SYNAPSE 61:375–390 (2007)

Spiking Neurons, Dopamine, and Plasticity: Timing Is Everything, but Concentration Also Matters

JEAN-PHILIPPE THIVIERGE,[1]* FRANÇOIS RIVEST,[2] AND OURY MONCHI[3]
[1]*Département de Physiologie, Université de Montréal, Québec, Canada*
[2]*Département d'informatique et de recherche opérationnelle, Université de Montréal, Québec, Canada*
[3]*Centre de Recherche de l'Institut Universitaire de Gériatrie, Université de Montréal, Québec, Canada*

KEY WORDS spike-timing-dependent plasticity; corticostriatal synapses; dopamine

ABSTRACT While both dopamine (DA) fluctuations and spike-timing-dependent plasticity (STDP) are known to influence long-term corticostriatal plasticity, little attention has been devoted to the interaction between these two fundamental mechanisms. Here, a theoretical framework is proposed to account for experimental results specifying the role of presynaptic activation, postsynaptic activation, and concentrations of extracellular DA in synaptic plasticity. Our starting point was an explicitly-implemented multiplicative rule linking STDP to Michaelis-Menton equations that models the dynamics of extracellular DA fluctuations. This rule captures a wide range of results on conditions leading to long-term potentiation and depression in simulations that manipulate the frequency of induced corticostriatal stimulation and DA release. A well-documented biphasic function relating DA concentrations to synaptic plasticity emerges naturally from simulations involving a multiplicative rule linking DA and neural activity. This biphasic function is found consistently across different neural coding schemes employed (voltage-based vs. spike-based models). By comparison, an additive rule fails to capture these results. The proposed framework is the first to generate testable predictions on the dual influence of DA concentrations and STDP on long-term plasticity, suggesting a way in which the biphasic influence of DA concentrations can modulate the direction and magnitude of change induced by STDP, and raising the possibility that DA concentrations may inverse the LTP/LTD components of the STDP rule. **Synapse 61:375–390, 2007.**

INTRODUCTION

Over the past decade, significant advances have been made in understanding the mechanisms underlying long-term potentiation (LTP) and depression (LTD) in corticostriatal synapses (Gurney et al., 2004; Mahon et al., 2004). Recent evidence suggests that the specific spike timing of corticostriatal activation can influence both the magnitude and direction of change in synaptic efficacy (Fino et al., 2005). In addition, long-term plasticity in corticostriatal synapses is also mediated by interactions with nigrostriatal cells, in particular through the levels of extracellular dopamine (DA) released around the precise time of corticostriatal activation (Reynolds and Wickens, 2002). However, at the current time, both empirical and theoretical investigations into the possible interactions of spike-based and neurotransmitter-based mechanisms are scarce.

In order to address this issue, the current paper presents a theoretical framework that captures synaptic plasticity in corticostriatal synapses. This framework includes two separate models. First, a model of spike-timing-dependent plasticity (STDP; Abbott and Nelson, 2000; Gerstner and Kistler, 2002; Legenstein et al., 2005; for a review see Dan and Poo, 2004) captures the influence of precise spike timing on synaptic plasticity. Second, a model of the temporal dynamics of DA release is implemented through Michaelis-Menton equations (MM; Montague et al., 2004a; Nicholson, 1995, 2001; Schonfub, 2001; Venton et al., 2003).

To capture the complex synaptic interactions between STDP and the dynamics of extrasynaptic DA, we take as starting point an explicitly-implemented multiplicative rule, such as that employed in temporal-

Contract grant sponsor: Fonds de Recherche en Santé du Québec (FRSQ).

*Correspondence to: Jean-Philippe Thivierge, Département de Physiologie, Université de Montréal, C.P.6128 Succ. Centre-ville, Montréal, Québec, Canada H3T 1J4. E-mail: jean-philippe.thivierge@umontreal.ca

Received 3 August 2006; Accepted 6 December 2006

DOI 10.1002/syn.20378

Published online in Wiley InterScience (www.interscience.wiley.com).

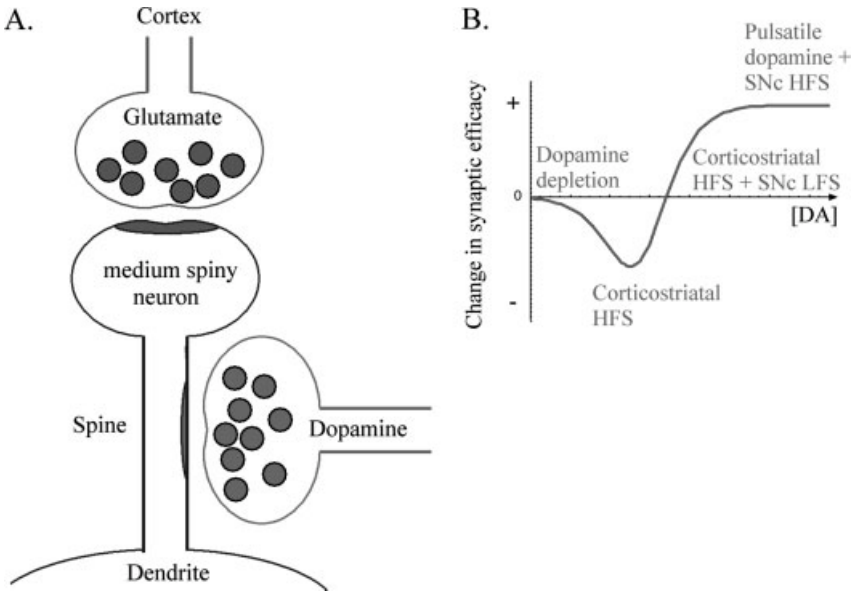

Fig. 1. A schematized corticostriatal synapse depicting the influence of DA on plasticity. In **A**, a cortical cell terminal releases glutamate (Glu) onto the dendritic spine of a medium spiny neuron of the striatum. DA cells of the SNc also project onto medium spiny cells, and, upon stimulation, release vesicular DA into the same synapse. In **B**, DA concentrations relate to change in corticostriatal synaptic efficacy in a biphasic fashion according to different experimental manipulations summarized on the figure. The *x*-axis represents DA concentrations present at the time of corticostriatal activation (adapted from Reynolds and Wickens, 2002).

difference learning (TD, Sutton and Barto, 1998; Montague et al., 2004b). While it remains an open question whether such a rule will ultimately explain all aspects of corticostriatal plasticity, it is consistent with a large body of experimental work demonstrating that presynaptic activity, postsynaptic activity, and DA concentrations are all required to induce changes in synaptic efficacy (Reynolds and Wickens, 2002). Given the range of experimental evidence that can be captured with a multiplicative rule, our goal is to establish the computational consequences of implementing such a rule in a temporally precise model of corticostriatal plasticity, one that includes both STDP and DA fluctuations.

The main goal of the simulations proposed here is to capture the body of experimental evidence on conditions leading to LTP and LTD, as well as to provide a quantitative fit to the STDP data provided by Fino et al. (2005). In an effort to uncover the general principles behind corticostriatal plasticity, we report on the successes and failures of different variants of simulations, including different neural coding schemes (spike-based vs. voltage-based) and learning rules for combining neural activity and DA (multiplicative vs. additive rules). A more specific aim of the current work is to evaluate these different variants with respect to their ability to capture the well-documented relationship between DA concentrations and synaptic plasticity (Reynolds and Wickens, 2002).

Through simulations of induced activity, several novel predictions can be derived from this conjecture. In particular, the influence of DA fluctuations on plasticity may interact with STDP manipulations such that DA concentrations may influence the magnitude of change in synaptic efficacy induced by STDP. Furthermore, different DA concentrations may inverse the typical STDP rule, inducing LTP in conditions that normally elicit LTD, and vice versa, as hinted at by recent experimental work (Fino et al., 2005).

In the following sections, we first review some of the main empirical findings where neural activity and neurotransmission are combined. We then describe a modeling framework that proposes to capture these findings. Third, we employ a voltage-based version of the framework to model the effects of DA concentrations on synaptic efficacy, comparing results between a multiplicative and an additive version of the framework. Fourth, a spike-based version of the model is used to simulate plasticity under various frequencies of stimulation. Finally, we provide a quantitative fit to the STDP data of Fino et al. (2005), and describe novel predictions relative to the effect of DA concentrations on STDP.

A three-way synaptic rule

At the presynaptic level, striatal cells receive activation from the cortex, through corticostriatal pathways (Fig. 1A). At the postsynaptic level, the striatum activates medium spiny neurons that constitute its primary output. Striatal neurons also receive input from the dopaminergic neurons of the substantia nigra pars compacta (SNc) and adjoining midbrain areas (Bjorklund and Lindvall, 1986; West et al., 2003). Several experiments, both in vivo and in vitro, agree that a three-way interaction involving presynaptic depolarization, postsynaptic depolarization, and certain levels of DA concentration are required to induce plasticity in corticostriatal circuits (Reynolds and Wickens, 2002). First, in the absence of any evoked activity, no plasticity is induced, and corticostriatal responses remain stable over time (Walsch, 1993). Second, neither presynaptic activity alone (Choi and Lovinger, 1997; Calabresi

et al., 1992) nor postsynaptic activity alone (Calabresi et al., 1999) can induce plasticity. Rather, a combination of both is required. Third, depleting DA prevents both LTP and LTD forms of plasticity (Calabresi et al., 1992; Centonze et al., 1999; Tang et al., 2001). Taken together, these results suggest a three-way rule for synaptic plasticity, one that includes presynaptic activity, postsynaptic activity, and certain levels of DA concentration (Miller et al., 1981; Schultz, 1998). It is likely that both D_1 and D_2 receptors play a role in such rule (Reynolds and Wickens, 2002). For instance, stimulation of D_1 both enhances L-type high-voltage-activated Ca^{2+} spikes through PKA-dependent pathways but reduces them through PKC-dependent signaling pathways (Young and Yang, 2004). Thus, through the same receptor subtype, DA can differentially affect the very same postsynaptic ion channel.

Manipulations involving individual factors in this three-way rule have shed some light on their influence with respect to synaptic plasticity. However, the resulting nonlinear interactions can be challenging to integrate in a coherent account. For instance, while there are some reports that a combination of pre- and postsynaptic activity enables LTD (Calabresi et al., 1992; Lovinger et al., 1993; Walsh, 1993; Wickens et al., 1996), there are also reports that it can enable LTP (Charpier and Deniau, 1997; Reynolds and Wickens, 2000), or produce no change (Akopian et al., 2000; Partridge et al., 2000; Spencer and Murphy, 2000). Might these seemingly contradictory results actually reflect important principles of plasticity? Assuming that this is indeed the case, can these principles then be captured through a theoretical model?

Synaptic activity and spike-timing-dependent plasticity

A likely possibility is that the above results can be explained by taking into account the precise timing of pre- and postsynaptic spiking potentials, through STDP. On the basis of this idea, synaptic efficacy can be affected in different ways according to whether a presynaptic spike closely precedes or follows an excitatory postsynaptic potential (EPSP). Synaptic rules founded on the principle of spike timing have been found in numerous neural circuits, including retinotectal connections (Zhang et al., 1998), neocortex (Markram et al., 1997), cerebellum (Bell et al., 1997), cochlear nucleus (Tzounopoulos et al., 2004), visual cortex (Froemke and Dan, 2002; Froemke et al., 2005; Sjostrom et al., 2001), barrel cortex (Egger et al., 1999; Feldman, 2000), and cultured hippocampal neurons (Bi and Poo, 1998; Debanne et al., 1994; Debanne et al., 1998; Magee and Johnston, 1997). In a majority of these systems (see the discussion section for other cases), LTP is produced if the presynaptic signal precedes postsynaptic response, and LTD is produced in the inverse scenario where the presynaptic signal follows the postsynaptic response. If the signals occur more than a few milliseconds apart, no change in synaptic efficacy occurs.

