ARTICLE Communicated by Anthony Burkitt

Stochastic Perturbation Methods for Spike-Timing-Dependent Plasticity

Todd K. Leen
leent@ohsu.edu
Robert Friel
rfriel@gmail.com
Department of Biomedical Engineering, Oregon Health & Science University, Portland, OR 97239, U.S.A.

Online machine learning rules and many biological spike-timing-dependent plasticity (STDP) learning rules generate jump process Markov chains for the synaptic weights. We give a perturbation expansion for the dynamics that, unlike the usual approximation by a Fokker-Planck equation (FPE), is well justified. Our approach extends the related system size expansion by giving an expansion for the probability density as well as its moments. We apply the approach to two observed STDP learning rules and show that in regimes where the FPE breaks down, the new perturbation expansion agrees well with Monte Carlo simulations. The methods are also applicable to the dynamics of stochastic neural activity. Like previous ensemble analyses of STDP, we focus on equilibrium solutions, although the methods can in principle be applied to transients as well.

1 Introduction

Both online machine learning algorithms and biological learning mediated by spike-timing-dependent plasticity (STDP) generate Markov processes. While the average dynamics of the synaptic strengths may describe salient phenomena associated with STDP, they are incomplete. A complete account specifies the probability densities for the synaptic weights. Similarly, the dynamics and function of stochastic neural activity can be described in mean terms by firing rates. However this addresses neither fluctuations (Cowan, 1991; Ohira & Cowan, 1993; Buice, Cowan, & Chow, 2010) nor important correlations (such as firing synchrony; (Gray & Singer, 1989). As for synaptic dynamics, stochastic neural activity dynamics requires a probabilistic description.

The goal in analysis of both machine and biological learning is to discern how learning rules affect circuit level and behavioral function. (Machine learning has a synthesis component; we design algorithms to achieve particular ends. But analysis is part of the design process.) The average dynamics,

obtained by replacing the learning rule by its expectation, may adequately describe salient phenomena. However, the inherent fluctuations in noisy learning create scatter around the average dynamics. (Furthermore, in nonlinear systems, the fluctuations couple to the mean behavior and can shift the average dynamics. See the example in section 3.2.2.)

At convergence, scatter about the optimal weights compromises the ability of the system to carry out its functional goals accurately. This is well understood in adaptive filter theory, where the trade-off between adaptation speed and asymptotic error in the weights is discussed (White, 1989; Kuan & Hornik, 1991; Haykin, 1991). Stochastic approximation methods anneal the learning rate (i.e., they reduce the step size) at late times in order to damp out the fluctuations. The annealing schedule is designed so that the parameters converge in mean square and with probability one to local optima (Robbins & Monro, 1951; Kushner & Clark, 1978; Fabian, 1968; Ljung, 1977). There is at present no evidence that biological systems governed by STDP contain a mechanism for modulating the step size to control equilibrium fluctuations. There is experimental evidence in mormyrid electrosensory lateral line lobe (ELL) that the probability of broad dendritic spikes, which mediate plasticity, is modulated by descending control (N. Sawtell, personal communication, 2010). However such modulation would not change the width of the asymptotic distribution as would a modulation of the step size.

Late in the 1990s, neuroscientists discovered that synaptic plasticity is strongly dependent on the relative timing between pre- and postsynaptic events (Bell, Han, Sugawara, & Grant, 1997; Markram, Lubke, Fortscher, & Sakmann, 1997; Bi & Poo, 1998; Han, Grant, & Bell, 2000). The timing dependence observed in such spike-timing-dependent-plasticity (STDP) renders obsolete previous models in which synaptic change depends only on spike rates in the pre- and postsynaptic cells (Sejnowski, 1977; Bienenstock, Cooper, & Munro, 1982).

While rate-dependent models of plasticity involve many spike events over which variations are naturally averaged, spike-timing-dependent models involve only a few (typically two) presynaptic and postsynaptic events, leaving random variations important. Variability in the amount of potentiation or depression observed at fixed relative timing between the pre- and postsynaptic events contributes to these random variations (van Rossum, Bi, & Turrigiano, 2000). At the circuit level, fluctuations in relative timing between pre- and postsynaptic events, arising, for example, from uncontrolled inputs to cells and expressed through the firing probability, contribute to random variations as well (Roberts & Bell, 2000). Stochastic frameworks are hence natural for mathematical modeling of STDP and its consequences in adaptive neural circuitry. For many STDP learning rules reported in the literature, the dynamics instantiate a Markov process (Bell, Han et al., 1997; Markram et al., 1997; Bi & Poo, 1998; Han et al., 2000; van Rossum et al., 2000); changes in the synaptic weight depend through the learning rule only on the current weight and a set of random variables that determine the transition probabilities.

Early experiments revealing STDP showed the effects of relative timing between pairs of pre- and postsynaptic events. More recent experiments show exceptions in the pairwise spike dependency. Pfister and Gerstner (2006a, 2006b) discuss experimental evidence for effects of event triplets and firing rates. They propose a (deterministic) modeling framework in which the weight change is a functional of the pre- and postsynaptic spike trains and explore specific (low-dimensional, phenomenological) mathematical models based on pairwise and triplet interactions. Morrison, Diesmann, and Gerstner (2008, and references therein) review a range of STDP phenomenology—including pairwise, triplet, and quadruplet interactions and effects of average firing rates—and discuss (low-dimensional, deterministic) modeling approaches.

In addition to multiple-spike dependence, there are complicating intracellular effects. Ismailov, Kalikulov, Inoue, and Friedlander (2004) found that otherwise similar pyramidal cells in slice preparations from guinea pig visual cortex exhibit different calcium kinetic profiles and that these differences reliably predict whether the cells undergo potentiation or depression when subject to identical spike pairing protocols. If the calcium kinetic profile within a given cell can change in time, the resulting transition probabilities would depend on this additional state information, complicating the description.

In this article, we restrict attention to pairwise rules that establish a Markov process. The dynamics of such systems are described by a master equation (for continuous-time formulations) or by a Chapman-Kolmogorov equation (for discrete-time formulations; Gardiner, 2009). Unfortunately, these formulations are intractable for all but the simplest cases. The most common analytic approach is to use, in their place, a Fokker-Planck equation (Risken, 1989; Gardiner, 2009).

