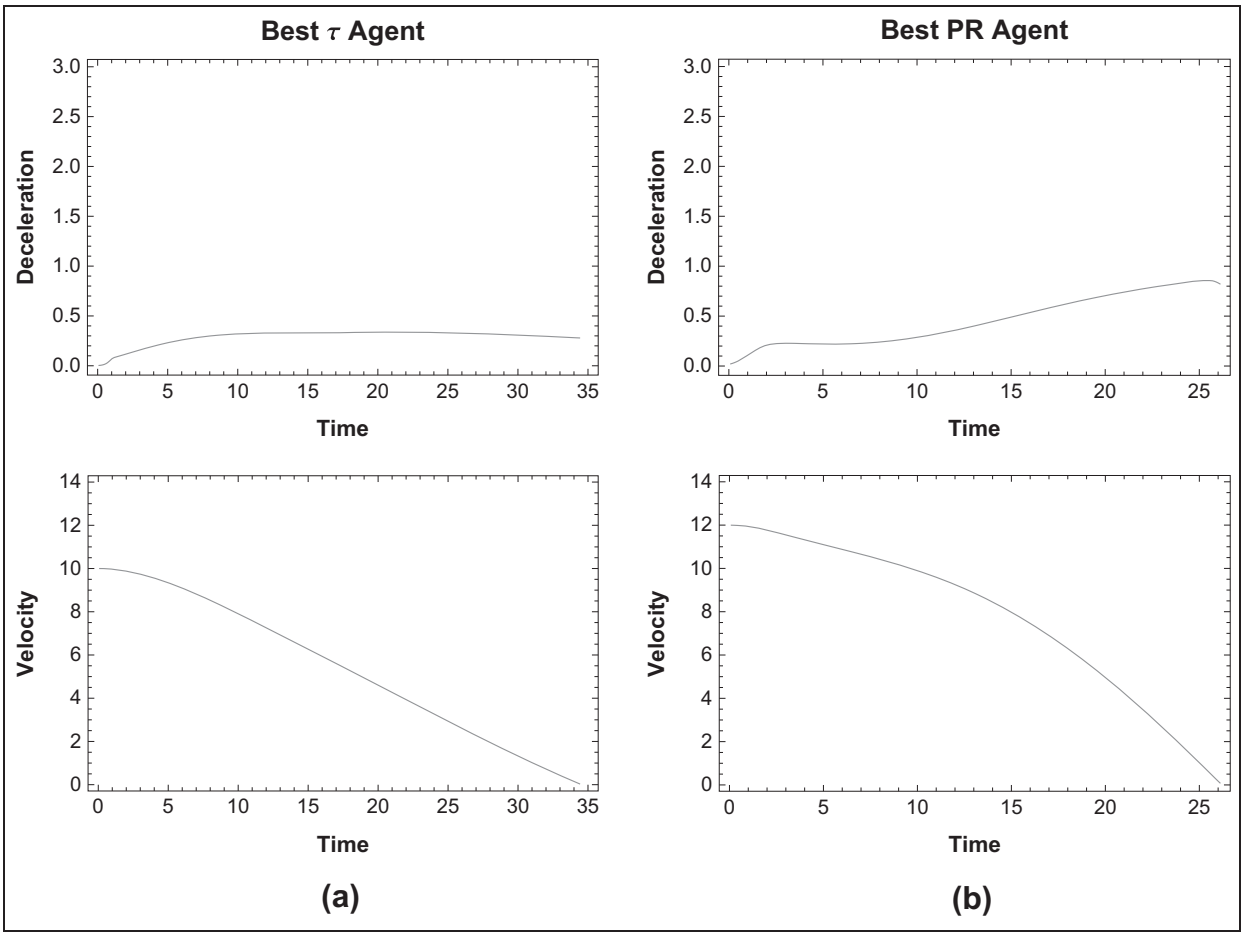

Figure 7. Sample velocity and deceleration trajectories of (a) the best τ agent and (b) the best *PR* agent in Experiment 2.

1993; Kimura & Nakano, 1998). This means that even though the best τ and PR agents evolved in Experiment 2 received only τ and PR as the visual information, they could, in theory, derive $\dot{\tau}$ and PR fr the information that they received and use them to solve the task. Therefore, we analyzed both the $\dot{\tau}$ and PR trajectories of these agents.