In corticostriatal synapses, the idea that the precise timing of spiking activity may have something to do with different effects of plasticity had been proposed some years ago, albeit with no precise mechanism specified (Calabresi et al., 1999; Calabresi et al., 1992; Choi and Lovinger, 1997). More recent work, however, has provided evidence of an STDP rule in corticostriatal synapses (Fino et al., 2005). In this study, STDP was obtained by inducing a low frequency stimulation, consisting of 600 stimuli at 1 Hz, separated by a one second interval. The resulting time window for plasticity was on the order of roughly 40 ms. That is, plasticity could be induced if the pre- and postsynaptic signals occurred within 40 ms of one another. One difference between the STDP found in corticostriatal synapses and that found in other systems is that an inversed rule was found: LTD was induced if a presynaptic potential followed the postsynaptic potential, and LTP was induced in the opposite case. Unfortunately, this study did not investigate the influence of high frequency stimulation (HFS) on plasticity; as will be explored, interesting predictions for this condition can be derived from the modeling framework proposed here.

Effects of dopaminergic concentrations on plasticity

Despite the potential of STDP to explain findings relative to plasticity in corticostriatal projections, it cannot, in its basic form, account for the effect of DA concentrations on synaptic plasticity. As shown experimentally, electrode stimulations of SNc cells can alter DA concentrations at the corticostriatal synaptic junction (for a review, see Reynolds and Wickens, 2002). In turn, the resulting extracellular DA can enable both LTD and LTP, depending on its concentration (Calabresi et al., 1992; for a review, see Reynolds and Wickens, 2002). Low levels of DA during corticostriatal activation induce depression (Kerr and Wickens, 2001), whereas higher doses induce potentiation (Wickens et al., 1996; Reynolds et al., 2001). Thus, the level of DA released around the time of corticostriatal activation is a critical determinant of the direction of synaptic modification. This phenomenon can be summarized through a nonlinear biphasic function that relates DA concentrations to plasticity (Fig. 1B); how this function can emerge from a computational account is the central theme of the current paper.

Modeling the interaction between activity and DA concentrations

Crucial to understanding plasticity in corticostriatal circuits is a means of capturing the interplay between

synaptic activity and the levels of DA present around the time of this activity. Widely-used theoretical models such as TD learning (Sutton and Barto, 1998) suggest that a multiplicative interaction captures a large number of findings in DA-dependent learning. However, it remains a matter of debate whether such multiplicative account is indeed plausible. A large part of this debate can be attributed to methodological issues. Multiplicative interactions seem to be contradicted by experiments where the combined action of pre- and postsynaptic cells was able to induce plasticity, seemingly without the influence of DA (Akopian et al., 2000; Charpier and Deniau, 1997; Partridge et al., 2000; Spencer and Murphy, 2000). However, it is likely that small concentrations of DA were present in these experiments; indeed, traces of DA can be detected in vitro by high-pressure liquid chromatography following HFS of the corticostriatal pathway (Calabresi et al., 1995). Together with evidence that near-total DA prevents the induction of either LTD (Calabresi et al., 1992) or LTP (Centonze et al., 1999), such evidence seems to vindicate the three-way multiplicative rule (Wörgötter and Porr, 2005). While the validity of a multiplicative rule is still debated in experimental work, it is also the focus of much theoretical investigation. At the current time, it remains unclear whether a model based on a multiplicative three-way interaction can account for experimental evidence where precise temporal aspects of synaptic activity and DA neurotransmission are manipulated. Here, we address this issue in a framework that multiplicatively combines the coactivation of corticostriatal cells with MM equations of dopaminergic fluctuations.

METHODS

To capture the synaptic activity of neurons, the proposed model employs leaky integrate-and-fire (LIF) equations (Gütig and Sompolinsky, 2006; Vogels and Abbott, 2005; for applications of LIF neurons to models of the basal ganglia, see Ashby et al., 2005; Brown et al., 1999), combined with a Hebbian-based rule for plasticity, and an extension of MM equations that has been successful in capturing amperometric DA data (Venton et al., 2003). While the proposed equations are phenomenological, and do not represent details of ionic exchanges across the cellular membrane, they nonetheless represent some of the essential characteristics of neuronal communication. Two versions of the LIF equations are possible, namely a voltage-based and a spike-based model, which we now describe in turn; a description of the different ways in which these equations were employed in simulations may be found in Appendix A.

Voltage-based membrane potentials

In the voltage-based version of LIF neurons, the firing rate of a presynaptic cortical neuron is obtained as follows:

$$A_{\text{pre}}(t) = \frac{1}{1+\exp\left(-\beta V_{\text{pre}}(t)\right)} + 0.5, \tag{1}$$

where t is a time-step, β is a parameter, and $V_{\text{pre}}(t)$ represents the presynaptic membrane potential of a cortical neuron. This potential is obtained in a different way for cortical input cells on the one hand:

$$V_{\text{pre}}(t) = V_{\text{rest}} + I_{\text{pre}}(t) + \lambda_{\text{pre}}, \tag{2}$$

and striatal output cells on the other hand:

$$V_{\text{post}}(t) = w(t)A_{\text{pre}}(t) + V_{\text{rest}} + I_{\text{post}}(t) + \lambda_{\text{post}}. \tag{3}$$

The membrane potential fluctuates according to the incoming activation $A_{\text{pre}}(t)$ from cortical neuron, weighted by the efficacy of the connections $w(t)$ between the cortical and striatal neurons. In addition, membrane potentials can be modulated by external currents $I_{\text{pre}}(t)$ and $I_{\text{post}}(t)$, as well as normalized stochastic noise λ_{pre} and λ_{post} in the range [0,1] (the model makes no distinction between intrinsic and extrinsic noise). Finally, V_{rest} corresponds to the membrane resting potential.

Spike-based membrane potentials

In a different version of the model, LIF neurons can represent the postsynaptic activity of a cell through a spike-based signal. Adapting Eq. 3 to this signal, and assuming instantaneous action potentials, the spike-based membrane potentials of striatal cells become

$$V_{\text{post}}(t) = w(t)\sum_{t_{\text{pre}}} K\left(t - t_{\text{pre}}\right) + V_{\text{rest}} + \lambda_{\text{post}}, \tag{4}$$

where t_{pre} denotes the spike times of the cortical afferent and $K(t - t_{\text{pre}})$ is the normalized potential contributed by each incoming spike (Gütig and Sompolinsky, 2006):

$$K\left(t - t_{\text{pre}}\right) = V_0\left(\exp(-(t - t_{\text{pre}})/\tau) - \exp\left(-(t - t_{\text{pre}})/\tau_S\right)\right). \tag{5}$$

A neuron fires whenever the membrane potential ($V_{\text{pre}}(t)$ or $V_{\text{post}}(t)$) reaches the threshold (Θ_{pre} or Θ_{post}) from below, for instance:

$$V_{\text{post}}(t) = \theta_{\text{post}} \quad \text{and} \quad \frac{d}{dt}V_{\text{post}}(t) > 0 \Rightarrow t = t_{\text{post}}. \tag{6}$$

Here, t_{post} represents the firing times of a postsynaptic striatal neuron. If a pulse is triggered, the membrane potential is reset to its resting state V_{rest}, and held there for a time length of T_{refract} corresponding to the absolute refractory period. A similar rule as Eq. 6 applies to V_{pre}.

Spike-timing-dependent plasticity

In the spike-based model described above, the precise temporal order of pre- and postsynaptic spike arrivals plays a central role in determining whether LTP or LTD is induced, leading to an STDP learning rule. This rule for updating synaptic connections can be expressed through the following model:

$$F(\Delta t) = \begin{cases} W_+ \exp(-\Delta t/\tau_+) & \text{if } \Delta t > \psi_{\text{LTP}} \\ -W_- \exp(\Delta t/\tau_-) & \text{if } \Delta t \le \psi_{\text{LTD}} \end{cases} \quad (7)$$

where $\Delta t = t_{\text{post}} - t_{\text{pre}}$ reflects the difference between the last spike arrival times of presynaptic (t_{pre}) and postsynaptic (t_{post}) cells. The parameters ψ_{LTP} and ψ_{LTD} (controlling the beginning in time of the LTP and LTD components) are set to zero for the preliminary simulations presented in this section (but see Fig. 6 for other results). The parameters W_+ and W_- (controlling the magnitude of change in synaptic efficacy) are provided in Appendix A. Finally, the parameters τ_+ and τ_- (controlling the time-course of plasticity) are defined in simulations (see Results section). These values do not reflect precise neurophysiological properties of the synapse; rather, they are meant to capture some general characteristics of plasticity. More detailed biophysical models of STDP are proposed elsewhere (e.g., Pfister et al., 2006), and are beyond the scope of the current paper.

Michaelis-Menton equations

In addition to accounting for STDP, our proposed framework incorporates equations for the effect of DA concentrations on synaptic plasticity. These equations are based on MM kinetics for the diffusion and reuptake of neurotransmitter in the extracellular environment. Such MM kinetics have been applied to modeling both tonic and phasic DA (Venton et al., 2003; Nicholson, 1995; Nicholson, 2001; Montague et al., 2004a). For instance, Montague et al. (2004a) proposes a "kick and relax" model based on MM kinetics. This model was successful at capturing all the dynamics of extracellular DA for repetitive stimuli, and is based upon some of the same basic MM equations as the model developed here. However, the kick and relax model does not capture long-term plasticity, which is the main goal of the current paper.

According to our proposed framework, this effect can be captured through a summation of the diffusion and reuptake of extracellular DA. Diffusion of DA in the extracellular space can be obtained using the following equation:

$$\Delta[D_n(t)]_{\text{diffusion}} = \frac{[D_{n-1}(t)]}{2} + \frac{[D_{n+1}(t)]}{2} - [D_n(t)], \quad (8)$$

where $n \in N$ indexes discrete bins in the extracellular microenvironment (the square brackets "[ ]" indicate a concentration of neurotransmitter). These bins divide the extracellular space around the cell into compartments of even size. A correspondence between amperometric data and a diffusion model similar to that of Eq. 8 has been established using extracellular bins of relatively small size (0.5 μm; Venton et al., 2003; see also Schmitz et al., 2001). The number of such bins employed in the model will alter the results obtained; as a general rule, it is best to keep this number small in order to prevent locations in the outer boundary from affecting results (Venton et al., 2003). Recordings made farther away from the diffusion site are less likely to be due to presynaptic release from a particular cell. A thorough exploration of bin sizes is beyond the scope of our present work; extensive discussions of MM dynamics are available elsewhere (e.g., Nicholson, 1995, 2001).

The following equation represents kinetics of the DA reuptake:

$$\Delta[D_n(t)]_{\text{uptake}} = \frac{V_{\max}\,[D_n(t)]}{K_m + [D_n(t)]}, \quad (9)$$

with constants $V_{\max}$ and K_{m}. The total change in extracellular DA is

$$\Delta[D(t)]_{\text{total}} = \left(\frac{1}{N}\sum_n \Delta[D_n(t)]_{\text{diffusion}}\right) - \left(\frac{1}{N}\sum_n \Delta[D_n(t)]_{\text{uptake}}\right) + [D(t)]_{\text{ext}}, \quad (10)$$

where $[D(t)]_{\text{ext}}$ is a concentration of DA released in the extracellular space by an adjacent SNc cell. Using the above equation, DA concentrations can be updated as follows:

$$[D(t)]_{\text{total}} = [D(t-1)]_{\text{total}} + \eta_D \Delta[D(t)]_{\text{total}}, \quad (11)$$

where η_D is a parameter controlling the rate of change of DA concentrations over time. In its current form, the proposed account does not differentiate between the influence of heteroreceptors and autoreceptors on diffusion and reuptake. In addition, further details would be required to capture differences between D_1 and D_2 receptors.