The Fokker-Planck equation (FPE) has been successfully applied to both machine and biological learning. Radons, Schuster, and Werner (1990) and Radons (1993) used the FPE to describe learning in backpropagation networks, including noise-induced hopping between basins of local optima, as did Orr and Leen (1993) and Hansen (1993). Der and Villmann (1993) successfully applied the FPE to describe the dynamics of Kohonen's self-organizing feature map, and Orr and Leen (1993) and Orr (1996) considered the behavior of the related K-means algorithm. In stochastic approximation algorithms, for which the learning rate is annealed, an FPE (with linear drift and constant diffusion—what van Kampen, 2007, refers to as the linear noise regime), is useful for analysis of asymptotic convergence. This approach to analyzing annealed algorithms, adopted ad hoc by Orr (1996) and rigorously derived from a fluctuation expansion by Leen (1998) and Leen, Schottky, and Saad (1998, 1999), provides a simple derivation of the classical theorems on asymptotic normality and the consequent bounds on convergence rates to local optima (Venter, 1967; Fabian, 1973). Results derived from such an FPE led to elegant and computationally efficient enhancements to stochastic gradient descent that automatically provide the

optimal asymptotic convergence rates (Leen & Orr, 1994; Orr & Leen, 1997). These algorithms, suggested directly by the ensemble dynamics, achieve optimal convergence rates more effectively and elegantly than attempts spanning several decades (Venter, 1967; Fabian, 1973; Darken & Moody, 1991, 1992).

More recently, theoretical neuroscientists have used the FPE to describe the dynamics of learning in systems governed by STDP. Câteau and Fukai (2003) use the FPE to determine the equilibrium distribution of synaptic strengths under learning rules qualitatively similar to those observed in rat hippocampus CA1 neurons. Van Rossum et al. (2000) use an equilibrium distribution obtained analytically from an FPE to predict the effects of the temporally antisymmetric spike-timing learning rule observed by Bi and Poo (1998). Their calculations give a unimodal but strongly skewed (nongaussian) distribution that agrees well with their Monte Carlo simulations and agrees qualitatively with experimental quantal amplitude distributions observed in pyramidal neurons. Rubin, Lee, & Sompolinsky (2001) also use the FPE to study temporally antisymmetric learning with a specific focus on how the way in which bounds on synaptic weights are expressed change the equilibrium distribution. Kepecs, van Rossum, Song, & Tegney (2002) use the FPE to reason about the effect of learning rules on equilibrium weight distributions arising from STDP. Masuda and Aihara (2004) use the FPE to analyze synaptic competition and the formation of functional clusters. Burkitt, Meffin, & Grayden (2004) use the FPE to predict equilibrium synaptic distributions for dynamics with a stable fixed point, and show good agreement with Monte-Carlo simulations. More recently, Babadi and Abbott (2010) used an FPE to analyze stability of STDP learning.

Despite its apparent utility, van Kampen claims in his seminal work (1976, 2007) that adoption of the FPE to approximate a jump process (as in stochastic learning) is deeply flawed unless rigorously derived by a proper limiting procedure. He says of its solutions that "any features they contain beyond that [predicted by linear drift and constant diffusion] are spurious" (van Kampen, 2007, p. 262) and that the general FPE does not give a proper stationary (equilibrium) solution. An FPE with linear drift and constant diffusion (what is called the linear noise regime) generates a gaussian equilibrium distribution. The skewed equilibrium distributions of synaptic weights predicted by the FPE for STDP learning rules with stable fixed points (van Rossum et al., 2000; Burkitt et al., 2004, for example) are, according to van Kampen, spurious—despite their agreement with Monte Carlo simulations. Similarly FPE predictions for basin hopping in machine learning depend strongly on nonlinearities in the drift coefficient (Radons et al., 1990; Radons, 1993; Orr & Leen, 1993; Hansen, 1993), and so are, according to van Kampen, also spurious—despite their agreement with Monte Carlo simulations. Heskes and Kappen (1993) brought van Kampen's arguments to the machine learning community and essentially put an end to the use of the nonlinear FPE in application to online learning.

Following van Kampen, they argue that only the lowest-order gaussian approximation to the FPE solution has any validity. However, they do not develop an extended analysis that goes beyond the gaussian approximation.

The FPE shows good predictive utility of distribution features that van Kampen claims are not trustworthy. His pronouncements are clearly heavy-handed. In our view, the fundamental problem is not that the FPE is not a useful approximation to a master equation, but rather that there is no apparent way to use it as the lowest-order approximation to a complete perturbation expansion. Such an expansion would enable one to correct its shortcomings and discuss its domain of applicability. Van Kampen (1976, 2007) derived his system-size expansion as a rigorous alternative to the FPE for predicting the dynamics of a master equation. However, it is underused and poorly understood (particularly beyond the lowest-order gaussian piece), in part because the development is difficult. Furthermore the higher-order corrections in the system size expansion address only the moments; van Kampen's approach does not give approximations for the probability density, which would provide useful visualization.

The motivation behind this article is twofold: (1) to provide rigorous and clearly derived computational tools to approximate probability distributions and moments for STDP learning rules and (2) to begin an even-handed critique of the FPE so that its application to theoretical neuroscience is not undermined as it was in machine learning. We give a perturbation approach appropriate for machine and STDP learning (with constant learning rates) and for stochastic neurodynamics. Our approach gives a clear, conceptually simple development and produces an expansion for the density as well as its moments. Our equations for the moments match van Kampen's (2007), but the density results here are novel, and to our knowledge, similar methods have not previously appeared in the physics literature. Departing with much of the stochastic process literature, we use a discrete time approach to avoid any possible confusion produced by regarding synaptic update processes in continuous time. (Our perturbation expansion is valid in continuous time with minimal changes, and of course the equilibrium distributions for both discrete and continuous time processes are identical.)

We apply our expansions to two observed STDP learning rules: the antisymmetric Hebbian learning observed by Bi and Poo (1998) and modeled by van Rossum et al. (2000) and the anti-Hebbian learning observed in weakly electric fish by Bell, Han et al. (1997), and modeled at the circuit level by Roberts and Bell (2000). For the Hebbian rule, we compare equilibrium densities and moments predicted by our perturbation expansion with those from the FPE, and with histograms from Monte Carlo simulations. For the physiological learning rate parameters, the FPE agrees very well with simulations. Large errors in the FPE predictions require potentiation and depression rate constants 100 times larger than physiologically observed. In this regime, the perturbation expansion agrees well with the simulations. In contrast to the van Rossum model, in the electric fish anti-Hebbian learning

rule, the FPE breaks down at the physiologically observed learning rates. The FPE fails to accurately predict the mean and variance and predicts the wrong sign for the skew. The perturbation expansion predicts the moments quite well.

2 Noisy Learning

2.1 Ensemble Dynamics. Machine learning by online algorithms and many observed biological learning rules mediated by spike-timing-dependent plasticity (STDP) can both be described by update equations of the form

$$w_{t+1} = w_t + \eta h(w_t, \psi_t), \tag{2.1}$$

where w_t is the synaptic weight at step t, η is the learning rate (constant for this discussion), $h(\cdot)$ is the learning rule, and ψ is a stochastic variable giving rise to the randomness in the dynamics. In the machine learning context, ψ_t might be the example chosen from the data to feed into the algorithm at time step t (Bishop, 2006). In STDP, randomness can arise out of fluctuations in the relative timing between the presynaptic and postsynaptic events or by noise present in the synaptic changes even at fixed relative timing. The vector variable ψ_t collectively carries all such effects.