As in Yilmaz and Warren (1995), before performing any analysis, we first windowed the trajectories to isolate the portions of the trials, in which deceleration was regulated continuously. To keep the total jerk low, the agents start braking slowly right at the beginning of the trials. For a brief period, the motor output of the agents is the same for all trials. It, then, starts to diverge. The point at which the deceleration trajectories start to diverge was taken to be the beginning of the window. The end point of the window was chosen based on the τ trajectories so that the final increase in τ to infinity together with the preceding small plateau was eliminated.

Sample τ, $\dot{\tau}$, and PR trajectories of the best τ and PR agent can be seen in Figures 8a and 8b, respectively. The dotted lines represent the beginning and end of the window. The analyses reported below were applied to the portion of the trajectories that fell within the window. For each agent, we first performed a linear regression analysis of the time series of τ to estimate the mean $\dot{\tau}$. As in Yilmaz and Warren (1995), for each trial, the slope of the regression line was taken as the mean $\dot{\tau}$ for that trial. For the best τ agent, the overall mean $\dot{\tau} = -0.57$ across all trials, with $SD = 0.068$. For the best PR agent, the overall mean $\dot{\tau} = -0.61$, with $SD = 0.038$. For both agents, linear regression resulted in R^2 values >0.96 in all trials. These results show that both agents had an overall mean $\dot{\tau}$ around -0.6, which is slightly smaller than $\dot{\tau} = -0.5$ as predicted by the constant $\dot{\tau}$ strategy.

We then performed the split-half analysis on the τ trajectories of the best τ and PR agents to determine whether $\dot{\tau}$ was constant or changing during approach. The split-half analysis with a paired t-test on the mean $\dot{\tau}$ values of the best τ agent revealed a significant difference between the $\dot{\tau}$ values in the first and second halves, $t(30) = -43.54$, $p < .001$, indicating that $\dot{\tau}$ was not held constant. The mean $\dot{\tau}$ in the first half was -0.70 with $SD = 0.065$ and the mean $\dot{\tau}$ in the second half was -0.46 with $SD = 0.070$ (Figure 9a). Consistent with the previous findings in the literature, the slope of the first

