

2009 Special Issue

Facilitating neural dynamics for delay compensation: A road to predictive neural dynamics?

Jaerock Kwon*, Yoonsuck Choe

Department of Computer Science and Engineering, Texas A&M University, College Station, TX 77843-3112, USA

ARTICLE INFO

Article history:

Received 5 March 2009

Received in revised form 17 March 2009

Accepted 21 March 2009

Keywords:

Goal-directed behavior

Prediction

Delay compensation

Facilitating neuronal dynamics

Pole balancing

Neuroevolution

ABSTRACT

Goal-directed behavior is a hallmark of cognition. An important prerequisite to goal-directed behavior is that of prediction. In order to establish a goal and devise a plan, one needs to see into the future and predict possible future events. Our earlier work has suggested that compensation mechanisms for neuronal transmission delay may have led to a preliminary form of prediction. In that work, facilitating neuronal dynamics was found to be effective in overcoming delay (the Facilitating Activation Network model, or FAN). The extrapolative property of the delay compensation mechanism can be considered as prediction for incoming signals (predicting the present based on the past). The previous FAN model turns out to have a limitation especially when longer delay needs to be compensated, which requires higher facilitation rates than FAN's normal range. We derived an improved facilitating dynamics at the neuronal level to overcome this limitation. In this paper, we tested our proposed approach in controllers for 2D pole balancing, where the new approach was shown to perform better than the previous FAN model. Next, we investigated the differential utilization of facilitating dynamics in sensory vs. motor neurons and found that motor neurons utilize the facilitating dynamics more than the sensory neurons. These findings are expected to help us better understand the role of facilitating dynamics in delay compensation, and its potential development into prediction, a necessary condition for goal-directed behavior.

© 2009 Elsevier Ltd. All rights reserved.

1. Introduction

Goal-directed behavior is a hallmark of intelligent cognitive systems. Therefore, understanding such behavior is not only important but also essential to scrutinize intelligence. However, it is not easy to directly investigate goal-directed behavior since there is an implied *agent* behind such behavior, and there is yet not a consensus on what constitutes an agent.

Thus, here we take a different approach to initiate a first step toward understanding goal-directed behavior. Our strategy is to focus on a precondition, or a necessary condition for goal-directed behavior, rather than trying to address the problem head-on.

The main question we will address here is how the precondition could have evolved. Once the prerequisite has evolved, it could have laid a critical stepping stone toward goal-directed behavior. We theorize that one important necessary condition of goal-directed behavior is prediction. Note that a goal is always defined as a future event. Thus, without the ability to anticipate future events, one may not be able to establish a goal. In order to anticipate, one needs to be able to predict. Consequently, by analyzing how prediction has evolved, we could shed light on

a potential evolutionary pathway toward goal-directed behavior. Also, we must note that prediction is increasingly being recognized as one of the core functions of the brain (Hawkins & Blakeslee, 2004; Llinás, 2002) (see also Carpenter and Grossberg (1992) and Li and Kozma (2003) on prediction in dynamic neural network architectures).

In our previous work (Lim & Choe, 2006a, 2006b, 2006c, 2008), we hypothesized that delay in the nervous system could have led to a delay compensation mechanism, which in turn could have further developed into a predictive function. First, let us take a look at neuronal delay in detail before investigating the predictive property of the delay compensation mechanism. Strictly speaking, representations of the present in the brain may not even be precisely aligned with the present in the environment. Our sensory information would reflect the past if the higher perceptual areas in the brain register the signal at the moment the signal is received. Consider visual processing. A series of steps is required for visual stimulus information to reach higher visual processing areas: photoreceptors, bipolar cells, ganglion cells, the lateral geniculate nucleus, the primary visual cortex, and beyond (Nijhawan, 2008). It could take in the range of 100 to 130 ms for the visual signal to arrive in the prefrontal cortex (in monkeys) (Thorpe & Fabre-Thorpe, 2001). In order to make up for the neuronal transmission delay, the brain should utilize information from the past and predict the current state.

* Corresponding author.

E-mail addresses: jrkwon@tamu.edu (J. Kwon), choe@tamu.edu (Y. Choe).




Fig. 1. Delay and delay compensation through facilitating neural activity. (a) The solid curve represents the original signal, and the dotted curve corresponds to the delayed signal (delayed by d). (b) The original signal can be extrapolated by facilitating the neural activity (further increasing when the signal is increasing, and further decreasing when the signal is decreasing). For example, an activation value b at time t (original signal from $t - d$, delayed by d) can be modulated down to a through facilitating dynamics, where the modulated value a is an approximation of the original signal at time t .




Fig. 2. Length of delay and required degree of facilitation for delay compensation. The solid curve represents the original signal, and the dotted curve the delayed signal (delayed by d). (a) Short delay requires only a moderate amount of facilitation to compensate for the delay. (b) More facilitation is needed as the length of delay between the original and the delayed signal becomes greater (the vertical arrows are longer in (b) than in (a)).