The update equation, 2.1, generates a Markov process on w; the weight w_{t+1} depends on only the weight w_t and the probability of the transition $w_t \to w_{t+1}$ (itself dependent on w_t and possibly on time). The probability density for the weights $P(w, t)$ thus evolves according to the Chapman-Kolmogorov equation (Gardiner, 2009),

$$P(w, t+1) = \int W(w|w')P(w', t)dw', \tag{2.2}$$

where $W(w|w')$ is the probability of making the transition from w' to w in one time step. The transition probability for fixed ψ is a Dirac-delta function satisfying the update equation, 2.1,

$$W(w|w', \psi) = \delta(w - w' - \eta h(w', \psi)),$$

and we obtain the unconditional transition probability by taking its expectation with respect to ψ:

$$\begin{aligned} W(w|w') &= E_\psi[W(w' \to w|\psi)] \\ &= E_\psi[\delta(w - w' - \eta h(w', \psi))]. \end{aligned} \tag{2.3}$$

The Kramers-Moyal (KM) expansion (Gardiner, 2009) of the Chapman-Kolmogorov equation, 2.2, is an infinite-order differential-difference equation that will serve as the starting point for our perturbation treatment and for comparison with the commonly adopted Fokker-Planck equation (FPE). Gardiner (2009) gives a derivation of the KM expansion, but for completeness we give a sketch here. To derive the expansion, substitute the transition probability, equation 2.3, into the Chapman-Kolmogorov equation, equation 2.2, leaving

$$P(w, t+1) = \int E_\psi \left[\delta(w - w' - \eta h(w', \psi))\right] P(w', t) dw'. \tag{2.4}$$

Next, Taylor-expand this about $\eta = 0$. The expansion of the Dirac delta function is handled using the theory of generalized functions (Gelfand & Shilov, 1964). Equivalently, multiply the result, equation 2.4, by a smooth test function of compact support (bump function) $\kappa(w)$ and integrate over w, leaving

$$E[\kappa, t+1] = \int E_\psi \left[\kappa(w' + \eta h(w', \psi)\right] P(w', t)\, dw'. \tag{2.5}$$

Next, Taylor-expand about $\eta = 0$,

$$\begin{aligned} E[\kappa, t+1] &= \int \sum_{j=0}^{\infty} \frac{\eta^j}{j!} E_\psi[h^j(w', \psi)] \frac{\partial^j \kappa(w')}{\partial w'^j} P(w', t) dw' \\ &= E[\kappa, t] + \int \sum_{j=1}^{\infty} \frac{\eta^j}{j!} \alpha_j(w') \frac{\partial^j \kappa(w')}{\partial w'^j} P(w', t) dw', \end{aligned} \tag{2.6}$$

having extracted the $j = 0$ term from the summation and defined the jump moments,

$$\alpha_j(w) \equiv E_\psi[h^j(w, \psi)]. \tag{2.7}$$

Next, integrate equation 2.6 by parts to take all derivatives off $\kappa(w')$ and rename the integration variable $w' \to w$,

$$\begin{aligned} E\left[\kappa(w), t+1\right] - E\left[\kappa(w), t\right] &\equiv \int \kappa(w) \left(P(w, t+1) - P(w, t)\right) dw \\ &= \int \kappa(w) \left(\sum_{j=1}^{\infty} \frac{(-\eta)^j}{j!} \partial_w^j \left(\alpha_j(w) P(w, t)\right) \right) dw, \end{aligned} \tag{2.8}$$

where $\partial_w^j \equiv \partial^j / \partial w^j$. Finally, use the fact that the support of κ is arbitrary to obtain the final form of the expansion:

$$P(w, t+1) - P(w, t) = \sum_{j=1}^{\infty} \frac{(-\eta)^j}{j!} \partial_w^j (\alpha_j(w) P(w, t)). \tag{2.9}$$

The first jump moment $\alpha_1(w)$, called the drift coefficient, captures the average dynamics.[1] The second jump moment $\alpha_2(w)$, called the diffusion coefficient, describes the tendency of the system to spread out over the configuration space.

The KM expansion, equation 2.9, is usually intractable. The common procedure is to retain only the first two terms of the infinite series, leaving the Fokker-Planck equation (FPE):

$$\begin{aligned} P(w, t+1) - P(w, t) = &-\eta \partial_w \left(\alpha_1(w) P(w, t)\right) \\ &+ \frac{\eta^2}{2} \partial_w^2 \left(\alpha_2(w) P(w, t)\right). \end{aligned} \tag{2.10}$$

The truncation of the KM expansion to the FPE is convenient and widely used for several reasons. First, the Pawula theorem (Risken, 1989) ensures that the resulting density will be nonnegative (which is not the case if a finite number of terms greater than two are retained from the KM expansion). Second, for one-dimensional problems, the FPE can be solved in closed form for the equilibrium probability density function, giving a useful visualization. Third, there is a large literature on solutions of the FPE.

However, truncation to the FPE is strictly justified only when arrived at by a limiting procedure in which the higher-order jump moments $\alpha_3, \alpha_4, \ldots$ necessarily vanish. This is the case, for example, in the derivation of Brownian motion from a constant-step-size random walk (Gardiner, 2009), so the FPE is valid there. In machine and biological learning systems (and in the treatment of fluctuations in neural activity), the truncation is generally not justified.

As an alternative to the FPE equation, we develop a perturbation expansion for the dynamics of $P(w, t)$. Our treatment follows the roots of van Kampen's (1976, 2007), system size expansion and so gives the same results for the moments. However, we give a much clearer development than previously available. We believe that the failure to adopt these expansions routinely is the result, in part, of confusing expositions.

[1] Indeed if $\alpha_1(w)$ is linear in w, the motion of the mean follows $E[w](t+1) = E[w](t) + \eta \alpha_1(E[w](t))$. If the drift is nonlinear, the motion of the mean depends on higher-order moments of the distribution.

Our perturbation solution goes beyond the moment expansions provided by van Kampen, giving approximations for the density as well. Thus, we provide visualization of the density on a more rigorous footing than the FPE. Finally, we retain a discrete-time formulation (rather than pass to continuous time, as is usually done), which may be easier to interpret with respect to transients in biological and machine learning. (The equilibrium distributions for both the continuous and discrete time formulations are identical.)

2.2 The Fluctuation Expansion. Following van Kampen, we decompose the weight trajectory w_t into the sum of a deterministic and a fluctuating piece,

$$w_t = \phi_t + \sqrt{\eta}\xi_t, \tag{2.11}$$

where ϕ_t is the deterministic trajectory and ξ_t are the fluctuations. In analogy with a simple random walk, we assume that the ensemble variance of w scales linearly with the step size η, hence the $\sqrt{\eta}$ multiplying ξ in equation 2.11. This scaling assumption is ultimately verified by confirming that the lowest-order contribution to the variance of ξ is indeed independent of η. In appendix A, we motivate the decomposition, equation 2.11, from the desire to build a complete perturbation expansion around the equilibrium solution to a Fokker-Planck equation. Van Kampen (1976, 2007) developed the decomposition, equation 2.11, to address fluctuations in chemical systems with many, but not infinitely many, molecules. In that framework, the variance of the fluctuations scales inversely with the volume—the system size. For us, the fluctuations arise from the random influences ψ in equation 2.1, and we refer to the procedure as a fluctuation expansion.