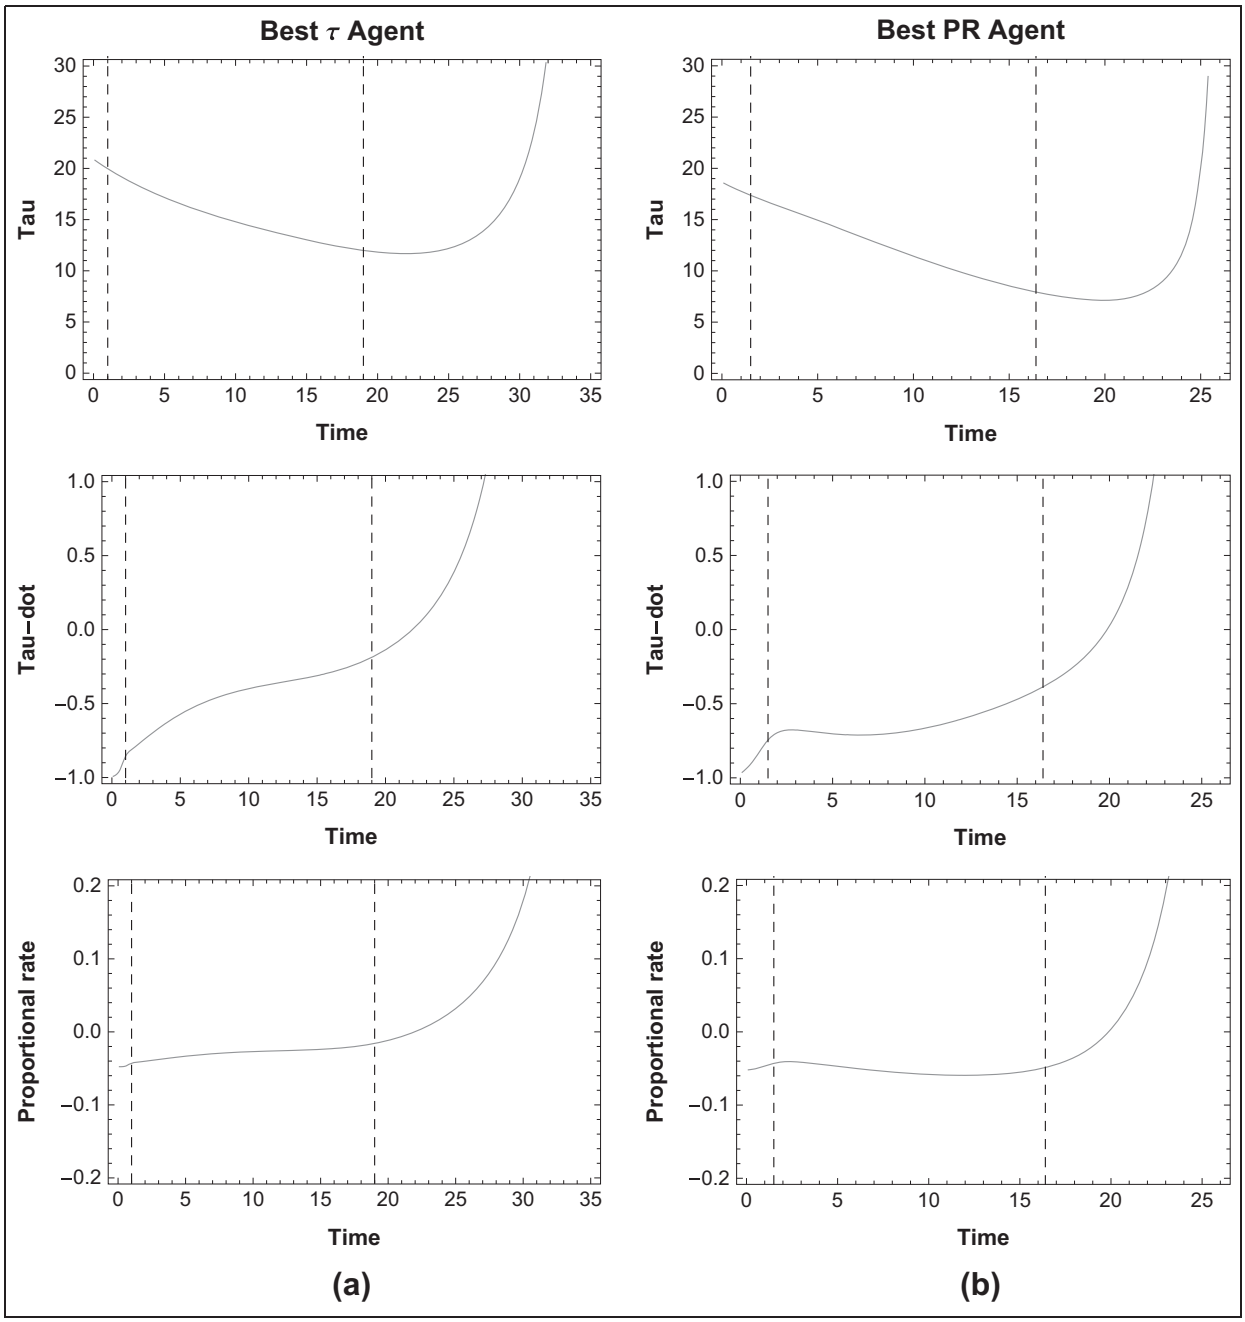

Figure 8. Sample τ, $\dot{\tau}$, and *PR* trajectories of (a) the best τ agent and (b) the best *PR* agent in Experiment 2. The dotted lines represent the beginning and end of the window. Statistical analyses were performed on the portions of the trajectories that fell within the window.

half was greater than that of the second half. These results indicate that the best τ agent did not use the constant $\dot{\tau}$ strategy to control braking.

The split-half analysis on the mean $\dot{\tau}$ values of the best PR agent also revealed a significant difference between the mean $\dot{\tau}$ values in each half, $t(28) = -15.14$, $p<.001$. The mean $\dot{\tau}$ in the first half was -0.69 with $SD = 0.024$ and the mean $\dot{\tau}$ in the second half was -0.52 with $SD = 0.055$ (Figure 9b). Again, the slope of the first half was greater than that of the second half. Finding a significant difference indicates that $\dot{\tau}$ was not held constant during approach, which means that the best PR agent did not use the constant $\dot{\tau}$ strategy, either.

We, then, performed the split-half analysis with a paired t-test on the mean PR values of the best τ and PR agents. For the best τ agent, the split-half analysis revealed a significant difference, $t(30) = -6.83$, $p<.001$. The mean PR values in the first and second halves were -0.06 (with $SD = 0.015$) and -0.05 (with $SD = 0.016$),