Some researchers probed this topic in terms of delay compensation (Lim, 2006; Lim & Choe, 2006b, 2006c) or prediction (Downing, 2005; Krekelberg & Lappe, 2000). Lim and Choe suggested a neural dynamic model for delay compensation using Facilitating Activity Network (FAN) based on short-term plasticity in the neuron known as *facilitating synapses* (Lim, 2006; Lim & Choe, 2006b). Facilitating synapses have been found at a single neuron level in which the membrane potential shows a dynamic sensitivity to the changing rate of the input (Liaw & Berger, 1999; Lim, 2006). As illustrated in Fig. 1, the original signal can be recovered from the delayed signal by using facilitating dynamics. According to the facilitation model, as Fig. 2 illustrates, higher facilitation rates are needed to effectively deal with longer delay. However, the FAN model turns out to have limitations, i.e., oscillation under high facilitation rate (see Section 3 for details). Furthermore, the analysis in Lim (2006) and Lim and Choe (2006c) did not consider differential utilization of facilitation among different neuron types within the context of the entire network (e.g., sensory neurons vs. motor neurons).

Here, we propose an improved dynamic model, *Neuronal Dynamics using Previous Immediate Activation value* (NDPIA) that solves the oscillation problem in FAN. In addition, we conducted experiments in less restricted conditions than in Lim (2006) and Lim and Choe (2006c): (1) input delay was applied to the system for the entire duration of each experiment, and (2) we extended the delay to twice the usual value compared to the earlier experiments with FAN, and analyzed the results from the increased delay.

To test NDPIA and to investigate the properties of the neuronal networks with the suggested neuronal dynamics, we employed a two degree-of-freedom (2D) pole-balancing (Gomez & Miikkulainen, 1998) agents with evolved recurrent neural networks as their controllers (cf. Gomez and Miikkulainen (2003) and Gomez (2003)). We used conventional neuroevolution to train the networks (see Section 3.2 for detailed justification, and R.G. Ward and R. Ward (2009) for successful use of such strategy in a different task domain).

Our main findings are as follows: (1) NDPIA can solve the oscillation problem in FAN during heightened facilitation. (2) Motor neurons in a NDPIA network tend to evolve high facilitation rates, confirming similar previous results with FAN. (3) Longer delay leads to higher facilitation rates. (4) Neural network controllers using NDPIA dynamics result in better performance in pole-balancing tasks than those based on FAN. (5) NDPIA networks show robust performance under extremely high facilitation rates, especially when only the motor neurons are facilitated. These results suggest that delay and facilitation rate must be positively correlated for effective compensation of delay, and the best part in the system to introduce such dynamics is the motor system.

Below, we first look into related research, then we analyze the limitations in the FAN dynamics. Then we will propose a new facilitating dynamics (NDPIA). Next, the 2D pole-balancing problem and evolutionary neural networks will be introduced. Finally we will present and analyze the results, followed by discussion and conclusion.




Fig. 3. Facilitating Neural Activity. (a) The immediate activation value $X(t)$ is modulated by the difference between $X(t)$ and the modulated activation value $A(t - 1)$ in the previous time step, with facilitation rate r . (b) The same principle can be applied to the decreasing activation case.




Fig. 4. Problems in facilitating dynamics of FAN. (a) When the activity is increasing, the immediate activation value $X(t)$ could be smaller than the modulated value $A(t - 1)$ from the previous time step, so the modulated value at the present $A(t)$ becomes smaller than the immediate value $X(t)$. This property of the conventional FAN model makes the system output to become unstable. (b) Basically, the same analysis can be applied in the case of decreasing activity. When the activity is decreasing, $X(t)$ is larger than $A(t - 1)$. Hence, $A(t)$ becomes even larger than $X(t)$.

2. Background

The activation level or the membrane potential of the postsynaptic neuron is modulated by the change in the rate of past activation. These dynamic synapses generate short-term plasticity, which shows activity-dependent decrease (depression) or increase (facilitation) in synaptic transmission (Fortune & Rose, 2001; Liaw & Berger, 1999). These activities occur within several hundred milliseconds from the onset of the stimulus (Liaw & Berger, 1999; Markram, 2003). Lim (2006), Lim and Choe (2006b, 2006c) investigated the relationship between these neuronal dynamics and delay compensation, and suggested that facilitating dynamics at a single neuron level may play an important role in the compensation of neuronal transmission delay.

How can such dynamics be realized in a neural network? We can begin with conventional artificial neural networks (ANNs), but ANNs lack such single neuron-level dynamics (note the adding recurrent connections can introduce a network-level dynamics). As we can see in Eq. (1), the activation values in conventional ANNs are determined by the instantaneous input value and the connection weights.

$$X(t) = g \left(\sum_{j=1}^m w_j X_j(t) \right) \quad (1)$$

where $g(\cdot)$ is a nonlinear activation function such as the sigmoid function, m is the number of neurons of the preceding layer, w_j is the connection weight, and X_j is an activation value from a neuron of the preceding layer (Lim, 2006; Lim & Choe, 2006b, 2006c). Eq. (1) shows that there is no room to consider the past values of X_j . Recurrent ANNs could be one simple solution for this, but the dynamics may not be fast enough to cope with input delays. Tan

and Cauwenbergh (1999) proposed a neural network based Smith predictor to compensate for large time delay; Miall and Wolpert (1996) used the Kalman filter in the internal forward model to predict the next state; and Lim and Choe (2006c) showed that facilitating neuronal dynamics at a single neuron level can play an important role in compensating for input delays.