We retain a discrete-time dynamical description complementary to the continuous-time version usually found in the literature. One can pass from the discrete-time random walk to a continuous-time master equation, in which $P(w, t+1) - P(w, t)$ is replaced by $dP(w, t)/dt$, by assuming the discrete-time steps are Poisson distributed in continuous time t (Bedeaux, Lakatos-Lindenberg, & Shuler, 1971). That is, the number of update steps taken in the continuous time interval $[0, t]$ is a random variable following the Poisson distribution. In the discrete-time formulation, the number of update steps is not random, but given deterministically by integer t. Under Poisson statistics, after a time period long enough to accumulate a few hundred to a few thousand update steps, the mean and actual number of updates will differ by only a few percent, and the discrete- and continuous-time formulations will be in good agreement.

To begin, we separate the dynamics of the deterministic trajectory ϕ_t from those of the fluctuations ξ_t. To achieve this, substitute the decomposition in

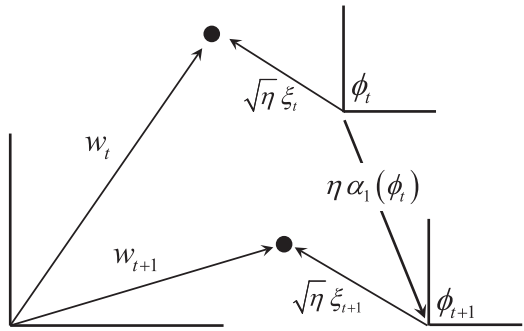

Figure 1: Coordinate transformation from the (fixed) w-frame to the moving ξ-frame.

equation 2.11 into the update equation, 2.1, leaving

$$\phi_{t+1} + \sqrt{\eta}\xi_{t+1} = \phi_t + \sqrt{\eta}\xi_t + \eta h(\phi_t + \sqrt{\eta}\xi_t, \psi_t). \tag{2.12}$$

Next, expand the last term on the right-hand side of equation 2.12 in powers of $\sqrt{\eta}$, and take the expectation of both sides with respect to ψ yielding

$$\phi_{t+1} + \sqrt{\eta}E[\xi_{t+1}|\xi_t] = \phi_t + \sqrt{\eta}\xi_t + \eta\alpha_1(\phi_t) + \mathcal{O}(\eta^{3/2}\xi_t), \tag{2.13}$$

having used the definition of α_1 from equation 2.7. Extracting the pieces of equation 2.13, that are independent of ξ_t yields the equation for the deterministic trajectory:

$$\phi_{t+1} = \phi_t + \eta\alpha_1(\phi_t). \tag{2.14}$$

Finally, to obtain the equation of motion for the fluctuations, substitute equation 2.14 into equation 2.12, yielding

$$\xi_{t+1} = \xi_t + \sqrt{\eta}\left[h(\phi_t + \sqrt{\eta}\xi_t, \psi_t) - \alpha_1(\phi_t)\right] \tag{2.15}$$

$$\equiv \xi_t + \sqrt{\eta}H(\phi_t + \sqrt{\eta}\xi_t, \phi_t, \psi_t), \tag{2.16}$$

where we have defined the fluctuation update function:

$$H(w, \phi, \psi) \equiv h(w, \psi) - \alpha_1(\phi). \tag{2.17}$$

The dynamics expressed by equation 2.1 are related to those expressed by equations 2.14 and 2.16 by the change of coordinates in equation 2.11. Figure 1 illustrates this. The coordinate frame in which w_t is measured is fixed. The deterministic trajectory ϕ_t defines the origin of the ξ-coordinates. This origin moves in time driven by $\alpha_1(\phi_t)$ according to equation 2.14. The fluctuation vectors $\sqrt{\eta}\xi_t$ are referenced to this moving frame. The figure

should help make clear that the term $-\alpha_1(\phi_t)$ in the equation of motion for the fluctuations, equation 2.15, compensates for the motion of the origin of ξ-coordinates. (For example, if $w_{t+1} = w_t + \eta\alpha_1(\phi_t)$, then the fluctuation is unchanged $\xi_{t+1} = \xi_t$.)

We can obtain a K-M expansion for the evolution of the probability density on the fluctuations ξ by steps analogous to those leading to equation 2.9. The result is

$$P(\xi, t+1) - P(\xi, t) = \sum_{j=1}^{\infty} \frac{(-\eta)^{j/2}}{j!} \partial_\xi^j \times (A_j(\phi_t + \sqrt{\eta}\xi, \phi_t) P(\xi, t)), \tag{2.18}$$

where the fluctuation jump moments are

$$A_j(w, \phi) = E_\psi[H^j(w, \phi, \psi)] = E_\psi[\{h(w, \psi) - \alpha_1(\phi)\}^j]. \tag{2.19}$$

(Compare equations 2.7 and 2.19.)

We complete the derivation by expanding the jump moments $A_j(w = \phi + \sqrt{\eta}\xi, \phi)$ on the right-hand side of equation 2.18 about $w = \phi$ and resumming the series to find

$$P(\xi, t+1) - P(\xi, t) = \sum_{k=2}^{\infty} \sum_{j=1}^{k} \frac{(-1)^j}{j!(k-j)!} \eta^{k/2} A_j^{(k-j)} \partial_\xi^j \times (\xi^{k-j} P(\xi, t)), \tag{2.20}$$

where we have used the definitions in equations 2.7, 2.17, and 2.19 to find

$$A_1(w = \phi, \phi) = E_\psi[H(w = \phi, \phi, \psi)] = E_\psi[h(\phi, \psi) - \alpha_1(\phi)] \equiv 0$$

and have defined

$$A_j^{(m)} \equiv \left. \frac{\partial^m A_j(w, \phi_t)}{\partial w^m} \right|_{w=\phi_t}, \tag{2.21}$$

the coefficients of the expansion of the jth jump moment about ϕ_t. Equations 2.14 and 2.20 together provide a complete description of the dynamics. For convenience, we reproduce them here:

$$\phi_{t+1} = \phi_t + \eta\alpha_1(\phi_t) \tag{2.22}$$

$$P(\xi, t+1) - P(\xi, t) = \sum_{k=2}^{\infty} \sum_{j=1}^{k} \frac{(-1)^j}{j!(k-j)!} \eta^{k/2} A_j^{(k-j)} \partial_\xi^j (\xi^{k-j} P(\xi, t)). \tag{2.23}$$

The right-hand side of equation 2.23 is identical to that in the continuous-time system size expansion (Gardiner, 2009) apart from the substitution of the jump moments of equation 2.19 for those in equation 2.7.