Putting it all together: A synaptic plasticity rule based on multiplicative interactions

In the voltage-based model, the following Hebbian-based multiplicative rule is proposed to link synaptic activity to DA fluctuations, and update connection efficacies accordingly:

$$\Delta w(t) = A_{\text{pre}}(t) A_{\text{post}}(t)\big([D(t)]_{\text{total}} - b\big), \quad (12)$$

where b is a baseline of DA concentration. In the spike-based model, this equation is modified by

replacing the terms $A_{pre}(t)$ and $A_{post}(t)$ (representing pre- and postsynaptic activity respectively) with a term $F(\Delta t)$ from Eq. 7 representing the STDP rule:

$$\Delta w(t) = F(\Delta t)\big([D(t)]_{total} - b\big). \quad (13)$$

The particular influence of DA concentrations on synaptic plasticity can be captured by setting $b > 0$. In this way, low concentrations of DA will have a negative impact on plasticity, while higher concentrations will have a positive impact (see Results section). Using the rule of Eq. 12 or Eq. 13, connection weights can be updated as follows:

$$w(t+1) = w(t) + \eta_w \ \Delta w(t), \quad (14)$$

where η_w is a learning rate parameter. A potential competitor to the multiplicative rule of Eq. 12 (voltage-based model) could be one involving an additive influence of DA on neural activity:

$$\Delta w(t) = A_{pre}(t)A_{post}(t) + \big([D(t)]_{total} - b\big). \quad (15)$$

Such a rule states that synaptic activity and DA concentrations contribute separately to plasticity, without interactions between them. Similarly, an additive rule can be derived for the spike-based model by modifying Eq. 13:

$$\Delta w(t) = F(\Delta t) + \big([D(t)]_{total} - b\big). \quad (16)$$

Simulations of induced activity and DA concentration

In the framework proposed here, various conditions of stimulations are obtained by different means for the voltage-based and spike-based models. For the former, it is the average amount of activity over time that is considered, and discrete values of $I_{pre}(t)$, $I_{post}(t)$, and $[D(t)]_{ext}$ in the range [0,...,200] are injected. For the latter, it is the frequency of stimulation that is considered, and induced stimuli (for both neural activity and DA release) are represented by a binary state of either 0 (absence) or 100 (presence). When such stimuli are injected in synaptic stimulation (e.g., $I_{pre}(t) = 100$), the result is a direct induction of spiking activity in the cell. Using this scheme, it is possible to precisely control the frequency of stimulation (i.e., the number of spikes per second), and induce phasic as well as tonic activation of the corticostriatal pathways (see Appendix A). While qualitatively similar results emerge from the voltage-based and spike-based models, both provide unique insights into the cellular mechanisms of corticostriatal plasticity.

RESULTS

Voltage-based simulations

Simulations with the voltage-based model involved stimulating cortical and striatal neurons through a presynaptic–postsynaptic train regime (where presynaptic activation preceded postsynaptic activation by a single time-step), and simultaneously inducing DA release (for details of methods see Appendix A).

Simulations performed in this fashion replicate known results relative to the influence of DA concentrations on synaptic plasticity (Fig. 2). In particular, a biphasic relationship between DA and synaptic plasticity emerged from the results (Fig. 2A), as characteristic of experimental data taken over a wide number of reports (Reynolds and Wickens, 2002). With DA depletion, there is no synaptic change possible. If concentrations of DA are low, LTD is produced. If these concentrations are gradually increased, eventually an equilibrium point is reached and no synaptic change is produced. If DA concentrations are increased further above this equilibrium point, LTP is produced. DA concentrations can be increased further, leading to further induction of LTP. However, past a certain level of concentration, the effect of DA on synaptic change reaches an asymptote beyond which no further increase in LTP is possible.

Taken as a whole, the voltage-based simulations described here provide a mechanistic account of the biphasic relationship between DA concentrations and synaptic activity (i.e., Fig. 1B), relying on the influence of induced activity on firing rates (modeled by Eq. 1; Fig. 2B) and DA concentrations (modeled by MM equations; Fig. 2C). By assuming a threshold delimiting DA concentrations that yield LTD vs. LTP (c.f., horizontal line in Fig. 2C), a multiplicative interaction between synaptic voltage and DA concentrations directly produces a biphasic relationship analogous to that observed experimentally (compare Fig. 1B with Fig. 2A).

Is the proposed multiplicative interaction between DA signals and synaptic activity the only possible account of empirical results, or might other types of interactions capture the results just as well? While it is impossible at present to rule out all possible types of interactions, it is nonetheless possible to eliminate some candidates, including an additive rule (Eq. 15). As simulations demonstrate, an additive effect does not provide the target results (Fig. 2D), because a complete deprivation of DA leads to LTD rather than no change. Given this result, we conjecture that the proposed multiplicative interaction constitutes a parsimonious and plausible account of available results. The additive model, at least in its current form, cannot account for these results.

Spike-based simulations

The above results using a voltage-based model of synaptic activity replicate some key findings associated with the influence of DA concentrations on plasticity. Can these same results emerge independently of the choice of model for synaptic activity? The goal of the current section is to show how similar effects also hold with a spike-based model. In addition, such model ena-

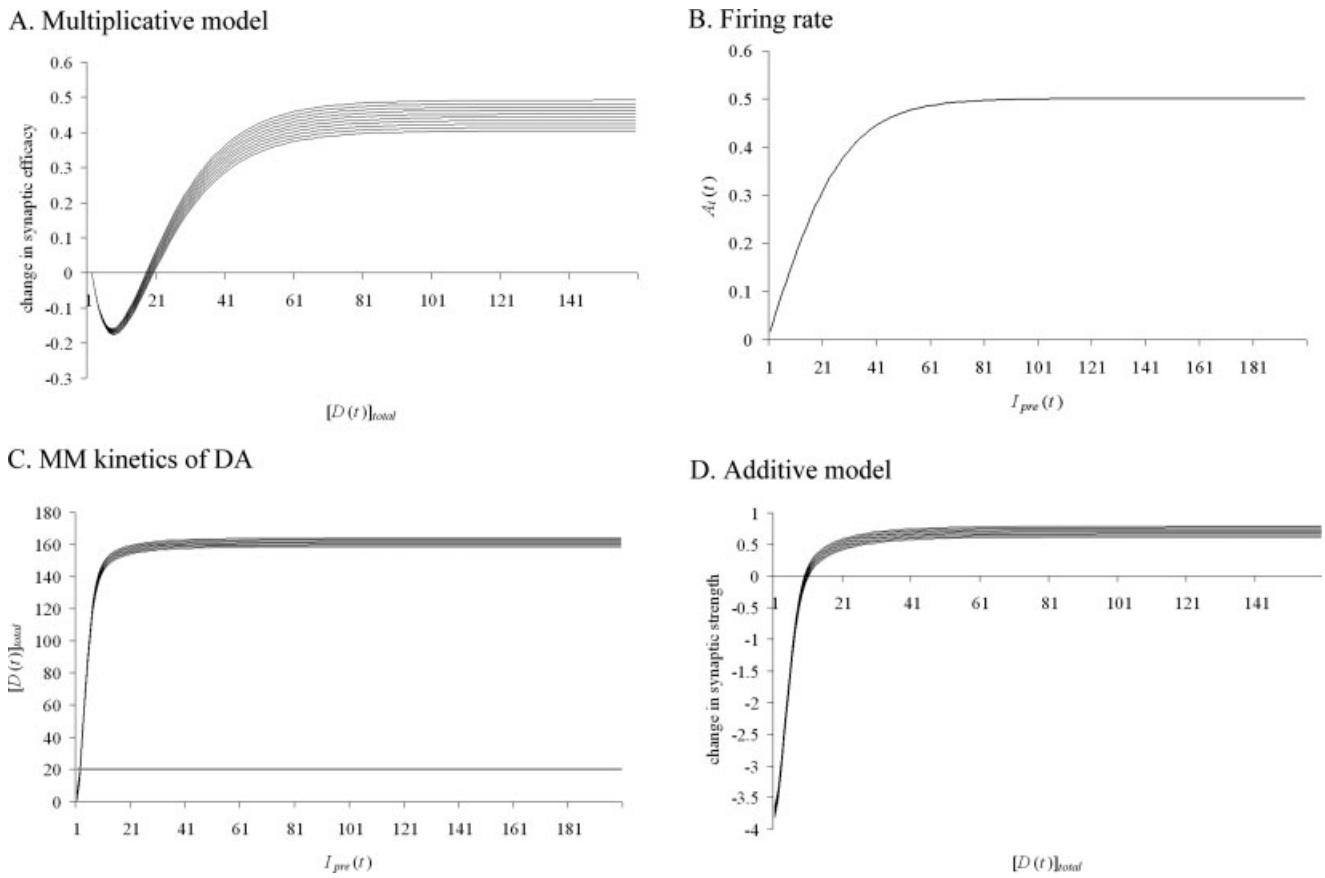

Fig. 2. Voltage-based model, where DA concentrations affect changes in synaptic efficacy, and induce both LTP and LTD. **A**: Multiplicative model of the effect of DA concentrations on plasticity. Injected current into the DA cell (see Appendix A): $I_{\text{pre}}(t) = I_{\text{post}}(t) = [1.0, 200]$. A presynaptic–postsynaptic activation regime is employed for corticostriatal cells. **B**: Firing rate of a stimulated cortical cell as a function of input currents $I_{\text{pre}}(t)$. **C**: Extracellular dopaminergic concentrations as modeled through MM kinetics. The gray line indicates the threshold b for LTD/LTP (see text). **D**: Additive model. (A-D) Different values of baseline DA concentrations are plotted, from 20.09 (higher values on all graphs) to 20.19 (lower values on all graphs), at intervals of 0.01.

bles the incorporation of precise temporal dynamics in spike patterns, based on STDP.

In simulations using the spike-based model, changes in synaptic efficacy can be modulated by controlling the frequency of stimulation of the corticostriatal cells (Fig. 3). By stimulating the cortical and striatal neurons using either HFS or low-frequency stimulation (LFS), different consequences for plasticity are obtained: while low-frequency stimulation (LFS) leads to depression, high-frequency stimulation (HFS) leads to potentiation. In addition, changes in synaptic efficacy can be modulated by inducing DA release in the SNc pathways. As in the voltage-based model, the spike-based model assumes a threshold delimiting the DA concentrations that result in LTD vs. LTP (c.f., solid horizontal lines of top figures in Fig. 3A, B). In LFS, because only low concentrations of extrasynaptic DA are released, this threshold is never reached, and the net effect of DA concentrations is LTD. Conversely, in HFS, higher concentrations of extrasynaptic DA are released, resulting in LTP.