In order to overcome the issues above, the activation value needs to be directly modulated as in the Facilitating Activity Network (FAN) model (Lim, 2006; Lim & Choe, 2006b, 2006c):

$$A(t) = X(t) + r \Delta(t) \quad (2)$$

where $A(t)$ is the modulated (facilitated or depressed) activation value at time t , $X(t)$ is the immediate activation value, r is a dynamic rate ($-1 \leq r \leq 1$), and $\Delta(t)$ is $X(t) - A(t - 1)$.

If $r \geq 0$, and if the signal increases for a while, the activation value is augmented by the difference $\Delta(t)$ of the immediate activation value $X(t)$ and the previous modulated activation $A(t - 1)$ with the rate r (see Fig. 3(a)). If $r \geq 0$, but if the signal decreases, the activation value is diminished by $\Delta(t)$, because it becomes a negative value in this case as shown in Fig. 3(b). This results in facilitation.

Suppose $r \leq 0$, and that the signal increases for a while, then the activation value is diminished by the difference $\Delta(t)$ between the immediate activation value and the previous modulated activation with the rate r . If the signal decreases for a while under the same condition, the amount of decrease becomes smaller than the immediate value by $\Delta(t)$ with the rate r , because r is a negative value and $\Delta(t)$ is a negative value as well, so $r\Delta(t)$ becomes a positive value. This makes the signal greater than the immediate signal. Furthermore it means that the signal is decreased less than what it is supposed to be. In other words, the modulated activation values can be considered within the range of $(X(t) - \Delta(t)) \leq A(t) \leq (X(t) + \Delta(t))$ (Lim, 2006) which means that the




Fig. 5. Proposed facilitation neural activity. (a) By using the previous immediate value $X(t - 1)$ instead of the modulated value $A(t - 1)$, the present modulated value is stabilized. (b) The same change is applied to the decreasing activation case.




Fig. 6. The proposed facilitation model solving the problem in FAN. (a) The modulated activation value $A(t)$ is guaranteed to be larger than the immediate activation value $X(t)$ as long as activation increases. (b) $A(t)$ is guaranteed to be smaller than $X(t)$ as long as activation decreases. Compare to Fig. 4.

present activation value could be diminished by $\Delta(t)$ (depressing dynamics) or augmented by $\Delta(t)$ (facilitating dynamics).

Neural networks using the FAN model showed not only better performance than conventional networks but they were also more robust in various delay conditions (Lim, 2006; Lim & Choe, 2006a).

3. Methods: Enhanced facilitating activity model

Even though Lim (2006), Lim and Choe (2006b, 2006c) paid attention to short-term synaptic plasticity, especially facilitating synapses, and suggested a compensation mechanism for neuronal transmission delay, several further challenges remain. As Fig. 2 illustrates, we need to use a higher facilitation rate as the delay increases. However, Lim and Choe did not investigate the effect of higher facilitating rates. When the FAN model is used with high facilitation rates, the modulated activation values become unstable/oscillatory. Furthermore, a systematic analysis of the neuronal dynamics in the network level is needed because Lim and Choe did not investigate differential utilization of facilitating dynamics dependent on neuron type.

Here, we propose an improved dynamics model to address these challenges. The previous FAN model turns out to have a limitation especially when longer delay is applied to the model. Below, we analyze the potential problems of the FAN model in detail and propose an enhanced model to deal with the problems. First, we expand Eq. (2) into:

$$A(t) = \left(\sum_{n=0}^{k-1} (-1)^n r^n (1+r) X(t-n) \right) + (-1)^k r^k A(t-k). \quad (3)$$

Now we can more clearly see that the current modulated activation value $A(t)$ is a function of $X(t - 1)$, $X(t - 2)$, $X(t - 3)$ and so on. The problem is that, given a positive dynamic rate r ,

$X(t - 1)$, $X(t - 3)$, etc. contribute negatively while $X(t - 2)$, $X(t - 4)$, etc. positively. These positive and negative components can give rise to abrupt oscillations in $A(t)$ that originally do not exist in the input signal.

To better illustrate the problem, let us take an example in the case of facilitating dynamics. As we can see in Fig. 4(a), even when $X(t)$ keeps increasing from $X(t - 1)$, the immediate activation value $X(t)$ could be smaller than the previous modulated value $A(t - 1)$. This is not desirable since $A(\cdot)$ will oscillate unlike $X(\cdot)$. The same phenomenon happens when the activity is decreasing as in Fig. 4(b).

3.1. Enhanced facilitating activity model

In order to address the above issue, we propose an improved neuronal dynamics model (NDPIA) which considers only the previous immediate activation value (see Figs. 5 and 6).

$$A(t) = X(t) + r(X(t) - X(t - 1)) \quad (4)$$

where $A(t)$ is the modulated (facilitated or depressed) activation value at time t , $X(t)$ is the immediate activation value, and r is the dynamic rate. The dynamic rate r can either facilitate or depress the activity, and it is not limited to $-1 \leq r \leq 1$, so that we can either facilitate or depress the immediate activation values as highly as we want. But practically, this value should not be too high.