2.3 Perturbation Solution for the Density and Its Moments. Here we give simple, clear, and direct perturbation expansions for the probability density of the fluctuations and its moments. For the remainder of the article, we confine attention to equilibrium solutions, for which

$$\begin{aligned}\phi_{t+1} &= \phi_t \equiv \phi_* \\ P(\xi, t+1) &= P(\xi, t) \equiv P(\xi).\end{aligned}$$

For equilibria, the left-hand side of equation 2.23 is zero, and we rewrite the right-hand side as a sum of differential operators times powers of η,

$$\eta \left(L_0 + \eta^{1/2} L_1 + \eta L_2 + \ldots\right) P(\xi) = 0, \tag{2.24}$$

where the operators L_j are defined by

$$L_p q(\xi) \equiv \sum_{j=1}^{p+2} \frac{(-1)^j}{j!(p+2-j)!} A_j^{(p+2-j)} \partial_\xi^j \left(\xi^{p+2-j} q(\xi)\right). \tag{2.25}$$

For reference, the first three are:

$$L_0 q(\xi) = -A_1^{(1)} \partial_\xi (\xi q) + \frac{1}{2} A_2^{(0)} \partial_\xi^2 q \tag{2.26}$$

$$L_1 q(\xi) = -\frac{1}{2} A_1^{(2)} \partial_\xi (\xi^2 q) + \frac{1}{2} A_2^{(1)} \partial_\xi^2 (\xi q) - \frac{1}{3!} A_3^{(0)} \partial_\xi^3 q \tag{2.27}$$

$$\begin{aligned}L_2 q(\xi) = &-\frac{1}{3!} A_1^{(3)} \partial_\xi (\xi^3 q) + \frac{1}{4} A_2^{(2)} \partial_\xi^2 (\xi^2 q) \\ &-\frac{1}{3!} A_3^{(1)} \partial_\xi^3 \left(\xi q\right) + \frac{1}{4!} A_4^{(0)} \partial_\xi^4 q.\end{aligned} \tag{2.28}$$

The system can settle into an equilibrium in an unbounded weight space when the deterministic trajectory in equation 2.22 approaches a fixed point ϕ_*,

$$\alpha_1(\phi_*) = 0, \tag{2.29}$$

that is asymptotically stable (Guckenheimer & Holmes, 1983):[2]

$$-2 < \eta \left.\frac{d\alpha_1}{dw}\right|_{w=\phi_*} < 0. \tag{2.30}$$

There may be multiple stable fixed points $\phi_i^*, i = 1 \ldots n_{fp}$ (with unstable fixed points in between) that satisfy the conditions of equations 2.29 and 2.30. In this case, equilibria with well-separated modes can be approximated by mixtures of the perturbation density we develop below.[3]

There are other types of equilibria as well. In systems with constraints bounding the weights $w \in \mathcal{R} \equiv [w_{\min}, w_{\max}]$ and no stable fixed point in $\mathcal{R}$, equilibria may be found with density peaks near either or both bounds depending on α_1. If $\alpha_1 > 0, \forall w \in \mathcal{R}$, the upper bound $w_{\max}$ is a stable fixed point (of ϕ_t), and the equilibrium density will be peaked up about $w_{\max}$. Similarly, if $\alpha_1 < 0, \forall w \in \mathcal{R}$, the lower bound $w_{\min}$ is a stable fixed point, and the equilibrium density will be peaked up about $w_{\min}$. If there is a single unstable fixed point at $w_* \in \mathcal{R}$, with $\alpha_1(w_*) = 0$, and $\alpha_1^{(1)}(w) > 0, \forall w \in \mathcal{R}$ (as, e.g., in Câteau & Fukai, 2003, and Kepecs et al., 2002), then α_1 drives density to both boundaries, and the equilibrium will be peaked up at both boundary points. This description of equilibrium types in terms of the fixed points of the drift $\alpha_1(w)$ is fundamental. It accounts for the effect of the various multiplicative and additive learning rules (together with spike timing distributions) on equilibrium distributions discussed by Kepecs et al. (2002), Câteau and Fukai (2003), Burkitt et al. (2004), and Morrison et al. (2008).

In what follows, we assume that there is a single, isolated, asymptotically stable fixed point $\alpha(\phi_*) = 0$ with no constraints that bound the synaptic weight. Note that from equations 2.17 and 2.29, at $\phi = \phi_*, H = h$. It follows from equations 2.7 and 2.19 that $A_j(w) = \alpha_j(w)$ at $w = \phi_* + \sqrt{\eta}\xi$. Consequently, in the operators of equation 2.25, the jump moment derivatives are identical to those for the continuous time case $A_j^{(i)} = \alpha_j^{(i)}$, and the equilibria of the discrete-time and continuous-time systems are identical, as one would expect.

[2]For a continuous-time system, the discrete-time map, equation 2.22, is replaced by a differential equation, and the condition for asymptotic stability of the fixed point is familiar $\alpha_1'(\phi_*) < 0$. For a discrete-time map, as here, the linearization gives $(\phi_t - \phi_*) = (1 + \eta\alpha_1'(\phi_*))^t(\phi_0 - \phi_*)$ and convergence of the linearization (the condition for asymptotic stability) requires $\lim_{t\to\infty}(1 + \eta\alpha_1'(\phi_*))^t = 0$, hence equation 2.30.

[3]If several of the fixed points ϕ_i^* are absorbing states, then the proportions of each mode in the equilibrium mixture will depend on the initial conditions. See, for example, Orr and Leen (1993).

To develop a perturbation expansion for the equilibrium density satisfying equation 2.24, we write

$$P(\xi) = P^{(0)}(\xi) + \eta^{1/2}P^{(1)}(\xi) + \eta P^{(2)}(\xi) + \ldots, \tag{2.31}$$

where the $P^{(k)}(\xi)$ are unknown functions (independent of η) that we will solve for. This decomposition, equation 2.31, and the operator expansion in equation 2.24, are the critical elements that provide a clear development for the approximations to the density and its moments.