A summary of the effect of stimulation frequency on synaptic plasticity is shown in Figure 4. These results describe how different frequencies of stimulation in the spike-based model can account for the characteristic curve relating DA concentrations to change in synaptic efficacy in various experiments, as captured with the voltage-based model described above (Fig. 2A). In the multiplicative model (Fig. 4A), a low frequency of stimulation (i.e., below approximately 20 Hz) leads to a net effect of LTD; higher frequencies are required for the expression of LTP. These results are in agreement with experimental evidence showing an equilibrium point between LTD and LTP at 20 Hz stimulation (Reynolds and Wickens, 2002). The multiplicative model can be compared to an additive model (Fig. 4B). The main difference between the two resides in conditions of LFS, where the additive model predicts LTD rather than no change in synaptic efficacy, as corroborated experimentally (Calabresi et al., 1992; Reynolds and Wickens, 2002), and consistent with the multiplicative model. In general, the simulation results obtained with the spike-based model are quite noisy when compared to the voltage-based account, in part because the latter embodies a smooth approximation of a process that is in reality subject to large amounts of noise.

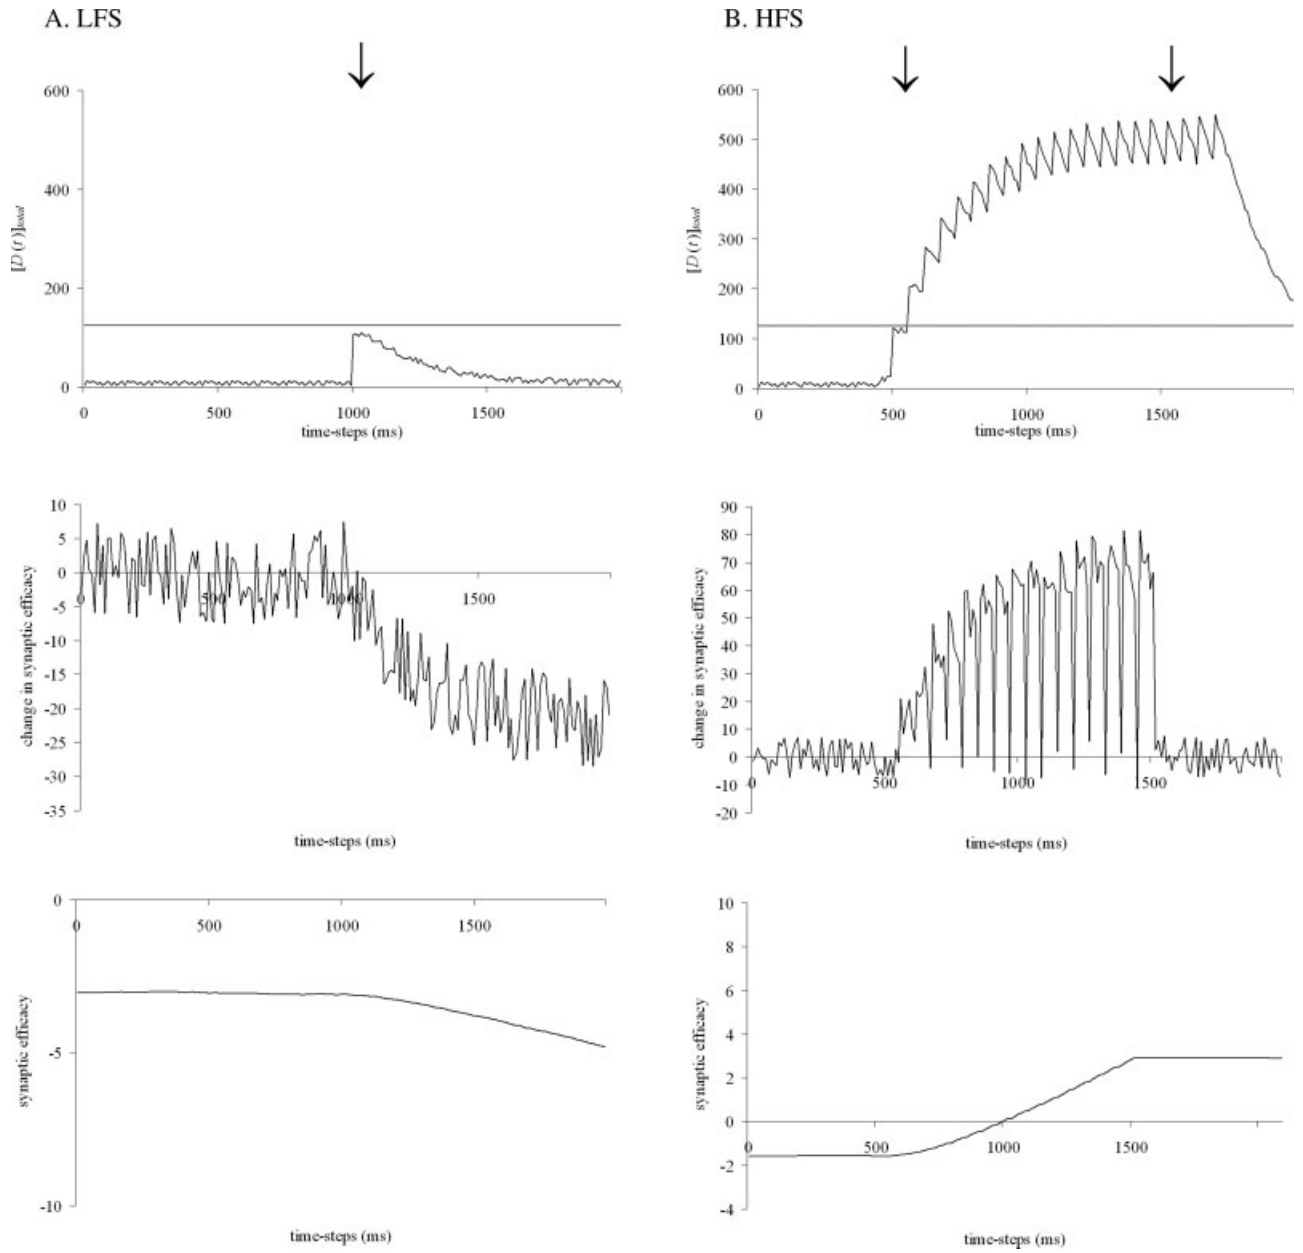

Fig. 3. Spike-based model of the influence of stimulation frequency on plasticity in corticostriatal synapses. **A**: LFS (single stimulation indicated by the arrow in the top figure). LFS induces DA release in the extracellular space, but remains below the threshold for expression of LTP (top figure, horizontal line). As a result, the net impact on plasticity is LTD (bottom figure). **B**: HFS (stimulation at 100 Hz; arrows indicate the start and end of stimulation in the figure). Induction of HFS pushes DA concentrations above the LTP threshold (top figure), and therefore leads to LTP (bottom figure). Corticostriatal spike trains followed a pre-post regime.

In the above simulations, induced DA release and corticostriatal stimuli are delivered at the same frequency (LFS or HFS). However, it is also possible to test conditions under which DA release and corticostriatal stimuli do not follow the same frequency. As suggested by experimental evidence (Reynolds and Wickens, 2000; c.f., Fig. 1B), corticostriatal HFS in the absence of SNc stimulation leads to LTD. However, when a similar manipulation is combined with LFS of the SNc pathways, no change in synaptic efficacy is induced. These results are captured through simulations performed with a multiplicative model where the induction of DA release is reduced or eliminated altogether (Fig. 5). If corticostriatal HFS is combined with an elimination of all induced DA release (Fig. 5A), DA levels never reach the baseline concentrations required for the expression of LTP; as a result, depression is expressed. A corticostriatal HFS combined with low frequency induction of DA release (Fig. 5B) produces different results: because DA concentrations now oscillate around the LTP threshold, both LTP and LTD are expressed (albeit only slightly), with a net effect of little or no change in plasticity.

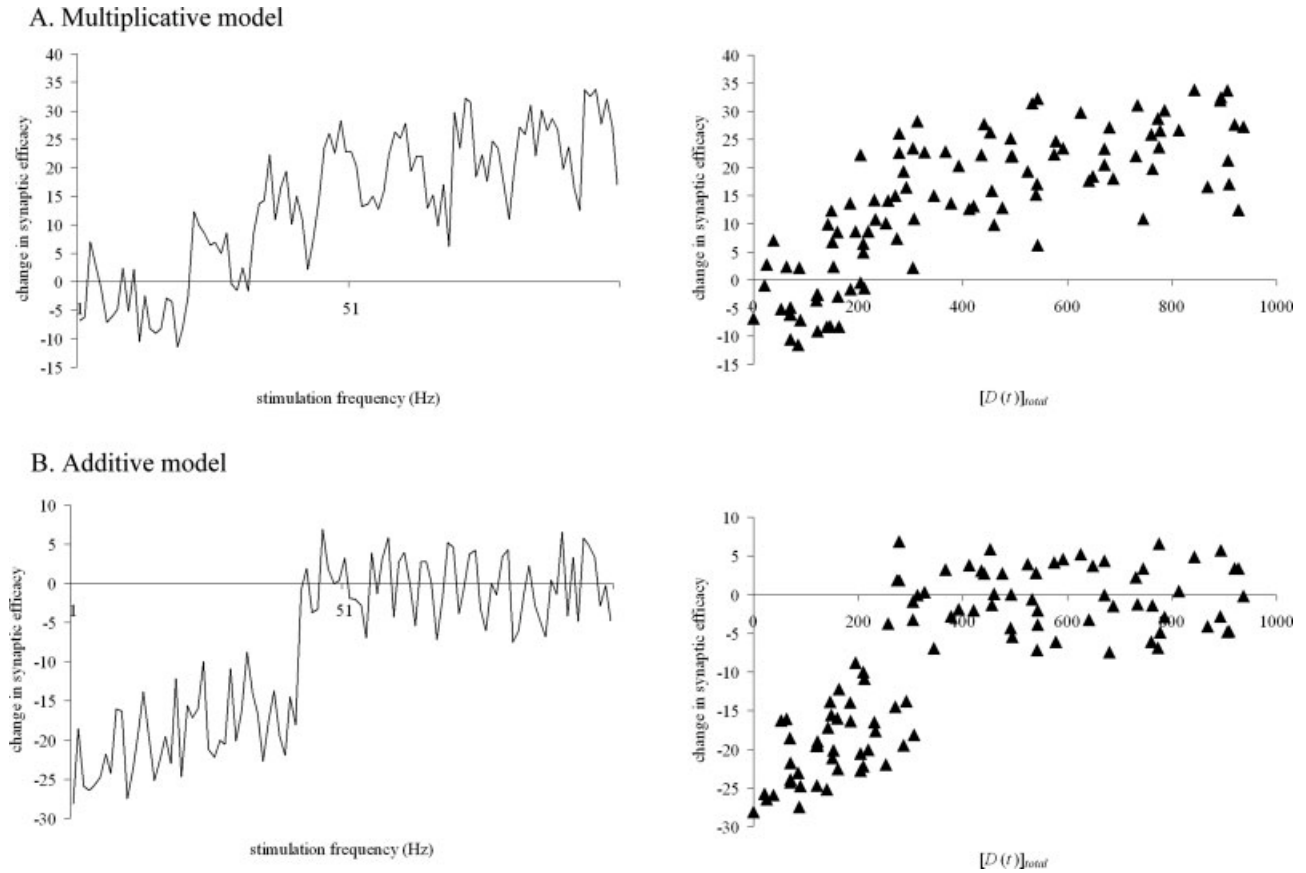

Fig. 4. Summary of the effect of stimulation frequency and DA concentrations on synaptic plasticity. With a multiplicative model, LFS induces LTD, while HFS induces LTP (**A**, left figure). This effect can be linked to the release of extracellular DA (A, right figure). With an additive model, lack of corticostriatal activity leads to strong LTD (**B**, left figure). In this model, DA depletion leads to LTD, a result inconsistent with experimental evidence (B, right figure). Both multiplicative and additive models presented are generated using the spike-based model.