As we have shown in Eq. (3), the effect of $X(t - (n + 1))$ disappears very quickly as n increases and r is less than 1. NDPIA accounts for the current and the previous immediate activation values. So in order to consider the previous activation values such as $A(t - 1)$ and $A(t - 2)$ prior to the immediate one, we used recurrent neural networks in the present paper, and the context inputs that are simply feedback from the hidden layer could make up for the effect of older past activation values.




Fig. 7. Two degree-of-freedom pole-balancing task. The cart (gray disc) with an upright pole attached to it must move around on a 2D plane while keeping the pole balanced upright. The cart controller receives the location (x, y) , the pole angle (θ_x, θ_y) , and their respective velocities as the input, and generates the force in the x and the y direction. Note that x and y are projected angles to the x - z and the y - z plane respectively.

3.2. 2D pole-balancing problem with delayed inputs

We tested our new facilitating dynamics in recurrent neural network controller for a 2D pole-balancing task (Fig. 7). The state of the cart (the gray disc on the bottom Fig. 7) with a pole on top is characterized by the following physical parameters: The cart position in the plane (x, y) , the velocity of the cart (\dot{x}, \dot{y}) , the angle of the pole from the vertical in the x and the y directions (θ_x, θ_y) , and their angular velocities $(\dot{\theta}_x, \dot{\theta}_y)$ (Gomez & Mäikkulainen, 1998).

To test our facilitation dynamics, we employed the 2D pole-balancing problem, following Lim and Choe (2006c). The differences from Lim and Choe (2006c) are as follows. First, we tested with delays in all inputs, and the delay was applied during the entire test period in all experiments. In Lim and Choe (2006c), delay was applied either to a subset of the input for the entire duration, or to all the inputs only for a limited time period during each trial. Second, we evolved the controllers under no delay condition and tested them with up to two-step delay with increased facilitation rate. Longer delay may not be acceptable because of the high possibility of phase difference (see Section 5 for a discussion). In Lim and Choe (2006c) only one-step delay was investigated for measuring the performance of controller networks. Third, we used conventional GA to evolve the controllers instead of Enforced SubPopulation algorithm (ESP) (Gomez, 2003; Gomez & Mäikkulainen, 1998). The main reason for using conventional GA was to have a clearly separated role for the sensory and the motor neurons, to investigate the differential utilization of facilitating dynamics in these neuron types.

The cart controller applies force to the cart on a flat surface to balance the pole (the pole must remain within $\pm 15^\circ$ of the vertical). The force was applied in both the x and the y directions at a 0.1 s interval. If the controller balances the pole more than 5000 steps (1 step = 100 ms), we consider it as a success. The fitness function returned the number of steps the agent balanced the pole within $\pm 15^\circ$ from the vertical and stayed inside a $3 \text{ m} \times 3 \text{ m}$ area (each axis ranging from -1.5 m to 1.5 m). We used recurrent neural networks to control the cart (Fig. 8). See Section 3.3 for details on the neural network controller. Fifty recurrent networks were evaluated in each generation, and to avoid situations where some neurons evolve to have accidentally good fitness values, we used the roulette wheel sampling method (Buckland, 2002; Ghanea-Hercock, 2003; Yao, 1999). We used a pole length of 0.5 m tilted 1° from the vertical towards the $+y$ direction with $(\theta_x, \theta_y) = (0, 0)$ in the initial state. Force within the range of -10 N to 10 N was applied to the cart at a time step of 0.1 s, based on the output (F_x, F_y) .




Fig. 8. Recurrent neural network for 2D pole balancing. The signal flow inside each neuron is shown in the box. Z^{-1} means unit delay.

With this setup, we (1) investigated the effect of dynamic rates in a single neuron level by evolving the rates from depressing to facilitating property, (2) compared the performance between FAN and NDPIA, and (3) showed that facilitating motor neurons are better at coping with longer delays than facilitating sensory inputs or both sensory and motor neurons.

3.3. Neuroevolution

Lim (2006) and Lim and Choe (2006c) investigated dynamic activation rates in a single neuron level, however, they have not tested the effect of facilitation in different parts of the network (e.g., by neuron type).

We evolved controllers having recurrent neural networks with dynamics neuronal activities. These activity rates were evolved to range across $-1 \leq r \leq 1$ which means it could be facilitating or depressing.

We used a recurrent neural network with eight input nodes, three context input nodes, three hidden neurons, and two output neurons in order to control the cart in the plane ($3 \text{ m} \times 3 \text{ m}$). Fig. 8 shows the recurrent network that we used in the experiments. Input nodes correspond to the cart position (x, y) , the velocity of the cart (\dot{x}, \dot{y}) , the angle of the pole from the vertical in the x and the y directions (θ_x, θ_y) , and their angular velocities $(\dot{\theta}_x, \dot{\theta}_y)$. The hidden layer activations are fed back as contextual input, with a unit delay. Output neurons F_x and F_y represent the force in the x and the y direction, respectively. Each neuron's immediate activity is calculated by Eq. (1), and subsequently facilitated using the dynamics in Eq. (4).