To solve for the $P^{(k)}(\xi)$ we substitute the series in equation 2.31 into 2.24 and collect terms of like powers of $\eta^{1/2}$ to obtain

$$\begin{aligned} &L_0P^{(0)} + \eta^{1/2}\left(L_1P^{(0)} + L_0P^{(1)}\right) \\ &\quad + \eta\left(L_2P^{(0)} + L_1P^{(1)} + L_0P^{(2)}\right) + \mathcal{O}(\eta^{3/2}) = 0. \end{aligned} \tag{2.32}$$

Since equation 2.32 holds regardless of the value of η, it must be true that the terms of equal order in η separately vanish. Hence we obtain a succession of equations:

$$L_0P^{(0)}(\xi) = 0, \tag{2.33}$$

$$L_0P^{(1)}(\xi) = -L_1P^{(0)}(\xi), \tag{2.34}$$

$$L_0P^{(2)}(\xi) = -L_1P^{(1)}(\xi) - L_2P^{(0)}(\xi). \tag{2.35}$$

$$\vdots$$

We are going to solve each of equations 2.33 to 2.35, and so forth, in order. The zeroth-order density $P^{(0)}(\xi)$ satisfies equation 2.33, which carries the linear part of the drift and the constant part of the diffusion coefficients (see equation 2.26). This describes the equilibrium of an Ornstein-Uhlenbeck process (Gardiner, 2009). The solution is a gaussian with mean zero and variance

$$\sigma_0^2 = \frac{\alpha_2^{(0)}(\phi_*)}{2|\alpha_1^{(1)}(\phi_*)|}. \tag{2.36}$$

This is the lowest-order piece—what van Kampen (2007) calls the linear noise approximation—of the expansion. By equation 2.31, in the limit $\eta \to 0$, the density is dominated by $P^{(0)}(\xi)$. Thus, the equilibrium solution about an asymptotically stable fixed point ϕ_* for any learning rule approaches this gaussian at low learning rates. Finally, since this lowest-order approximation to the variance, equation 2.36, is independent of η, the scaling of

fluctuations in the original decomposition is justified, as per the paragraph immediately following equation 2.11.

2.3.1 Perturbation Corrections to the Density. To obtain the higher-order corrections to the density $P^{(i)}(\xi)$, $i = 1, 2, \ldots$ we require the eigenfunctions of L_0 and its adjoint,[4]

$$L_0 f_k(\xi) = \lambda_k f_k(\xi), \tag{2.37}$$

$$L_0^\dagger g_k(\xi) = \lambda_k g_k(\xi), \tag{2.38}$$

$$\lambda_k = -k\left|\alpha_1^{(1)}\right|, \quad k = 0, 1, 2, \ldots. \tag{2.39}$$

The eigenfunctions of the adjoint operator $L_0^\dagger$ are Hermite polynomials,

$$g_k(\xi) = (2^k k!)^{-1/2} H_k\left(\sqrt{\frac{|\alpha_1^{(1)}|}{\alpha_2^{(0)}}}\xi\right), \tag{2.40}$$

(recall $\alpha_1^{(1)} < 0$), where (Abramowitz & Stegun, 1972)

$$H_k(x) \equiv e^{x^2}\left(-\frac{d}{dx}\right)^k e^{-x^2}. \tag{2.41}$$

The eigenfunctions of L_0 are Hermite polynomials times the gaussian stationary solution,

$$f_k(\xi) = \frac{1}{\sqrt{2\pi\sigma_0^2}} \exp\left(-\frac{x^2}{2\sigma_0^2}\right) g_k(\xi), \tag{2.42}$$

with σ_0^2 from equation 2.36. The forward and adjoint eigenfunctions form a complete, biorthogonal set,

$$(g_j, f_k) = \delta_{jk}, \tag{2.43}$$

under the standard $\mathcal{L}_2$ inner product (Gardiner, 2009). The f_k are similar to quantum mechanical harmonic oscillator wave functions. The first of these is the zeroth-order gaussian density $P^{(0)}(\xi) = f_0(\xi)$.

[4]Let (q, r) be the inner product between functions q and r. The adjoint $L^\dagger$ of an operator L satisfies $(q, Lr) = (L^\dagger q, r)$ for all q, r for which the inner product exists. For the standard $\mathcal{L}_2$ inner product used here, the adjoint is recovered through integration by parts. We write the operator $L_0^\dagger$ explicitly in equation 2.52.

To find the first order (in $\eta^{1/2}$) correction to the density, we expand it as a linear combination of eigenfunctions of L_0:

$$P^{(1)}(\xi) = \sum_{j=1}^{\infty} a_j f_j(\xi). \tag{2.44}$$

To find the coefficients a_j, first substitute this expansion into equation 2.34 and use the fact that the $f_j(\xi)$ are eigenfunctions of L_0, equation 2.37, to obtain

$$\sum_{j=1}^{\infty} a_j \lambda_j f_j = -L_1 P^{(0)}. \tag{2.45}$$

Next take the inner product of both sides with g_k and use the biorthogonality, equation 2.43, to find

$$a_k = -\frac{1}{\lambda_k} \left(g_k, L_1 P^{(0)} \right). \tag{2.46}$$

Thus, the first order (in $\eta^{1/2}$) correction to the density is

$$P^{(1)}(\xi) = -\sum_{k=1}^{\infty} \frac{1}{\lambda_k} \left(g_k, L_1 P^{(0)} \right) f_k(\xi). \tag{2.47}$$

With this correction in hand, the next is found following the same procedure on equation 2.35. The result is

$$P^{(2)}(\xi) = -\sum_{k=1}^{\infty} \frac{1}{\lambda_k} \left(g_k, L_1 P^{(1)} + L_2 P^{(0)} \right) f_k(\xi). \tag{2.48}$$

As we shall see, in practice, the sums over k in equations 2.47 and 2.48, and the higher-order analogs terminate at a finite number of terms. Thus, we can solve for the corrections $P^{(n)}$ at any order n with this procedure.

Once the desired correction terms are computed, the perturbation density $P(\xi)$ in equation 2.31 is assembled. Finally, the density in w is recovered from $P(\xi)$ as

$$P_w(w) = P\left(\xi = \frac{w - \phi_*}{\sqrt{\eta}} \right) \left| \frac{d\xi}{dw} \right|. \tag{2.49}$$

Calculating the corrections to the density requires evaluating inner products of the form

$$(g_k, L_p f_j) \equiv \int g_k(\xi) L_p f_j(\xi) d\xi.$$

We do not need to explicitly evaluate these integrals. Instead, we start by using the adjoint of L_p to rewrite them as

$$(g_k, L_p f_j) \equiv (L_p^\dagger g_k, f_j). \tag{2.50}$$

The action of the adjoint operators $L_p^\dagger$ on functions $q(\xi)$ is

$$L_p^\dagger q = \sum_{j=1}^{p+2} \frac{1}{j!(p+2-j)!} \alpha_j^{(p+2-j)} \xi^{p+2-j} \partial_\xi^j q. \tag{2.51}$$

For reference, we give the first three:

$$L_0^\dagger q(\xi) = \alpha_1^{(1)} \xi \partial_\xi q + \frac{1}{2} \alpha_2^{(0)} \partial_\xi^2 q, \tag{2.52}$$

$$L_1^\dagger q(\xi) = \frac{1}{2} \alpha_1^{(2)} \xi^2 \partial_\xi q + \frac{1}{2} \alpha_2^{(1)} \xi \partial_\xi^2 q + \frac{1}{3!} \alpha_3^{(0)} \partial_\xi^3 q, \tag{2.53}$$

$$L_2^\dagger q(\xi) = \frac{1}{3!} \alpha_1^{(3)} \xi^3 \partial_\xi q + \frac{1}{4} \alpha_2^{(2)} \xi^2 \partial_\xi^2 q + \frac{1}{3!} \alpha_3^{(1)} \xi \partial_\xi^3 q + \frac{1}{4!} \alpha_4^{(0)} \partial_\xi^4 q. \tag{2.54}$$