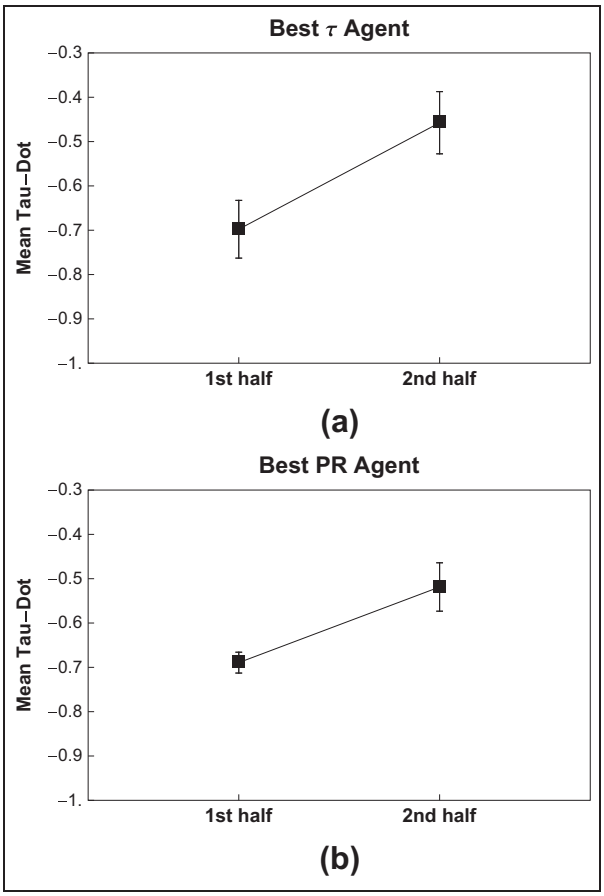

Figure 9. The mean $\dot{\tau}$ values for (a) the best τ agent and (b) the best *PR* agent in Experiment 2. Error bars represent $\pm$ 1 standard deviation.

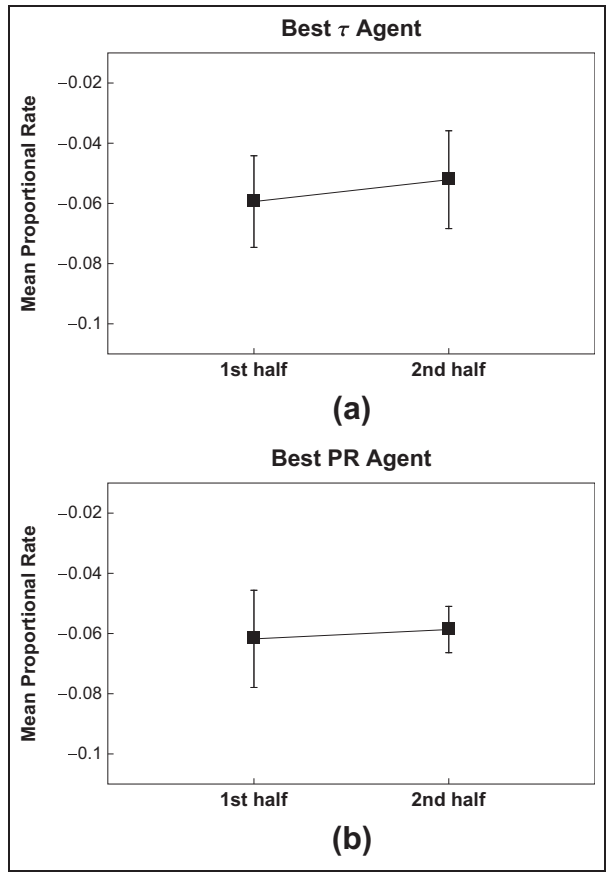

Figure 10. The mean *PR* values for (a) the best τ agent and (b) the best *PR* agent in Experiment 2. Error bars represent $\pm$ 1 standard deviation.

respectively (Figure 10). However, finding a significant difference between the mean PR does not completely rule out the use of proportional rate control, as proportional rate control allows a range of PR values that will result in successful braking. Given that the mean PR values in each half are very close to each other, it can still be argued that the best τ agent used a version of the proportional rate control, in which the PR was not held constant, but kept bounded within a region.

The split-half analysis on the mean PR values of the best PR agent reveled no significant difference, $t(28) = -1.503$, $p = .144$. The mean PR values in the first and the second halves were 0.062 (with $SD = 0.016$) and 0.059 (with $SD = 0.008$), respectively (Figure 10). The lack of significance indicates that the PR values were kept constant during approach, suggesting the best PR agent used proportional rate control.