Results of LFS corticostriatal simulations presented here may offer an explanation for the reversed STDP observed in Fino et al. (2005). In this experimental work, pulses were delivered at a 1 Hz frequency, separated by intervals of one second. As our simulations show (Fig. 3A), with such low frequency of stimulation, DA concentrations never cross the threshold above which LTP would be expressed. In addition, DA returns to a resting concentration after approximately 500 ms, as consistent with previous work (Venton et al., 2003). As a consequence of introducing an interval of one second between stimuli, DA would never accumulate above the required threshold for LTP. Next, we explore how the assumption of a negative effect of DA concentrations on plasticity can be used to provide a quantitative fit to the corticostriatal STDP data provided by Fino et al. (2005).

Quantitative fit to corticostriatal STDP

The data of Fino et al. (2005) reflect changes in synaptic efficacy that occur an estimated 45 min after induction of STDP. Here, we employ our multiplicative framework to account for this data. The model of STDP proposed in Eq. 7 has been employed to model plasticity in several other simulations (e.g., Bi and Wang, 2002; Bi and Poo, 1998; Song and Abbott, 2001). A particular consideration in extending these results to corticostriatal synapses is that the STDP function reported by Fino et al. (2005) is reversed when compared to that obtained in other brain systems (Fig. 6E, filled triangles). In this reversed STDP, LTD is expressed when $\Delta t > 0$, which means that the presynaptic spike followed the postsynaptic spike (i.e., $t_{\text{pre}} < t_{\text{post}}$). Conversely, LTP is expressed when $\Delta t < 0$, that is, when the presynaptic spike preceded the postsynaptic spike (i.e., $t_{\text{pre}} > t_{\text{post}}$). The standard STDP rule predicts just the reverse of these effects; we propose that a term relative to the negative influence of DA concentrations on synaptic plasticity is required to capture this finding.

In addition to capturing this reversed STDP, one challenge in adapting Eq. 7 to corticostriatal synapses is that our model must allow for a relatively large time window between the LTP and LTD portions of the STDP function (i.e., ± 10 ms), as is reflected in Fino

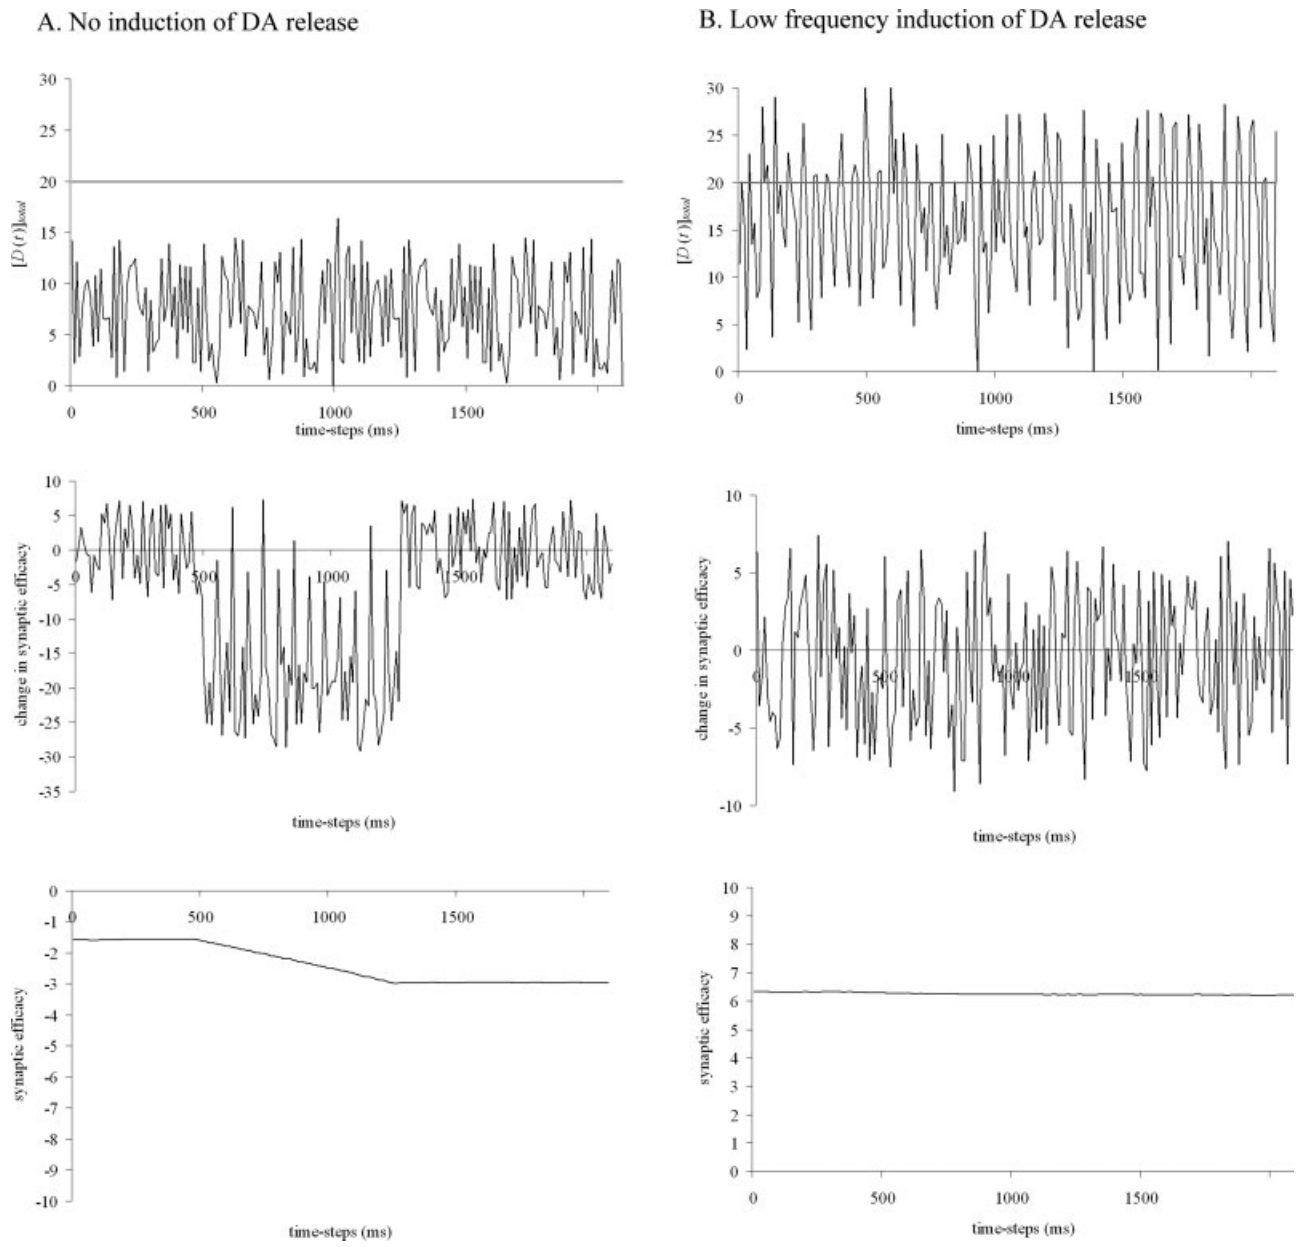

Fig. 5. Spike-based model varying the frequency of induced DA release. In **A**, HFS corticostriatal stimulation in combined with no induction of DA release ($[D(t)]_{ext} = 0$ throughout). This manipulation leads to low levels of extracellular DA, below the threshold for expression of LTP (top figure, gray line). This results in a net effect of LTD (middle and bottom figures). **B** is the same as A, but induces a low frequency of DA (20 Hz LFS with $[D(t)]_{ext} = 20$). As a result, DA concentrations sometimes cross the threshold for LTP (top figure), and changes in synaptic efficacy oscillate between LTP and LTD (middle and bottom figures).

et al.'s (2005) data. In order to account for this observation, the STDP model proposed in Eq. 7 is employed, with values of the parameters ψ_{LTP} and ψ_{LTD} that are increased above zero in order to widen the distance between the LTP and LTD components. The value of $F(\Delta t)$ in Eq. 7, reflecting a measure of distance between pre- and postsynaptic spikes, is used to determine the change in synaptic efficacy (Eq. 13). For the purposes of the preliminary quantitative analysis presented here, we set the model to represent a negative influence of DA on plasticity by making $[D(t)]_{total} - b = -10$, thus leading to the following weight update:

$$\Delta w(t) = -10\, F(\Delta t). \tag{17}$$

While further DA measurements would be required to obtain more precise values for these parameters, our current instantiation suggests that, whatever these values, DA should have a negative influence on plasticity if the reversed STDP rule is to be observed. Such