(At fixed points satisfying equations 2.29 and 2.30, the jump moment derivatives satisfy $A_i^{(j)} = \alpha_i^{(j)}$ as for the forward operators. We have made this substitution in equations 2.51 to 2.54.) Recall that the g_k are Hermite polynomials of order k. The adjoint operators $L_p^\dagger$ contain terms of the form $\xi^n \partial_\xi^m$ with $n + m = p + 2$ and $1 \le m \le p + 2$. Using the well-known recursion relations on Hermite polynomials (Abramowitz & Stegun, 1972, for example), we can write $L_p^\dagger g_k$ as a finite sum of several different Hermite polynomials $g_l(\xi)$. One finds

$$\begin{aligned} L_p^\dagger g_k = {} & a_{k+p} g_{k+p} + a_{k+p-2} g_{k+p-2} \\ & + a_{k+p-4} g_{k+p-4} + \ldots a_{k-p-2} g_{k-p-2}, \end{aligned} \tag{2.55}$$

where the coefficients a_i involve jump moments and their derivatives and numerical factors, and the series on the right-hand side terminates at $g_0(\xi)$. The biorthogonality condition, equation 2.43, makes the final evaluation of equation 2.50 simple. The inner products, equations 2.50, are thus nonzero

on a set of the eigenfunctions f_j of number equal to the number of terms in equation 2.55. Hence, the sums over k in equations 2.47 and 2.48, and their higher-order analogs terminate after a finite number of terms.

2.3.2 Perturbation Corrections to the Moments. The lowest-order solution $P^{(0)}(\xi)$ is a gaussian with mean zero and variance given by equation 2.36. The corrections to any of the moments can be calculated order-by-order in $\eta^{1/2}$ as follows. We define a perturbation series for the *kth* moment by multiplying both sides of equation 2.31 by ξ^k and integrating to obtain

$$M_k(\eta) = M_k^{(0)} + \eta^{1/2} M_k^{(1)} + \eta M_k^{(2)} + \cdots, \tag{2.56}$$

where the nth-order correction to the kth moment is

$$M_k^{(n)} \equiv \int \xi^k P^{(n)}(\xi)\, d\xi. \tag{2.57}$$

To find the perturbation equations for the moments, return to the succession of equations 2.33 to 2.35. Multiply both sides of each by ξ^k, and use the definition of the adjoint of L_p to move their action off the $P^{(k)}(\xi)$ and onto the ξ^k. This leaves the following set of equations:

$$0 = \int P^{(0)} L_0^\dagger \xi^k d\xi, \tag{2.58}$$

$$0 = \int P^{(1)} L_0^\dagger \xi^k + P^{(0)} L_1^\dagger \xi^k d\xi, \tag{2.59}$$

$$0 = \int P^{(2)} L_0^\dagger \xi^k + P^{(1)} L_1^\dagger \xi^k + P^{(0)} L_2^\dagger \xi^k d\xi, \tag{2.60}$$

and so forth. These yield a set of algebraic equations for the corrections that do not require explicit calculation of the density corrections $P^{(i)}$. To obtain the explicit form for each of the terms in equations 2.58 to 2.60, we evaluate the action of the adjoint operators, equation 2.51 on ξ^k, multiply by $P^{(i)}$, and integrate, leaving

$$\int P^{(i)}(\xi) L_p^\dagger \xi^k d\xi = \sum_{j=1}^{\min(k,p+2)} \binom{k}{j} \frac{1}{(p+2-j)!} \alpha_j^{(p+2-j)} M_{p+2+k-2j}^{(i)}. \tag{2.61}$$

For a given k, each of equations 2.58 to 2.60, and their extensions can be solved in turn for the $M_k^{(i)}$. Equation 2.58 gives the zero-order moments (also available directly from the gaussian stationary state). Equation 2.59 gives the first-order corrections. Equation 2.60 gives the second-order corrections,

and so forth. It is straightforward to verify that to solve for $M_k^{(n)}$, one needs only the corrections

$$M_{k+i}^{(n-i)}, M_{k+i-2}^{(n-i)}, \ldots, M_{k-i-2}^{(n-i)}, \quad \text{for } i = 0, 1, 2, \ldots, n.$$

Thus, the perturbation framework yields a set of equations that provide the equilibrium moments in increasing order of $\eta^{1/2}$. The explicit perturbation corrections to the density $P^{(i)}(\xi)$ are not required to calculate the corrections to the moments $M_k^{(i)}$.

For completeness, we give the first few corrections below. The zeroth-order density is a gaussian with zero mean and variance given by equation 2.36. It is properly normalized ($M_0^{(0)} = 1$), and the higher-order moments follow from the usual relations for a gaussian. In all orders, the normalization of the perturbation density, equations 2.31, is unity (and independent of the learning rate η), reflected by the result $M_0^{(n)} = 0, n > 0$. The first few corrections are

$$\begin{aligned}
M_2^{(1)} &= M_1^{(2)} = M_3^{(2)} = 0, \\
M_1^{(1)} &= -\frac{\alpha_1^{(2)}}{2\alpha_1^{(1)}} M_2^{(0)}, \\
M_3^{(1)} &= -\frac{1}{3\alpha_1^{(1)}} \left(\frac{3}{2}\alpha_1^{(2)} M_4^{(0)} + 3\alpha_2^{(0)} M_1^{(1)} + 3\alpha_2^{(1)} M_2^{(0)} + \alpha_3^{(0)} \right), \\
M_2^{(2)} &= -\frac{1}{2\alpha_1^{(1)}} \left(\alpha_1^{(2)} M_3^{(1)} + \frac{1}{3}\alpha_1^{(3)} M_4^{(0)} + \alpha_2^{(1)} M_1^{(1)} + \frac{1}{2}\alpha_2^{(2)} M_2^{(0)} \right). \quad (2.62)
\end{aligned}$$

One assembles the terms $M_k^{(0\ldots n)}$ into the series in equation 2.56. Finally, one forms the moments of the weight w from those of the fluctuation using the original decomposition in equation 2.11, $E[w^k] = E[(\phi_* + \sqrt{\eta}\xi)^k]$.

It is important to note that perturbation expansions are typically asymptotic but may not be convergent at a particular value of η (Bender & Orszag, 1999). Hence, the accuracy of our expansion improves as $\eta \to 0$, but retaining more terms in the perturbation expansion at a given η may not improve its accuracy.