To sum up, contrary to the evidence provided by Yilmaz and Warren (1995), our results provide no evidence in favor of a $\dot{\tau}$-based strategy such as the constant $\dot{\tau}$ strategy. A main assumption made in our simulations is that only one optical variable is available to be used to control braking. Under such circumstance, our results support the hypothesis that the optical variable proportional rate is alone sufficient to control braking successfully and proportional rate control is used in continuously regulated braking. Indeed, two different proportional rate control strategies were observed: (1) the weak version, in which the PR was not held constant, but kept bounded within a region and (2) the strong version, in which the PR was held constant to stop successfully at the target.

5 General discussion

The field of evolutionary robotics has emerged from the need to develop an automatic design process for building control systems for robots. However, researchers have recently begun to use the evolutionary robotics techniques as a modeling methodology for studying cognition and behavior. Research in the last decade has proven that these techniques can be useful tools in the study of cognition and behavior in different ways. For instance, they can provide existence proofs. Izquierdo-Torres and Harvey (2007) challenged the view that synaptic plasticity is necessary for learning to occur by evolving CTRNNs with fixed weights that can produce Hebbian-like learning. The analysis of the network

revealed that the learning behavior arises from the interaction between slow and fast acting components of the network. Izquierdo-Torres and Harvey's (2007) results do not, by all means, indicate that the same mechanism is also responsible for learning in biological systems. They also do not undermine the importance of synaptic plasticity. However, they open up new perspectives on the mechanisms underlying learning and can help us formulate new theories and hypotheses that can be tested empirically.

Using an approach that they call "comparative psychology of evolved agents and people,"Ward and Ward (2008) investigated the psychological principles underlying selective attention by employing evolutionary robotics simulations and experiments involving human subjects. Their work demonstrated the presence of reactive inhibition in both the evolved agents and people, suggesting that it is an essential mechanism for selective attention.

The evolutionary robotics techniques have also been applied to model experimental paradigms from psychology. Based on Auvray, Lenay, and Stewart's (2009) minimalist perceptual crossing experiment, Di Paolo, Rohde, and Iizuka (2008) developed a series of evolutionary robotics simulations to investigate the dynamics of social interaction. They successfully replicated the results of the original study, indicating that the recognition of the presence of another's agency does not necessarily require complex cognitive mechanisms and simple solutions can emerge from the dynamics of the interaction process itself. Their results also generated additional hypotheses that can be tested empirically.

In light of these studies, the present work utilized evolutionary robotics techniques as a hypothesis generator to investigate the visual information and the control strategies used in the context of a driving-like braking task. Based on a widely used experimental paradigm from psychology, a series of evolutionary robotics simulations were developed to aid in resolving the existing questions as to the control strategies underlying continuously regulated braking and to investigate the information variable(s) and the control strategy(ies) used in impulsive braking. Model agents were evolved to solve a braking task in a simple 2D environment containing one stationary object. The paper focused on the in-depth behavioral analyses of the evolved agents to uncover the optical variables and the control strategies used by the agents to solve the braking task.

Based on the geometry of the looming objects and the cost function used in the simulations, it is also possible to design an optimal feedback controller for the braking task being investigated in this paper. Providing extra constraints on the derived controller can result in different solutions and much can be learned from such an analysis. However, designing an optimal controller requires choosing a class of controller to utilize. This choice can greatly influence the results, constituting additional assumptions of the model. Furthermore, only the simplest such classes can be analytically optimized. We chose to evolve CTRNNs to let the agents find their own control strategies, and thus, to explore a larger space of possible braking strategies with minimal a priori assumptions.

The results of our study can be summarized under five main findings. The first main finding is that behaviors similar to human behavior were observed only when the constraints were selected correctly. Our results confirm that time and jerk are important constraints that shape braking behavior. More specifically, time is a major constraint on impulsive braking and jerk is a key constraint on continuously regulated braking. If we consider the cases in which constraints that underlie an observed behavior are not clear or known, our results suggest that evolutionary robotics simulations can be useful tools in determining the necessary and/or sufficient constraints for a given behavior.

The second main finding of our simulations is that when only one optical variable is available to control braking, τ and PR are the most efficient optical variables to solve a driving-like braking task investigated in this paper, as only these two variables gave rise to successful braking performance in both experiments. When the task demands were increased by adding a jerk term, the $\dot{\tau}$ agents could not perform the task successfully. Why did $\dot{\tau}$ agents fail to perform the task while both τ and proportional rate agents were very successful? We argue that the answer lies in the nature of the optical variables. While approaching an object with a constant velocity, both τ and PR evolve over time, resulting in a change in the visual information. However, $\dot{\tau}$ stays constant at -1.0 unless braking is initiated. This implies that $\dot{\tau}$ agents had to "act" to experience change in the visual information, which probably made the task "harder" for the $\dot{\tau}$ agents, as a result, gave rise to a poor performance. The failure of the $\dot{\tau}$ agents in Experiment 2 can be taken as evidence indicating when isolated from other optical variables, $\dot{\tau}$ alone is not sufficient to control braking.