In training these nonlinear controllers, neuroevolution methods proved efficient (Gomez, 2003; Gomez & Mäikkulainen, 2003). Unlike Gomez (2003) and Gomez and Mäikkulainen (2003), we used a conventional neuroevolution method instead of ESP. The chromosome encoded the connection weights between input nodes and hidden layer neurons, and between hidden layer neurons and output neurons. In the experiment of the evolution of dynamic activation rates, we additionally included a dynamic rate parameter in the chromosome. Crossover occurred with probability 0.7 and the chromosome was mutated by ± 0.3 (perturbation rate) with probability 0.2. These parameters were determined empirically.

4. Experiments and results

First, we tested whether NDPIA helps fix the unstable-ness/oscillation problem in FAN, using a fixed time series as input. Next, to test the rest, we evolved recurrent neural networks




Fig. 9. Facilitation under fast signal change condition. This graph shows a small portion of the function $f(t) = 2 \times \exp(-t) \times \sin(t)$ in the interval $[0..20]$ (dotted line, simulating a delayed signal). The x axis represents time t and the y axis the activation value $f(t)$. Facilitation rate of $r = 0.8$ was used in this example. (a) When the signal changes quickly (increasing leg, from time 0 to 7), the FAN results in jagged oscillation. Note that when the signal change is slow relative to the facilitation rate (decreasing leg, from time 7 to 20), the oscillation disappears. (b) The proposed method eliminates the oscillation problem in (a). Note that due to the facilitation, the resulting curve (solid line) appears shifted to the left (i.e., we can say that delay was compensated).




Fig. 10. Facilitation under slow signal change condition. The same function as in Fig. 9 is shown, but in a different interval $[20:40]$. A higher facilitation rate of $r = 0.9$ was used, to demonstrate the oscillation problem in FAN even under slow signal change conditions. (a) Even when the signal changes slowly, if r is high, FAN results in oscillatory activation. (b) The proposed dynamics again eliminates the oscillation problem, with the same delay compensation property as in Fig. 9.

with FAN or NDPIA dynamics, where the connection weights and also the dynamic rates were allowed to evolve. The networks were trained in a 2D pole-balancing task, and then tested with added delay in the sensory signals (the input).

4.1. Enhanced facilitating activity model

With NDPIA, we were able to correct the oscillation problem in FAN, discussed in Section 2. We compared the two models with a signal taking a simple functional form: $f(t) = 2 \times \exp(-t) \times \sin(t)$ (t is time). Figs. 9 and 10 show portions of the function where the change in the signal is either fast or slow. First we observed the part of the signal where the signal changes rapidly. Here we used a dynamic activation rate of 0.8 (facilitation). In Fig. 9(a), which shows FAN, the immediate activation value cannot keep up with the modulated one, thus oscillation occurs in the modulated activation values. In other words, facilitation did not occur properly in this case. Fig. 9(b) shows that using NDPIA these oscillations can be removed. Even when the signal changes slowly, if the facilitation rate is high enough (0.9 was used in this case), the oscillation would occur again (Fig. 10). Facilitation rate of 0.9 might seem too high, but as we examined in Section 1, high facilitation rates can be necessary. As before, FAN results in oscillation (Fig. 10(a)) while NDPIA results in no oscillation (Fig. 10(b)).

4.2. Evolved dynamic activation rates

In the recurrent neural network controller, the hidden units receive direct sensory input, so we can say these are sensory neurons. In a similar manner we can consider the output neurons as motor neurons since they are directly coupled to the cart motion. Our main question here is if all types of neurons (sensory or motor) evolve to utilize facilitating dynamics under delayed input conditions. Furthermore, we question if increase in input delay leads to stronger facilitating dynamics.

In order to test these, we encoded into the chromosome the dynamic activation rates of sensory neurons (hidden) and motor neurons (output) as well as the synaptic weights in the connections in the neural network controller.

Fig. 11 shows the distribution of evolved dynamic rates of top 5 individuals from 25 separate populations (each population had 50 individuals), under three different delay conditions (0, 1, and 2). The motor neurons exhibit higher utilization of facilitating dynamics (high dynamic rate) compared to sensory neurons, when the delay is high (Fig. 11(c)). The cumulative distribution of the dynamic rate shows more clearly the positive correlation between increasing input delay and higher dynamic rate in motor neurons (Fig. 11(e)) but not in sensory neurons (Fig. 11(d)). In sum, motor neurons are more likely to be facilitated, and increasing input delay leads to higher facilitation.




Fig. 11. Dynamic rate distribution. The distribution of dynamic rates for the controllers trained with different delay conditions are shown (top 5 individuals from the last generation in 25 separate evolutionary trials). (a–c) shows the dynamic rate distribution under different delay conditions. As the delay increases, motor neurons increasingly utilizes higher dynamic rate (i.e., facilitation). (d–e) shows the cumulative distribution, directly comparing the different delay conditions. Only for the motor neurons (e), increasing delay shifts the cumulative distribution toward the right (i.e., at higher facilitating rates the mean increases from -0.196 to -0.184 to 0.724).