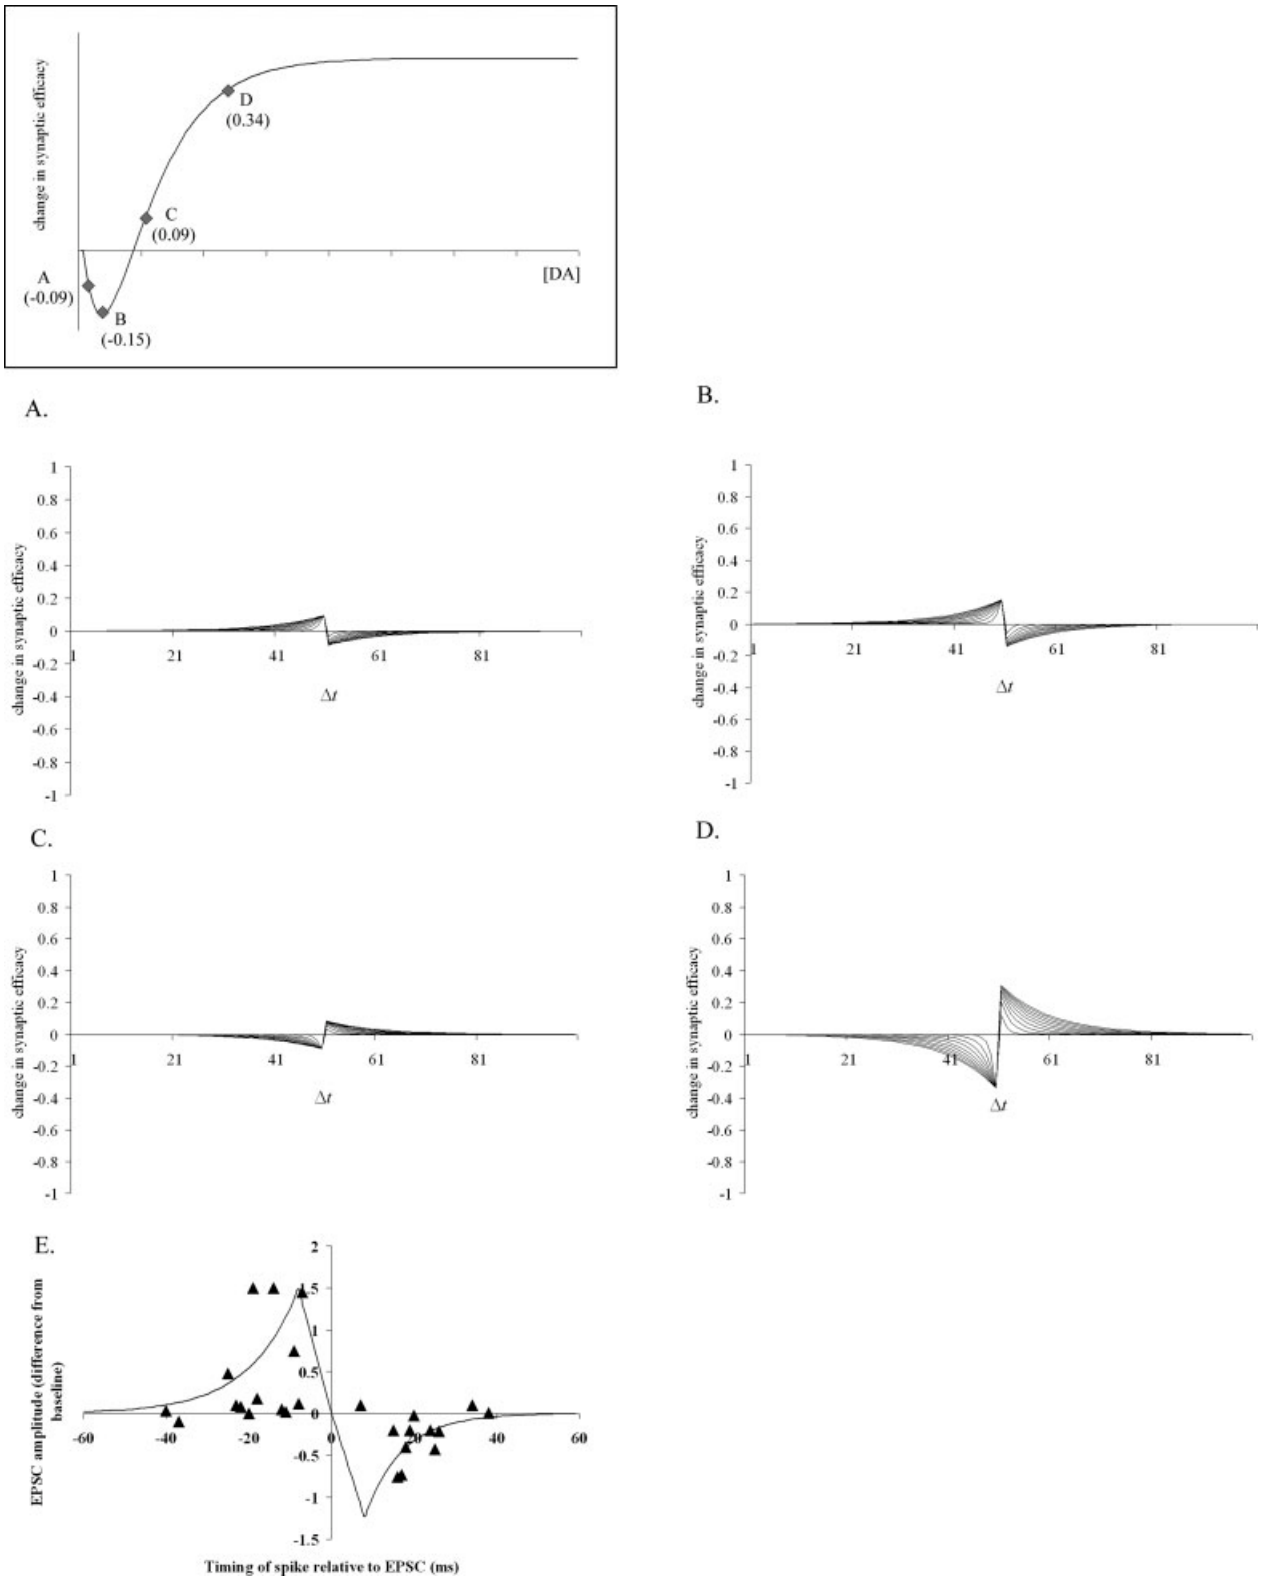

Fig. 6. Proposed effect of DA on spike-timing dependent plasticity. **A–D:** Plots were generated through Eq. 7 for STDP. Different values of time constants (τ_+ and τ_-) are plotted in each figure, in the range [1,10] at intervals of 1.0 mV (higher values leading to a wider window for plasticity). Inset: filled diamonds represent different values of $[D(t)]_{\text{total}}$ (c.f., Eq. 12) used to generate plots A–D. **E**: Data from Fino et al. (2005) representing different recorded striatal cells (filled triangles) is modeled through STDP (solid line; Eq. 7 combined with 17). The data represents a difference from baseline condition in excitatory postsynaptic current (EPSC); this is captured in the model through adjustments in synaptic efficacy ($\Delta w(t)$). Model parameters: $\tau_+ = 9$, $\tau_- = 12$, $W_+ = 3$, $W_- = 0.29$, $\psi_{\text{LTP}} = 8$, and $\psi_{\text{LTD}} = -7.3$.

negative influence is likely to result from low concentrations of extrasynaptic DA, as demonstrated both experimentally (c.f., Fig. 1B) and computationally through the results presented here (c.f., Fig. 2A).

Using Eqs. 7 and 17, a reasonable fit to the data of Fino et al. (2005) is obtained by selecting appropriate parameters in the model (τ_+, τ_-, W_+, W_-, ψ_{LTP}, ψ_{LTD}) through an exhaustive search through the space of possible values, performed independently for each parameter (Fig. 6E). This search minimized the sum of squared errors between each individual data point and the corresponding value generated by the model. Performing a search through the parameter space of Eqs. 7 and 17 was essential to fitting the STDP data, due to the model's sensitivity to its parameters.

While the model does not provide a perfect quantitative fit to the data—due to factors including potential sources of noise in the data—it was nonetheless able to capture the main qualitative characteristics of variation in synaptic efficacy following induction of STDP. More precise quantitative fits would only be expected to provide refinements to these results, which already demonstrate the relevance of the proposed model in capturing corticostriatal STDP. Another caveat is the small number of data points obtained in the Fino et al. (2005) study; however, we consider it likely that further data would only corroborate the presence of an STDP rule in corticostriatal synapses. Along with previous successes in modeling STDP in several systems (e.g., Bi and Poo, 1998; Song and Abbott, 2001), our results promote further the proposed STDP rule (c.f., Eqs. 7 and 17) as an ubiquitous principle for plasticity across many brain systems.

Inducing STDP under various conditions of DA concentration

While the previous section provided a quantitative fit to experimental data using a fixed value of DA, it is also possible to vary the influence of DA concentrations, and evaluate the computational consequences of a multiplicative interaction between DA concentrations and STDP. We investigate the effects of DA concentrations on the STDP rule by setting arbitrary values of $[D(t)]_{\mathrm{total}}$ in Eq. 13 (Fig. 6, inset), and inducing an STDP stimulation. While the values chosen for $[D(t)]_{\mathrm{total}}$ are not meant to fit particular empirical data, they nonetheless capture the essence of the relationship between DA concentrations and change in synaptic efficacy (Fig. 1B). As results of simulations suggest, DA may be able to drastically alter the form of the STDP rule, potentially reversing the conditions under which LTP and LTD are expressed (Figs. 6A–D). Indeed, our simulations suggest that the biphasic influence of DA concentrations on plasticity (Fig. 6, inset) can modulate both the direction and magnitude of change in synaptic efficacy induced by STDP. In this way, moderate to high concentrations of DA could produce a similar form of plasticity as observed across many other brain centers (e.g., Zhang et al., 1998), where a pre-post regime induces potentiation, and a post-pre regime induces depression. Conversely, lower concentrations of DA may inverse this rule, and induce depression in pre-post regimes, and potentiation in post-pre regimes.

DISCUSSION

Our computational integration of STDP and DA kinetics suggests that the characteristic influence of DA concentrations on plasticity may be the consequence of a multiplicative rule. When taken together, the various aspects of the proposed framework were able to capture a wide range of seemingly contradictory experimental results on conditions leading to the expression of LTP and LTD. For instance, our simulations suggest that a combination of pre- and postsynaptic activity can result in LTD (Calabresi et al., 1992; Lovinger et al., 1993; Walsh, 1993; Wickens et al., 1996), LTP (Charpier and Deniau, 1997; Reynolds and Wickens, 2000), or no change in synaptic efficacy (Akopian et al., 2000; Partridge et al., 2000; Spencer and Murphy, 2000), depending on the concentrations of DA released around the time of corticostriatal activation. Qualitatively, the results of simulations were not found to depend on the type of neural model employed (i.e., spike-based vs. voltage-based), but failed to emerge if the multiplicative rule was replaced with an additive model. These results constitute a step forward in demonstrating the ubiquity of multiplicative interactions between DA concentrations and synaptic activity.

The proposed simulations raise the intriguing yet empirically-motivated possibility that the STDP rule may be reversed when DA extrasynaptic concentrations are below certain levels. The conjectured STDP "flip" under low DA concentrations explains the results of Fino et al. (2005), who reported a reverse STDP in corticostriatal synapses. In this experiment, induced stimulations were performed under LFS, with 600 stimuli at 1 Hz, each separated by a one second delay. According to our conjecture, this manipulation may have led to low extrasynaptic DA concentrations, thus resulting in the inversed STDP observed in their data as well as in our model (Figs. 6B–C). At the current time, this explanation remains speculative and awaits further experimental validation. A fundamental question that is not addressed here is how certain concentrations of DA are tied to the expression of LTD and LTP, and in particular what is the molecular manifestation of a threshold in concentration where the net effect of DA concentration on synaptic plasticity changes from LTD to LTP (e.g., Fig. 2C, horizontal line). Speculatively, there may be a certain level of DA concentration above which a high enough postsynaptic influx of

Ca^{2+} is induced (Young and Yang, 2004), in turn responsible for the expression of LTP. With lower DA concentrations, lower levels of Ca^{2+} may be induced, leading to LTD. This possibility leaves unanswered the question of why and how a particular threshold of DA concentration is responsible for these events. Interestingly, reversed STDP (also termed "anti-STDP") has been reported in other systems, including the electrosensory system of a weakly electric fish (Bell et al., 1997), the cerebellum (Wang et al., 2000), and the dorsal cochlear nucleus of the brain stem (Tzounopoulos et al., 2004). In these systems, the hypothesized role of anti-STDP is to equalize synaptic efficacies. Because anti-STDP directly counters STDP, it may act only under certain limited conditions, and potentially be triggered by neuromodulatory processes (Goldberg et al., 2002; Rumsey and Abbott, 2004). Results of our theoretical approach suggest, for the first time, that a balance between STDP and anti-STDP could be under the influence of extrasynaptic DA concentrations present around the time of corticostriatal activation. How DA concentrations can be employed to control the joint operation of STDP and anti-STDP in order to promote both associative learning and synaptic homeostasy, as well as the functional consequences of this joint operation, are currently being investigated using the framework proposed here.

The function proposed here to relate DA concentrations to changes in synaptic efficacy may be employed to guide further experimentation on the conditions under which the STDP rule may be reversed. One of the main predictions to follow from simulations is that HFS would induce a standard (i.e., non-reversed) STDP, when compared to the LFS condition. Interestingly, this effect could potentially be addressed with current experimental technologies; for instance, through in vitro patch-clamping combined with amperometric measurements of evoked DA (e.g., Lavin et al., 2005; Venton et al., 2003).

An assumption of the proposed model is that DA concentrations can affect the plasticity resulting from the timing of pre- and postsynaptic spikes relative to one another. The results captured by our model are directly dependent upon this assumption. If, for instance, DA could only influence pre- and postsynaptic spikes independently of one another, higher concentrations would shift the STDP rule out of a time window where plasticity is possible. These results would stand in stark contrast to empirical evidence showing that high concentrations of DA (i.e., through HFS) induce LTP. Thus, the model proposed here argues for the necessity of DA concentrations to affect the relative timing of pre- and postsynaptic spikes in order to capture the data presented.