The third major finding of the present work is that two different classes of impulsive braking behaviors were observed: (1) discrete impulsive braking, in which the brake is used in an on-or-off fashion; and (2) oscillatory impulsive braking, in which brake is used in a periodic fashion. There were two versions of oscillatory impulsive braking. In the first version, the oscillations in deceleration created oscillation with increasing amplitudes in τ, $\dot{\tau}$ and PR. In the second one, the oscillations in deceleration created oscillation with decreasing amplitudes in τ, $\dot{\tau}$, and PR. Both of these oscillatory braking behaviors have not been reported in the literature before.

As previously indicated, no studies have yet investigated the optical variable(s) and the control strategy(ies) used in impulsive braking. The fourth main

finding of our work is that one promising impulsive braking strategy is to use τ to initiate and terminate braking. The best τ agent in Experiment 1 was found to be using the weak version of the constant τ strategy proposed by Bingham (1995). Even though, this is a discrete impulsive braking strategy, it is very efficient in that it brought the agent to a stop right at the object.

Finally, the fifth main finding of our simulations is that our results cast doubt on the use of the constant $\dot{\tau}$ strategy to control braking. Contrary to the evidence provided by Yilmaz and Warren (1995), none of the evolved agents was found to be using a $\dot{\tau}$-based strategy such as the constant $\dot{\tau}$ strategy. Our results support the hypothesis that the proportional rate control is used when regulating the brake continuously. Two different proportional rate control strategies were observed: (1) the weak version, in which the proportional rate was not held constant during braking, but kept bounded within a region; and (2) the strong version, in which the proportional rate was held constant to stop successfully at the target.

One of the criticisms of using evolutionary robotics simulations in the study of human behavior and cognition is that these evolved agents are too simple to tell us anything about human behavior. This is a fair criticism. However, what is generally ignored in this criticism is that CTRNNs are universal approximators of smooth dynamics (Funahashi & Nakamura, 1993; Kimura & Nakano, 1998). This means that the mathematics underlying these simple brain–body–environment systems are quite powerful. Our results show that once the constraints on behavior are selected correctly and the appropriate fitness measure is chosen, it is possible to capture important aspects of human behavior even with these simple agents, which have "brains" consisting of four neurons.

One important aspect of continuously regulated braking that was not captured by our simulations is the initiation of braking. Our results do not provide an answer to the question when humans initiate braking. We argue that this is probably due to the minimum jerk criterion included in the fitness measure. The minimum jerk criterion caused the agents to initiate braking as soon as trials started. In our simulations, the fitness measure minimized the total jerk in a trial. This means that the agents were penalized even for the slightest jerk. The effect of the jerk in the fitness measure can be alleviated by: (1) punishing the agents if the jerk is higher than a certain threshold value; or (2) penalizing agents based on the maximum jerk in a trial. With these modifications, it might be possible to capture the initiation of braking.

As stated above, our results revealed a second type of impulsive braking in which brake is used in a periodic fashion, resulting in oscillations in τ, $\dot{\tau}$, and proportional rate. Then, the question is: what causes those oscillations? How and when do they arise? Relatedly, those oscillations should disappear when switching from impulsive braking to continuously regulated braking. Then, what causes those oscillations to disappear? Such questions can be answered by analyzing the underlying dynamics of the agent–environment system, using tools from dynamical systems theory. The results obtained from such analysis may provide insights into the control of braking in humans and help us form new hypotheses that can be tested in further experiments.

Unlike the natural environment we live in, which provides abundant information, the simulated environment of the agents provided only one type of information. Our aim was to isolate each type of optical variable and to investigate what kind of braking strategies emerge for each of these variables. Our results indicated that proportional rate is alone sufficient to control braking successfully whereas $\dot{\tau}$ is not. However, allowing agents to receive more than one optical variable could yield different results. For example, if agents receive $\dot{\tau}$ together with τ, it might be possible to observe a $\dot{\tau}$-based strategy as proposed by Yilmaz and Warren (1995). Therefore, the next step in the analysis is to evolve agents that receive a combination of optical variables and to investigate what kind of braking strategies emerge from combinations of different optical variables.