4.3. Pole-balancing performance in FAN and NDPIA

In this experiment, we compared the effectiveness of FAN and NDPIA in the 2D pole-balancing task, under various delay conditions.

To test the effectiveness of facilitation under a strictly controlled environment, we set the facilitation rate to a fixed value of 0.7, with different types of neurons being facilitated in three different sets of experiments. We trained (evolved) 60 FAN networks and 60 NDPIA networks under no delay. Among the 60 networks, 20 were evolved with facilitated motor neurons, 20 networks were evolved with facilitated sensory neurons, and 20 remaining networks were evolved with facilitation in both sensory neurons and motor neurons for each facilitation model (FAN and NDPIA respectively).

For testing, we put the evolved networks under no delay, one-step delay, and two-step delay environment to see the effect of dynamic activation rates. Fig. 12 shows that NDPIA has better performance than FAN in most cases, especially under longer delays. Note that if both sensor and motor neurons are facilitated at the same time, no significant difference is found (Fig. 12(c)). There was not much difference when there was no delay in the input (see delay 0 cases in Fig. 12), but the difference of performance becomes clear when motor neurons were facilitated as delay increases (see Fig. 12(b)).

4.4. Pole-balancing performance under extremely high facilitation rates in different neuron types

In this experiment, we investigated the effect of extremely high dynamic activation rates under long delay. Controller networks with NDPIA maintained their performance under longer delay, with a fairly high facilitation rate of 0.7, especially for the motor neuron only facilitation (Fig. 12). How would the performance change if we push the facilitation rate to an even higher value? We

tested how extreme facilitation like that affects performance when different types of neurons are facilitated: sensory, motor, or both. This is an interesting question since longer delay might necessitate higher facilitation rate.

When either the sensory neurons or the motor neurons were facilitated with a high facilitation rate of 0.7, the performance remained high (Fig. 13 (a) and (b)), but the performance degraded when both neuron types were facilitated at that rate (Fig. 13 (c)). As facilitation rate was further increased to even higher values (1.2 to 1.5 to 2.0), performance started to degrade for the case where sensory neurons were facilitated (Fig. 13(a)), but it was not the case when only motor neurons were facilitated (Fig. 13(b)). The case with both sensory and motor neurons facilitated showed consistently low performance regardless of the facilitation rate. These results suggest that motor neurons could be the best type to facilitate at higher rates, for the compensation of longer delays.

5. Discussion

The main contribution of our work is to have shown the link between facilitating neuronal dynamics and delay compensation in a systematic study. In particular we have shown that facilitation is more effective in motor neurons, and that longer input delay leads to higher rates of facilitation. We have also improved the previous facilitation model (FAN) (Lim, 2006; Lim & Choe, 2006b, 2006c) so that higher facilitation rates can be used without side effects (oscillation). As a consequence, our new approach allowed our model to deal with longer delay applied over the entire duration of each trial.

One of the main results of our investigation was that facilitating dynamics is more effective in counteracting delay in certain classes of neurons (i.e., motor neurons). This could be due to two high-level reasons: (1) local connection topology, and (2) series of delay compensation happening in a chain of neurons, and both could cause over-compensation. First, if the local topology of the neuron has a recurrent link, then facilitation could be amplified,




Fig. 12. Comparison of FAN and NDPIA under different delay conditions and for different types of neurons facilitated. (a) Top five individuals from the final generation of those with facilitated sensory neurons are shown, for delay 0 (top) to delay 2 (bottom). The number of pole-balancing steps in FAN (*) and NDPIA (°) are plotted. NDPIA shows a slight advantage under longer delay conditions (bottom row). (b) The same information is plotted as in (a), for top 25% individuals with facilitated motor neurons. NDPIA has better performance than FAN under longer delays (middle and bottom rows). (c) The same information is plotted as in (a), for top 25% individuals with facilitated sensory and motor neurons. There is no significant difference between FAN and NDPIA.




Fig. 13. The performance of NDPIA in two-step delay under high facilitation rates. Performance of top 5 controllers under extremely high facilitation rates are shown for facilitated (a) sensory neurons, (b) motor neurons, and (c) sensory and motor neurons. With a facilitation rate of 0.7, there is minimal degradation of performance in (a) and (b), but much degradation in (c). Also, (b) is the only case where very high facilitation rate (=2.0) does not affect the performance. Thus, when very high facilitation is needed (e.g., to compensate for longer delays), motor neurons should be facilitated.

thus leading to over-compensation. This is why we believe the sensory neurons did not do well compared to the motor neurons (recall that the sensory neurons are the hidden units that receive not only the input but also the context input from the recurrent connections). Second, if a chain of neurons originating from the input and ending in the output are all facilitated, again it could lead to over-compensation. This could be the reason why the networks with both sensory and motor neurons facilitated performed poorly. These observations can lead to concrete predictions that can be experimentally validated: (1) Facilitating neural dynamics, when

employed to perform delay compensation, may be more prevalent where the local connection topology is feedforward. (2) Within a sensorimotor processing chain, facilitating dynamics may be found in only few parts of the chain.