One question that has not been fully addressed empirically concerns the origins of an equilibrium point where, at certain mid-level concentrations of DA, no synaptic change is induced (c.f., Fig. 1B). As a potential answer, Reynolds and Wickens (2002) suggested that it is "ongoing activity in the corticostriatal and nigrostriatal pathways [that] maintains an equilibrium level of synaptic efficacy at the zero-crossing point." However, it is difficult to relate this proposition to the actual DA concentrations inducing plasticity, let alone to a mechanism by which it would occur. One direct answer to this issue comes from the proposed interaction between synaptic activity and DA. As modeled (c.f., Eqs. 12–13), an equilibrium point between LTP and LTD is reached when the DA signal itself reaches the threshold b. Given a value of $[DA(t)]_{\text{total}} = b$, it is straightforward to see that the change in synaptic efficacy will fall to zero (i.e., $\Delta w = 0.0$), because it is the result of a multiplication between the DA signal and synaptic activity. Hence, the model predicts that an equilibrium point between LTP and LTD is attained when the DA signal's net effect on plasticity is null. One drawback of this prediction is that it may be difficult to validate empirically without a clear conception of the independent effect of DA concentrations on plasticity. Validation of the proposed hypothesis may thus require several steps of investigation. At the current time, most investigations of corticostriatal plasticity have not performed direct measurements of extracellular DA concentrations (e.g., Kerr and Wickens, 2001; Reynolds et al., 2001; Wickens, 1996); this would be essential for a full validation of the simulations proposed here.

Through future theoretical investigations, it could be possible to explore other impacts of DA on STDP, including modulations in time constants τ_+ and τ_-, affecting the size of the time window within which plasticity is possible (Fig. 6). With larger values of DA (positive or negative), the size of this plasticity window will be larger. As a result, LTP and LTD may be obtained between cells whose spike arrival times are further apart. The narrowing and widening of the plasticity window could be accounted for by known links between STDP and voltage-gated Ca^{2+} channels. According to the standard model of long-term plasticity, moderate Ca^{2+} levels above baseline can induce LTD, while sufficiently high levels can induce LTP (Bi and Rubin, 2005; c.f., Karmarkar and Buonomano, 2002). A proposed role of DA could be to modulate the plasticity window by controlling the time-course of Ca^{2+} activation. This conjecture is likely dependent upon two (or more) separate mechanisms for LTP and LTD (see Bi and Rubin, 2005), which implies that DA concentrations would need to act independently on each of these mechanisms. Previous theoretical work has recognized these possible links between neurotransmission and modulation in plasticity window (Senn et al., 2000), albeit not in the specific context of corticostriatal synapses.

While the proposed model constitutes an advancement in understanding the interactions between synaptic activity and neurotransmission, further efforts

are required to understand the contribution of other possible factors influencing corticostriatal plasticity. These include, for instance, the activation of group I metabotropic glutamate receptors responsible for LTD (Calabresi et al., 1999; Dos Santos Villar and Walsh, 1999; Gubellini et al., 2001). In addition to DA and STDP factors contributing to the different results of LTP and LTD, other likely contributors have been identified, including the anatomical location of recorded and stimulated neurons, and certain developmental changes associated with corticostriatal plasticity (Reynolds and Wickens, 2002). Given certain deliberate limits in the scope of the proposed model, more complete accounts will be required to incorporate these factors computationally (for a review of biophysical models of STDP, see Pfister et al., 2006; Senn et al., 2000).

Further theoretical efforts are also required to establish links between the current model and investigations of reinforcement learning. In this field, one of the most studied models is TD learning (Houk et al., 1995; for a review, see Montague et al., 2004b). TD models of corticostriatal circuits envision DA as a signal of error in reward prediction (Hollerman and Schultz, 1998). Not unlike the model presented here, TD learning represents synaptic plasticity as a three-way multiplicative rule involving presynaptic (cortical) activity, postsynaptic (striatal) activity, and phasic variations in DA (Suri and Schultz, 1999). The most salient difference between our model and TD is that the latter describes plasticity along a coarser time-scale. As a result, TD is not amenable to STDP, a phenomenon that requires neural activity at a fine-grained timing. Further theoretical investigations are required to link the proposed framework to TD learning. Bridging the two models would help strengthen links between the biokinetics of DA diffusion on the one hand, and reinforcement learning approaches on the other. One possible starting point could be the work of some authors suggesting that STDP rules, such as the one proposed here, embody a form of spike-based reinforcement learning (Roberts, 1999; Rao and Sejnowski, 2001; Xie and Seung, 2004). The upshot of such an investigation may be a demonstration that the DA concentrations induced experimentally (e.g., Fig. 1B) can be expressed naturally in a behaving animal according to particular task demands and reward schedules.

Further simulations should also address other aspects of the study by Fino et al. (2005) that are beyond the scope of the current work. For instance, anti-Hebbian plasticity was obtained in conditions where postsynaptic cells were maintained at their resting potentials. These results remain to be captured through some variation of the rules proposed here; among other considerations, such rule would need to differentiate between pre- and postsynaptic mechanisms for plasticity.

Much additional work will also be required in order to fully elucidate the question of whether a multiplicative rule is implemented at the cellular level; here, our goal was simply to investigate the computational consequences of such rule, and perform some preliminary comparisons to an additive model. Further comparisons along similar lines will need to be performed in order to rule out (or perhaps validate) alternative explanations.

Despite its speculative nature, our proposal may open promising lines of investigation in the study and treatment of neurological diseases, given that pathological modulations in DA may lead to opposite effects on plasticity when compared to normal cases. For instance, in Parkinson's disease (PD), there is typically a marked depletion of DA in the dorsal striatum (Kish et al., 1988). So far, computational efforts to model PD (e.g., Monchi et al., 2000) have not thoroughly explored the possibility that an equilibrium in DA neurotransmission may be required for normal cellular communication and plasticity.

In sum, while many aspects of the proposed interactions between synaptic activity and DA remain speculative, the current work nonetheless captures an important nonlinear function relating DA concentrations to changes in synaptic efficacy (see Fig. 2A). Further, it is argued that a basic multiplicative rule lies as the basis of this interaction, and that this rule is independent of the type of neural model employed. The proposed account leads to several predictions regarding the influence of DA on STDP that are directly testable using available technologies. Because of the simplicity, robustness, and ease of interpretation of this rule, there exists the prospect that it could serve as basis for several other findings relating synaptic activity and neurotransmission.

ACKNOWLEDGMENTS

This research was supported by a doctoral fellowship from the Fonds Québécois de Recherche sur la Nature et les Technologies (FQRNT) to J.P.T. and a doctoral studentship from the CIHR New Emerging Team in Computational Neuroscience to F.R. This work benefited from discussions with Louis-Eric Trudeau (Université de Montréal) and Peter Schizgal (Concordia University).

REFERENCES

Abbott LF, Nelson SB. 2000. Synaptic plasticity: Taming the beast. Nat Neurosci 3:1178–1183.

Akopian G, Musleh W, Smith R, Walsh JP. 2000. Functional state of corticostriatal synapses determines their expression of short- and long-term plasticity. Synapse 38:271–280.

Ashby FG, Ell SW, Valentin VV, Casale MB. 2005. FROST: A distributed neurocomputational model of working memory maintenance. J Cog Neurosci 17:1728–1743.

Bell CC, Han VZ, Sugawara Y, Grant K. 1997. Synaptic plasticity in a cerebellum-like structure depends on temporal order. Nature 387:278–281.

Bi G-Q, Poo MM. 1998. Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. J Neurosci 18:10464–10472.

Bi G-Q, Rubin J. 2005. Timing in synaptic plasticity: From detection to integration. Trends Neurosci 28:222–228.