Finally, the third step in the analysis is to allow agents to receive all available information. After evolving these agents, biological lesioning and informational lesioning analyses (Keinan, Meilijson, & Ruppin, 2003; Aharonov, Segev, Meilijson, & Ruppin, 2003) can be carried out to investigate whether it is possible to observe a tendency in the evolved agents to rely on certain optical variables even if all the variables are available. We believe that such analyses can provide more insights into the optical variables and the control strategies used in visually guided braking.

In addition to isolating the optical variables, another assumption made in the evolutionary robotics simulations reported in this paper is the absence of biological noise in extracting the information variables from the optic flow. Introducing noise into the optical variables and allowing the noise to propagate in the system could affect the behavior of the agents and lead to changes in the observed braking strategies. Another possible future direction that can be pursued both experimentally and using evolutionary robotics simulations is to include acceleration in the task. Finding support in favor of the proportional rate control in such a study can provide further evidence that proportional rate control is used in visually guided braking.

Acknowledgments

This work was supported in part by NSF grant IIC-1216739 to RDB.

Note

1. Even though, in theory, the fitness values change between 0.0 and 1.0, the evolutionary search selects for the best performing agents, which causes our data set to be skewed towards the higher end of the fitness distribution, making it non-normal. Therefore, a non-parametric test was used to test the effect of the different optical variables on the fitness values of the evolved agents.

References

Aharonov, R., Segev, L., Meilijson, I., & Ruppin, E. (2003). Localization of function via lesion analysis. *Neural Computation*, *15*, 885–913.

Anderson, J., & Bingham, G. P. (2010). A solution to the online guidance problem for targeted reaches: Proportional rate control using relative disparity τ. *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.

Auvray, M., Lenay, C., & Stewart, J. (2009). Perceptual interactions in a minimalist virtual environment. *New Ideas in Psychology*, *27*, 32–47.

Beer, R. D. (1995). On the dynamics of small continuous-time recurrent neural networks. *Adaptive Behavior*, *3*(4), 471–511.

Beer, R. D. (2003). The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior*, *11*, 209–243.

Beer, R. D. (2014). Dynamical systems and embedded cognition. In K. Frankish, & W. Ramsey (Eds.), *The Cambridge handbook of artificial intelligence* (pp. 128–148). Cambridge, UK: Cambridge University Press.

Bingham, G. P. (1995). The role of perception in timing: Feedback 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.

Di Paolo, E. A., Rohde, M., & Iizuka, H. (2008). Sensitivity to social contingency or stability of interaction? Modelling the dynamics of perceptual crossing. *New Ideas in Psychology*, *26*(2), 278–294.

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: Hougton Miffin.

Holland, J. (1975). *Adaptation in natural and artificial systems*. Ann Arbor, MI: University of Michigan Press.

Iizuka, H., & Di Paolo, E. A. (2007). Minimal agency detection of embodied agents. In F. Almeida e Costa, L. M. Rocha, E. Costa, I. Harvey, & A. Coutinho (Eds.), *Advances in artificial life: Proceedings of the 9th European Conference on Artificial Life* (pp. 485–494). Berlin, Germany: Springer-Verlag.

Izquierdo, E., Harvey, I., & Beer, R. D. (2008). Associative learning on a continuum in evolved dynamical neural networks. *Adaptive Behavior*, *16*(6), 361–384.

Izquierdo, E. J., & Beer, R. D. (2013). Connecting a connectome to behavior: An ensemble of neuroanatomical models of *C. elegans* klinotaxis. *PLoS Computational Biology*, *9*(2): e1002890. doi:10.1371/journal.pcbi.1002890.

Izquierdo-Torres, E., & Harvey, I. (2007). Hebbian learning using fixed weight evolved dynamical 'neural' networks. In H. Abbass, M. Bedau, S. Nolfi, & J. J. Wiles (Eds.), *Proceedings of the First IEEE Symposium on Artificial Life* (pp. 394–401). Honolulu, HI: IEEE.

Keinan, A., Meilijson, I., & Ruppin, E. (2003). Controlled analysis of neurocontrollers with informational lesioning. *Philosophical Transactions of the Royal Society of London A*, *361*, 2123–2144.

Kimura, M., & Nakano, R. (1998). Learning dynamical systems by recurrent neural networks from orbits. *Neural Networks*, *11*, 1589–1599.

Lee, D. N. (1976). A theory of visual control of braking based on information about time- to-collision. *Perception*, *5*, 437–459.

Nolfi, S., & Floreano, D. (2000). *Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines*. Cambridge, MA: MIT Press.

Phattanasri, P., Chiel, H., & Beer, R. D. (2007). The dynamics of associative learning in evolved model circuits. *Adaptive Behavior*, *15*, 377–396.