As we have seen already in Fig. 11, in order to compensate for longer delay, higher facilitation should be used. However, since higher facilitation in longer delay may cause higher error rates, this facilitation dynamics is applicable only when the delay is within a certain bound. In other words, if the signal changes more rapidly than the delay duration (or equivalently if the delay duration is

longer than the time scale of signal change), our approach may not work well. In order to deal with such situations, sensorimotor anticipation (Gross, Heinze, Seiler, & Stephan, 1999) or sensory prediction (Downing, 2005; Gross et al., 1999) may be needed. To address the challenge of longer delay, internal representations, internal models, or forward models (Oztopa, Wolpert, & Kawato, 2005; Webb, 2004) can also be used. These works suggest the use of anticipation for the future inputs or states in a higher level than at the level of neuronal circuits.

In the beginning, we started out with the insight that predictive function could have originated from delay compensation mechanisms. What our results show is that an intimate relationship exists between longer delay and higher facilitatory dynamics (i.e., the compensation mechanism), and that such compensation mechanisms can emerge during the process of evolution. Since the final outcome of delay compensation is the estimation of the present state (based on information from the past), one might argue that it is not prediction. However, the task itself can be systematically mapped to that of prediction, since it all boils down to the estimation of the state in the relative future, whether that future is now or whenever (see Werbos (2009)). An interesting future direction is to see if actual predictive capability can evolve in a similar simulated evolution environment, when the task itself requires prediction, rather than just delay compensation. We expect facilitating dynamics to again play an important role in such a case.

Another matter of debate is the biological significance of the proposed facilitating dynamics. As we briefly mentioned in the beginning, part of the motivation for this work came from the facilitating synapses reported in the experimental literature (Markram, Wang, & Tsodyks, 1998), which provides the feasibility. Also, in return, our proposed mechanism assigns a specific role regarding the function of such synapses. In a similar line, we can ask whether the proposed method is biologically feasible, especially within an individual's lifetime. Although our experiments were done using simulated evolution, it could be seen as just another optimization process, so there is no reason why such a delay compensation mechanism can be developed over time within a single individual's lifetime.

Finally, we would like to take the discussion further and speculate on the role of prediction in brain function. Prediction is receiving increasing attention as a central function of the brain (Hawkins & Blakeslee, 2004; Llinás, 2002) (also see Werbos (2009)). The brain is fundamentally a dynamical system, and understanding the dynamics can lead to deep insights into the mechanisms of the brain. For example, according to Kozma and Freeman (2009), the brain state trajectory transitions back and forth between chaotic high-dimensional attractors to periodic low-dimensional attractors. An interesting property of these two different attractors is that for the chaotic attractor, predicting the future state in the state trajectory may be difficult compared to that of the periodic attractor. Such predictive function could form an important necessary condition for more complex and sometimes subjective phenomena as consciousness or self-awareness (Kwon & Choe, 2008). It could also be thought of as the “unconscious” process discussed in Gazzaniga (1998) that drives brain function, that is later confirmed or described by consciousness, post hoc. Prediction can also be useful in other ways, including predicting the upcoming input, based on the current model of the world (Perlovsky, 2009). This kind of mechanism, coupled with emotional circuits (e.g., Levine (2009)), could serve as a fundamental component in goal-directed behavior. In sum, how such humble delay compensation mechanisms as presented in this paper can develop into a fully functioning predictive system, laying the foundation for high-level cognitive processes, is an important future question to be addressed.

6. Conclusions

In this paper, we proposed an improved facilitating dynamics, NDPIA, to address shortcomings in the previous FAN model. We showed that our approach overcomes the limitations and results in higher performance in a standard 2D pole-balancing task, with longer delay in the input during the entire duration of the task. More importantly, we have found that facilitating dynamics is the most effective in motor neurons, and increasing input delay leads to higher utilization of facilitating dynamics. Our findings are expected to help us better understand the role of facilitating dynamics in delay compensation, and its potential development into prediction, a necessary condition for goal-directed behavior.

Acknowledgments

This is an expanded version of a preliminary report that was presented by Kwon and Choe (2007) at the INNS Texas/MIND/UTA Conference on Goal-Directed Neural Systems (2006). We would like to thank the three anonymous reviewers for their constructive comments.