Bi G-Q, Wang H-X. 2002. Temporal asymmetry in spike timing-dependent synaptic plasticity. Physiol Behav 77:551–555.
Bjorklund A, Lindvall O. 1986. Catecholaminergic brain stem regulatory systems. In: Mountcastle VB, Bloom FE, Geiger SR, editors. Handbook of physiology: The nervous system, Part I, Vol. 4. Bethesda: American Physiological Society. p155–235.
Brown J, Bullock D, Grossberg S. 1999. How the basal ganglia use parallel excitatory and inhibitory learning pathways to selectively respond to unexpected reward cues. J Neurosci 19:10502–10511.
Calabresi P, Maj R, Pisani A, Mercuri NB, Bernardi G. 1992. Long-term synaptic depression in the striatum: Physiological and pharmacological characterization. J Neurosci 12:4224–4233.
Calabresi P, Centonze D, Gubellini P, Marfia GA, Bernardi G. 1999. Glutamate triggered events inducing corticostriatal long-term depression. J Neurosci 19:6102–6110.
Calabresi P, Fedele E, Pisani A, Fontana G, Mercuri NB, Bernardi G, Raiteri M. 1995. Transmitter release associated with long-term synaptic depression in rat corticostriatal slices. Eur J Neurosci 7:1889–1894.
Centonze D, Gubellini P, Picconi B, Calabresi P, Giacomini P, Bernardi G. 1999. Unilateral dopamine denervation blocks corticostriatal LTP. J Neurophysiol 82:3575–3579.
Charpier S, Deniau JM. 1997. In vivo activity-dependent plasticity at cortico-striatal connections: Evidence for physiological long-term potentiation. Proc Natl Acad Sci USA 94:7036–7040.
Choi S, Lovinger DM. 1997. Decreased probability of neurotransmitter release underlies striatal long-term depression and postnatal development of corticostriatal synapses. Proc Natl Acad Sci USA 94:2665–2670.
Dan Y, Poo MM. 2004. Spike timing-dependent-plasticity of neural circuits. Neuron 44:23–30.
Debanne D, Gahwiler BH, Thompson SM. 1994. Asynchronous pre- and postsynaptic activity induces associative long-term depression in area CA1 of the rat hippocampus in vitro. Proc Natl Acad Sci USA 91:1148–1152.
Debanne D, Gahwilder BH, Thompson SM. 1998. Long-term synaptic plasticity between pairs of individual CA3 pyramidal cells in rat hippocampal slice cultures. J Physiol (Lond) 507:237–247.
Dos Santos Villar F, Walsh JP. 1999. Modulation of long-term synaptic plasticity at excitatory striatal synapses. Neurosci 90:1031–1041.
Egger V, Feldmeyer D, Sakmann B. 1999. Coincidence detection and changes of synaptic efficacy in spiny stellate neurons in rat barrel cortex. Nat Neurosci 2:1098–1105.
Feldman DE. 2000. Timing-based LTP and LTD at vertical inputs to layer II/III pyramidal cells in rat barrel cortex. Neuron 27:45–56.
Fino E, Glowinski J, Venance L. 2005. Bidirectional activity-dependent plasticity at corticostriatal synapses. J Neurosci 25:11279–11287.
Froemke RC, Dan Y. 2002. Spike-timing-dependent synaptic modification induced by natural spike trains. Nature 416:433–438.
Froemke RC, Poo MM, Dan Y. 2005. Spike-timing-dependent synaptic plasticity depends on dendritic location. Nature 434:221–225.
Gerstner W, Kistler W. 2002. Spiking neuron models. Cambridge: Academic.
Goldberg J, Holthoff K, Yuste R. 2002. A problem with Hebb and local spikes. Trends Neurosci 25:433–435.
Gubellini P, Saulle E, Centonze D, Bonsi P, Pisani A, Bernardi G, Conquet F, Calabresi P. 2001. Selective involvement of mGlu1 receptors in corticostriatal LTD. Neuropharmacology 40:839–846.
Gurney K, Prescott TJ, Wickens JR, Redgrave P. 2004. Computational models of the basal ganglia: From robots to membranes. Trends Neurosci 27:453–459.
Gütig R, Sompolinsky H. 2006. The tempotron: A neuron that learns spike timing-based decisions. Nat Neurosci 9:420–428.
Hollerman JR, Schultz W. 1998. Dopamine neurons report an error in the temporal prediction of reward during learning. Nat Neurosci 1:304–309.
Houk JC, Adams JL, Barto AG. 1995. A model of how the basal ganglia generate and use neural signals that predict reinforcement. In: Houk JC, Davis JL, Beiser DG, editors. Models of information processing in the basal ganglia. Cambridge, MA: MIT Press. p249–274.
Karmarkar UR, Buonomano DV. 2002. A model of spike-timing dependent plasticity: One or two coincidence detectors? J Neurophysiol 88:507–513.
Kerr JN, Wickens JR. 2001. Dopamine D-1/D-5 receptor activation is required for long-term potentiation in the rat neostriatum in vitro. J Neurophysiol. 85:117–124.
Kish SJ, Shannak K, Hornykiewicz O. 1988. Uneven pattern of dopamine loss in the striatum of patients with idiopathic Parkinson's disease. Pathophysiol Clin Implic 318:876–880.
Lavin A, Nogueira L, Lapish CC, Wightman RM, Phillips PE, Seamans JK. 2005. Mesocortical dopamine neurons operate in distinct temporal domains using multimodal signaling. J Neurosci 25:5013–5023.
Legenstein R, Naeger C, Maass W. 2005. What can a neuron learn with spike-timing-dependent plasticity? Neural Comput 17:2337–2382.
Lovinger DM, Tyler EC, Merritt A. 1993. Short- and long-term synaptic depression in rat neostriatum. J Neurophysiol 70:1937–1949.
Magee JC, Johnston D. 1997. A synaptically controlled, associative signal for Hebbian plasticity in hippocampal neurons. Science 275:209–213.
Mahon S, Deniau J-M, Charpier S. 2004. Corticostriatal plasticity: Life after the depression. Trends Neurosci 27:461–467.
Markram H, Lubke J, Frotscher M, Sakmann B. 1997. Regulation of synaptic efficacy by coincidence of postsynaptic Aps and EPSPs. Science 275:213–215.
Miller JD, Sanghera MK, German DC. 1981. Mesencephalic dopaminergic unit activity in the behaviorally conditioned rat. Life Sci 29:1255–1263.
Monchi O, Taylor JG, Dagher A. 2000. A neural model of working memory processes in normal subjects, Parkinson's disease, and schizophrenia for fMRI design and predictions. Neural Networks 13:953–973.
Montague PR, McClure SM, Baldwin PR, Phillips PE, Budygin EA, Stuber GD, Kilpatrick MR, Wightman RM. 2004a. Dynamic gain control of dopamine delivery in freely moving animals. J Neurosci 24:1754–1759.
Montague PR, Hyman SE, Cohen JD. 2004b. Computational roles for dopamine in behavioural control. Nature 431:760–767.
Nicholson C. 1995. Interaction between diffusion and Michaelis-Menten uptake of dopamine after iontophoresis in striatum. Biophys J 68:1699–1715.
Nicholson C. 2001. Diffusion and related transport mechanisms in brain tissue. Rep Prog Phys 64:815–884.
Partridge JG, Tang KC, Lovinger DM. 2000. Regional and postnatal heterogeneity of activity-dependent long-term changes in synaptic efficacy in the dorsal striatum. J Neurophysiol 84:1422–1429.
Pfister J-P, Toyoizumi T, Barber D, Gerstner W. 2006. Optimal spike-timing-dependent plasticity for precise action potential firing in supervised learning. Neural Comput 18:1318–1348.
Reynolds JNJ, Wickens JR. 2000. Substantia nigra dopamine regulates synaptic plasticity and membrane potential fluctuations in the rat neostriatum, in vivo. Neurosci 99:199–203.
Reynolds JNJ, Wickens JR. 2002. Dopamine-dependent plasticity of corticostriatal synapses. Neural Networks 15:507–521.
Reynolds JNJ, Hyland BI, Wickens JR. 2001. A cellular mechanism of reward-related learning. Nature 314:67–70.
Rao RPN, Sejnowski TJ. 2001. Spike-time-dependent Hebbian plasticity as temporal difference learning. Neural Comput 13:2221–2237.
Roberts PD. 1999. Computational consequences of temporally asymmetric learning rules. I. Differential hebbian learning. J Comput Neurosci 7:235–246.
Rumsey CC, Abbott LF. 2004. Equalization of synaptic efficacy and activity- and timing-dependent synaptic plasticity. J Neurophysiol 91:2273–2280.
Senn W, Markram H, Tsodyks M. 2000. An algorithm for modifying neurotransmitter release probability based on pre- and postsynaptic spike timing. Neural Comp 13:35–67.
Shaskan EG, Snyder SH. 1970. Kinetics of serotonin accumulation into slices from rat brain: Relationship to catecholamine uptake. J Pharmacol Exp Ther 175:404–418.
Schmitz Y, Lee CJ, Schmauss C, Gonon F, Sulzer D. 2001. Amphetamine distorts stimulation-dependent dopamine overflow: Effects on D2 autoreceptors, transporters, and synaptic vesicle stores. J Neurosci 21:5916–5924.
Schonfub D, Reum T, Olshausen P, Fischer T, Morgenstern R. 2001. Modelling constant potential amperometry for investigations of dopaminergic neurotransmission kinetics in vivo. J Neurosci Methods 112:163–172.
Schultz W. 1998. Predictive reward signal of dopamine neurons. J Neurophysiol 80:1–27.
Sjostrom PJ, Turrigiano GG, Nelson SB. 2001. Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. Neuron 32:1149–1164.
Spencer JP, Murphy KP. 2000. Bi-directional changes in synaptic plasticity induced at corticostriatal synapses in vitro. Exp Brain Res 135:497–503.
Song S, Abbott LF. 2001. Cortical development and remapping through spike timing-dependent plasticity. Neuron 32:339–350.
Suri RE, Schultz W. 1999. A neural network model with dopamine-like reinforcement signal that learns a spatial delayed response task. J Neurosci 91:871–890.

Sutton RS, Barto AG. 1998. Reinforcement learning: An introduction. Cambridge, MA: MIT Press.
Tang K, Low MJ, Grandy DK, Lovinger DM. 2001. Dopamine-dependent synaptic plasticity in striatum during in vivo development. Proc Natl Acad Sci USA 98:1255–1260.
Tzounopoulos T, Kim Y, Oertel D, Trussel LO. 2004. Cell-specific, spike timing-dependent plasticities in the dorsal cochlear nucleus. Nat Neurosci 7:719:725.
Venton BJ, Zhang H, Garris PA, Phillips PEM, Sulzer D, Wightman RM. 2003. Real-time decoding of dopamine concentration changes in the caudate-putamen during tonic and phasic firing. J Neurochem 87:1284–1295.
Vogels TP, Abbott LF. 2005. Signal propagation and logic gating in networks of integrate-and-fire neurons. J Neurosci 25:10786–10795.
Walsch JP. 1993. Depression of excitatory synaptic input in rat striatal neurons. Brain Res 608:123–128.
Wang SS, Denk W, Häusser M. 2000. Coincidence detection in single dendritic spines mediated by calcium release. Nat Neurosci 3:1266–1273.
West AR, Floresco SB, Charara A, Rosenkrank JA, Grace AA. 2003. Electrophysiological interactions between striatal glutamatergic and dopaminergic systems. Ann NY Acad Sci 1003:53–74.
Wickens JR, Begg AJ, Arbuthnott GW. 1996. Dopamine reverses the depression of rat corticostriatal synapses which normally follows high-frequency stimulation of cortex in vitro. Neurosci 70:1–5.
Wörgötter F, Porr B. 2005. Temporal sequence learning, prediction, and control: A review of different models and their relation to biological mechanisms. Neural Comp 17:245–319.
Xie X, Seung S. 2004. Learning in neural networks by reinforcement of irregular spiking. Phys Rev E 69:041909.
Young CE, Yang CR. 2004. Dopamine D1/D5 receptor modulates state-dependent switching of soma-dendritic Ca2+ potentials via differential protein kinase A and C activation in rat prefrontal cortical neurons. J Neurosci 24:8–23.
Zhang L, Tao HW, Holt CE, Harris WA, Poo M-M. 1998. A critical window for cooperation and competition among developing retinotectal synapses. Science 395:37–44.

APPENDIX A: DESCRIPTION OF SIMULATIONS AND PARAMETRIC VALUES

The proposed simulations involved computations of the membrane potential of two cells (cortical and striatal) as well as a measure of extracellular DA released from an adjacent SNc cell. Throughout, a time-constant of $\tau = 10$ ms was employed, implying that the dynamical values of the model were updated once every 10 ms.

Simulations using the voltage-based model were performed as follows. In order to obtain different DA concentrations, $[D(t)]_{\text{ext}}$ was varied by $[D(t)]_{\text{ext}} = t - 1$ for each discrete time-step t ranging from 1 to 200. In the spike-based model, stimulations modulated the number of consecutive time-steps during which stimulation was set to $[D(t)]_{\text{ext}} = 100\ \mu\text{M s}^{-1}$ (the default value was set to $[D(t)]_{\text{ext}} = 0.5\ \mu\text{M s}^{-1}$). In particular, HFS was obtained by a cortical stimulation and DA release at 100 Hz, both lasting one second; LFS was obtained by a single stimulation and single DA release (1 Hz). Through simulations, it was also possible to modulate the frequency of DA release independently of corticostriatal stimuli. Cortical and striatal cells were either stimulated through a pre-post (cortical followed by striatal) or a post-pre (striatal followed by cortical) spike train regime at every time-step, depending on the specific condition tested (see Methods). After the membrane potentials were computed, it was determined whether a spike was produced or not for each cell, according to whether the firing threshold was exceeded or not. If a spike was produced, the time t of spike arrival was stored. Finally, weight updates were computed.

Parameters for the voltage-based model are as follows: $\beta = 0.07$, $\tau = 10$ ms, $b = 20$, and $T_{\text{refract}} = 3$ ms. Initial weights w were set to random values between 0.0 and 1.0. Initial membrane potentials were set to $V_{\text{pre}}(t) = 1.0$ mV and $V_{\text{post}}(t) = 1.0$ mV. Parameters for the spike-based model were as follows: $\tau = 10$ ms, $\tau_S = 3.75$ ms, $b = 125$, $V_0 = 2.0$ mV; $V_{\text{rest}} = 0.5$ mV; $\theta_{\text{pre}} = 2.0$ mV, and $\theta_{\text{post}} = 1.5$ mV. Parameters for the MM equations were set to $V_{\max} = 0.08\ \mu\text{M s}^{-1}$ and $K_m = 0.3\ \mu\text{M s}^{-1}$. These values were adapted from results of striatal rat slices (Nicholson, 1995; Shaskan and Snyder, 1970). In the spike-based model, the parameters W_+ and W_- were set to fixed values of 0.2 and 0.3 respectively. The parameters τ_+, and τ_- are defined in simulations. The update rates for connection weights and DA concentrations are set to $\eta_w = 0.01$ and $\eta_D = 0.5$ respectively. Synaptic weights are bounded in the interval $[0.0, W_{\max}]$, where $W_{\max}$ is set to 1.0; this was implemented in the following way:

$$\text{if } w(t) > W_{\max}, w(t) = W_{\max},$$
$$\text{if } w(t) < 0.0, w(t) = 0.0.$$