Piaget, J. (1954). *The construction of reality in the child*. New York: Basic.

Ward, R., & Ward, R. (2008). Selective attention and the control of action: Comparative psychology of an artificial evolved agent and people. *Journal of Experimental Psychology: Human Perception and Performance*, *34*(5), 1165–1182.

Warren, W. H. (2006). The dynamics of perception and action. *Psychological Review*, *113*(2), 358–389.

Williams, P. L., Beer, R. D., & Gasser, M. (2008). An embodied dynamical approach to relational categorization. In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), *Proceedings of the 30th Annual Conference of the Cognitive Science Society* (pp. 223–228).

Wood, R., & Di Paolo, E. (2007). New models for old questions: Evolutionary robotics and the 'A Not B' error. In: F. Almeida e Costa, L. M. Rocha, E. Costa, I. Harvey, & A. Coutinho (eds.) *ECAL 2007. LNCS (LNAI)*, *vol. 4648*, pp. 1141–1150. Heidelberg, Germany: Springer.

Yilmaz, E. H., & Warren, W. H. (1995). Visual control of braking: A test of the τ hypothesis. *Journal of Experimental Psychology: Human Perception and Performance*, *21*, 996–1014.

About the Authors

Randall D Beer received his Ph.D. in computer science in 1989. From 1989 to 2006, he was a professor of electrical engineering and computer science, biology, and cognitive science at Case Western Reserve University. He spent the 1995–1996 academic year as a visiting scientist at the Santa Fe Institute, where he also served as an external faculty member for the next 6 years. In 2006, he joined the Cognitive Science program at Indiana University, where he is currently a professor of Informatics and Computing. Professor Beer's research is broadly concerned with understanding how coordinated behavior arises from the dynamical interaction of an animal's nervous system, its body, and its environment.

Geoffrey P Bingham received his Ph.D. in experimental psychology from the University of Connecticut in 1985, having spent a year 1982–83 studying at Uppsala University supported by a Fulbright Award and another year 1984 as a fellow in the Laboratory for Perceptual Robotics at the University of Massachusetts. From 1985 to 1988, he was an NIH NRSA postdoctoral fellow at Haskins Laboratories in New Haven, CT followed by a fellowship at the Institute for Advanced Study in the Social Sciences in Uppsala and a year as a visiting assistant professor at Trinity College, Hartford, CT. He joined the faculty at Indiana University in 1989, where he is now a professor in Psychological and Brain Sciences and the Cognitive Science Program. Professor Bingham's research addresses human visual and haptic guidance of action (locomotion, reaches-to-grasp, rhythmic coordination, long distance throwing), visual event perception, embodied memory, and 3D shape perception using optic flow and stereo-motion.

Didem Kadıhasanoğlu received her Ph.D. in cognitive science with a minor in cognitive psychology in 2012 from Indiana University, Bloomington, USA. During her Ph.D. studies, she worked under the supervision of Prof. Randall D Beer and Prof. Geoffrey P Bingham. She is now working as an assistant professor at the Department of Psychology at TOBB University of Economics and Technology, Turkey. Her current research focuses on studying human perception and action from a "situated,""embodied" and "dynamical" perspective, with a primary focus on the visual control of locomotion.

Appendix A

Table 2. The mean (with standard deviation), and the maximum final distances and the final velocities of the best agents in Experiments 1 and 2.

		Final distance		Final velocity	
		M (*SD*)	Max	*M* (*SD*)	Max
Experiment 1					
	Image size	2.87 (3.10)	19.10	1.27 (1.32)	3.62
	Image expansion rate	4.08 (2.92)	12.36	1.64 (0.41)	2.32
	Tau	0.14 (0.089)	0.28	0.30 (0.37)	0.98
	Tau-dot	2.41 (2.57)	11.76	0.13 (0.12)	0.35
	Proportional rate	0.03 (0.06)	0.4	0.0	0.0
Experiment 2					
	Image size	36.22 (24.83)	99.20	6.41 (2.54)	11.22
	Image expansion rate	16.91 (10.08)	44.00	4.65 (2.21)	9.19
	Tau	7.73 (5.36)	17.54	4.93 (2.11)	8.52
	Tau-dot	4.35 (2.83)	9.32	4.26 (3.31)	10.31
	Proportional rate	9.54 (6.20)	20.70	2.62 (1.19)	4.52

The mean final distance includes the final distances in successful trials together with premature stops and failure-to-reach trials; the mean final velocity includes the final velocities in crash trials and failure-to-reach trials.