References

- Buckland, M. (2002). *AI techniques for game programming*. Premier press.
- Carpenter, G. A., & Grossberg, S. (1992). A self-organizing neural network for supervised learning, recognition, and prediction. *IEEE Communications Magazine*, 30, 38–49.
- Downing, K. L. (2005). The predictive basis of situated and embodied artificial intelligence. In *GECCO'05: Proceedings of the 2005 conference on Genetic and evolutionary computation* (pp. 43–50). New York, NY, USA: ACM Press.
- Fortune, E. S., & Rose, G. J. (2001). Short-term synaptic plasticity as a temporal filter. *Trends in Neurosciences*, 24, 381–385.
- Gazzaniga, M. S. (1998). *The mind's past*. University of California Press.
- Ghanea-Hercock, R. (2003). *Applied evolutionary algorithms in Java*. Springer.
- Gomez, F. (2003). Robust non-linear control through neuroevolution. *Ph.D. thesis, Technical Report AI03-303*. Austin, TX: Department of Computer Science, The University of Texas at Austin.
- Gomez, F., & Miikkulainen, R. (1998). 2-D pole balancing with recurrent evolutionary networks. In *Proceeding of the international conference on artificial neural networks* (pp. 758–763).
- Gomez, F., & Miikkulainen, R. (2003). Active guidance for a finless rocket through neuroevolution. In *Proceedings of the genetic and evolutionary computation conference* (pp. 2084–2095).
- Gross, H.-M., Heinze, A., Seiler, T., & Stephan, V. (1999). Generative character of perception: A neural architecture for sensorimotor anticipation. *Neural Networks*, 12, 1101–1129.
- Hawkins, J., & Blakeslee, S. (2004). *On intelligence*. Times Books.
- Kozma, R., & Freeman, W. J. (2009). The KIV model of intentional dynamics and decision making. *Neural Networks*, 22(3), 277–285.
- Krekelberg, B., & Lappe, M. (2000). A model of the perceived relative positions of moving objects based upon a slow averaging process. *Vision Research*, 40, 201–215.
- Kwon, J., & Choe, Y. (2007). Enhanced facilitatory neuronal dynamics for delay compensation. In *Proceedings of the international joint conference on neural networks* (pp. 2040–2045). Piscataway, NJ: IEEE Press.
- Kwon, J., & Choe, Y. (2008). Internal state predictability as an evolutionary precursor of self-awareness and agency. In *Proceedings of the seventh international conference on development and learning* (pp. 109–114). IEEE, URL kwon.icdl08.pdf.
- Levine, D. S. (2009). Brain pathways for cognitive-emotional decision making in the human animal. *Neural Networks*, 22(3), 286–293.
- Li, H., & Kozma, R. (2003). A dynamic neural network method for time series prediction using the kiii model. In *Proceedings of the international joint conference on neural networks* (pp. 347–352).
- Liaw, J., & Berger, T. W. (1999). Dynamic synapse: Harnessing the computing power of synaptic dynamics. *Neurocomputing*, 26–27, 199–206.
- Lim, H. (2006). Facilitatory neural dynamics for predictive extrapolation. *Ph.D. thesis*. Texas A&M University.
- Lim, H., & Choe, Y. (2006a). Compensating for neural transmission delays using extrapolatory neural activation in evolutionary neural networks. *Neural Information Processing—Letters and Reviews*, 10, 147–161.
- Lim, H., & Choe, Y. (2006b). Delay compensation through facilitating synapses and stdp: A neural basis for orientation flash-lag effect. In *Proceedings of the international joint conference on neural networks* (pp. 8385–8392).
- Lim, H., & Choe, Y. (2006c). Facilitating neural dynamics for delay compensation and prediction in evolutionary neural networks. In *Genetic and evolutionary computation conference* (pp. 167–174). Nominated for Best Paper Award in the Artificial Life/Evolutionary Robotics/Adaptive Behavior track.

- Lim, H., & Choe, Y. (2008). Extrapolative delay compensation through facilitating synapses and its relation to the flash-lag effect. *IEEE Transactions on Neural Networks*, 19(10), 1678–1688.
- Llinás, R. R. (2002). *I of the vortex: From neurons to self*. MIT Press.
- Markram, H. (2003). *Elementary principles of nonlinear synaptic transmission*. London, UK: Springer-Verlag.
- Markram, H., Wang, Y., & Tsodyks, M. (1998). Differential signaling via the same axon of neocortical pyramidal neurons. *Proceedings of the National Academy of Sciences*, 95, 5323–5328.
- Miall, R. C., & Wolpert, D. M. (1996). Forward models for physiological motor control. *Neural Networks*, 9, 1265–1285.
- Nijhawan, R. (2008). Visual prediction: Psychophysics and neurophysiology of compensation for time delays. *Behavioral and Brain Sciences*, 31, 179–239.
- Oztopa, E., Wolpert, D., & Kawato, M. (2005). Mental state inference using visual control parameters. *Cognitive Brain Research*, 22, 129–151.
- Perlovsky, L. (2009). Language and cognition. *Neural Networks*, 22(3), 247–257.
- Tan, Y., & Cauwenbergh, A. V. (1999). Neural-network-based d-step-ahead predictors for nonlinear systems with time delay. *Engineering Applications of Artificial Intelligence*, 12, 21–35.
- Thorpe, S. J., & Fabre-Thorpe, M. (2001). Seeking categories in the brain. *Science*, 291, 260–263.
- Ward, R. G., & Ward, R. (2009). Representation in dynamical agents. *Neural Networks*, 22(3), 258–266.
- Webb, B. (2004). Neural mechanisms for prediction: Do insects have forward models? *Trends in Neurosciences*, 27, 278–282.
- Werbos, P. (2009). Intelligence in the brain: A theory of how it works and how to build it. *Neural Networks*, 22(3), 200–212.
- Yao, X. (1999). *Evolutionary computation theory and applications*. World scientific.