2.3.3 Computer Algebra Implementation. We implemented our perturbation expansions for the density functions and for the moments in the Mathematica computer algebra language.[5] This provides fast and accurate evaluation of the otherwise tedious, error-prone, and laborious expressions. For example, in Mathematica 8.0 running under Windows XP on an Intel

[5]Mathematica 8.0 (Wolfram Research, 2010).

Zeon 3.2 GHz processor with 3.0 GB RAM, computing the perturbation corrections to the density to fifth-, tenth-, and fifteenth-order required 8.5, 65, and 236 seconds, respectively. (These timing estimates are from a system with closed-form analytic expressions for the jump moments.)

3 Ensemble Synaptic Plasticity

We apply the methods from the section 2 to two observed biological STDP learning rules: antisymmetric Hebbian learning (Bi & Poo, 1998; van Rossum et al., 2000) and the anti-Hebbian learning observed in the electrosensory lateral line lobe (ELL) of weakly electric fish (Bell, Bodznick, Montgomery, & Bastian, 1997; Bell, Han et al., 1997; Han et al., 2000). We give theoretical results for both the probability density function and its moments, and compare them with results from Monte Carlo simulations.

3.1 Antisymmetric Hebbian Learning. Several researchers have observed antisymmetric Hebbian learning in which synapses are potentiated if the presynaptic spike precedes the postsynaptic spike and depressed if the presynaptic spike follows the postsynaptic spike (Markram et al., 1997; Bi & Poo, 1998; Zhang, Tao, Holt, Harris, & Poo, 1998). Van Rossum et al. (2000) applied a nonlinear Fokker-Planck equation to model the equilibrium synaptic weight distribution in one such learning rule. Their analysis draws on Bi and Poo's (1998) experimental results on cultured rat hippocampal cells for the particular learning rule form and modeling parameters. They use the learning rule (van Rossum et al., 2000)

$$\begin{aligned} w_{t+1} &= w_t + \eta(c_p + v w_t) e^{-\delta t/\tau}, \quad \delta t > 0, \\ w_{t+1} &= w_t + \eta(-c_d w_t + v w_t) e^{\delta t/\tau}, \quad \delta t \le 0, \end{aligned} \tag{3.1}$$

where the first expression describes potentiation and the second expression describes depression. The time of the postsynaptic spike minus the time of the synaptic event is denoted δt; c_p and c_d are constants determining the amount of potentiation and depression, respectively; τ is the time constant for the exponential learning window; and v is zero mean gaussian noise of variance σ_v^2 describing the experimentally observed fluctuations. The multiplicative scaling factor η facilitates application of the perturbation analysis. It will be set to unity at the end of the calculations.

To make analytic progress, van Rossum et al. (2000) modify the learning rule by replacing each exponential lobe in equation 3.1 with a unit height rectangular window function of width t_w as shown in Figure 2. With this simplification, the learning rule reduces to

$$\begin{aligned} w_{t+1} &= w_t + \eta(c_p + v w_t), 0 \le \delta t < t_w, \\ w_{t+1} &= w_t + \eta(-c_d w + v w_t), -t_w \le \delta t < 0. \end{aligned} \tag{3.2}$$

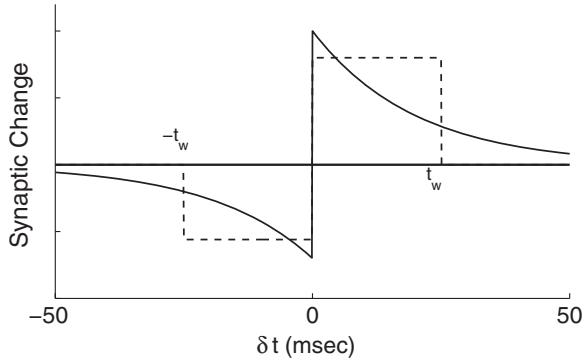

Figure 2: Time dependence of learning in the model of van Rossum (solid line) and its approximation by a rectangular learning rule (dashed line). The approximation, adopted by van Rossum, allows the Fokker-Planck equation to be solved in closed form.

Van Rossum et al. (2000) consider an integrate-and-fire neuron receiving background inputs plus incoming Poisson-distributed spikes on the single synapse w under study. This synapse undergoes depression with probability $p_d(w)$ (the probability that the incoming spike arrives in the depression window $-t_w \le \delta t < 0$). It undergoes potentiation with probability $p_p(w)$ (the probability that the incoming spike arrives in the potentiation window $0 \le \delta t < t_w$). (Note $p_d + p_p \le 1$.) Since an input spike cannot influence the generation of an earlier postsynaptic spike, the depression probability is independent of w. However, an incoming spike elevates the probability of a following postsynaptic spike, and so $p_p(w)$ increases with increasing w. The model captures this by taking

$$p_p(w) = p_d(1 + w/W_{tot}), \tag{3.3}$$

where W_{tot} is the effective synaptic weight of background inputs on the neuron. Van Rossum et al. (2000) assume that the synapse w contributes a minute fraction of the total inputs to the neuron, hence $w \ll W_{\text{tot}}$. We follow that here, taking, as they do,[6]

$$p_p = p_d \equiv p. \tag{3.4}$$

The equilibrium distribution and moments are determined by the jump moments through the Kramers-Moyal expansion, equation 2.9, or our

[6]We have derived the equilibrium moments without this assumption by a perturbation expansion in powers of $1/W_{tot}$.

fluctuation expansion of it. Under van Rossum's simplified model, the required jump moments are easy to evaluate. Combining their definition, equation 2.7, the learning rule, equation 3.2, and the condition on the depression and potentiation probabilities, equation 3.4, one finds (Adrian, 2008)

$$\begin{aligned}\alpha_n(w) &= p_p E_v[(c_p + vw)^n] + p_d E_v[(-c_d w + vw)^n] \\ &= p \sum_{k=0}^{\lfloor n/2 \rfloor} \binom{n}{2k} \sigma_v^{2k} (2k-1)!! \left(c_p^{n-2k} w^{2k} + (-c_d)^{n-2k} w^n \right), \end{aligned} \tag{3.5}$$

where $\lfloor q \rfloor$ is the largest integer equal to or less than q.

3.1.1 Equilibrium Density. As is typically done, van Rossum et al. (2000) truncate the Kramer-Moyal expansion, equation 2.9, at the second term and use the FPE for the ensemble dynamics. For their model, the drift and diffusion coefficients are (using equation 3.5)

$$\begin{aligned}\alpha_1(w) &= p(c_p - c_d w), \\ \alpha_2(w) &= p\left(c_p^2 + \left(c_d^2 + 2\sigma_v^2\right) w^2\right). \end{aligned} \tag{3.6}$$

For a one-dimensional system (as here), the equilibrium distribution can be reduced to quadrature (Risken, 1989) with the result (at $\eta = 1$) (van Rossum et al., 2000; Adrian, 2008)

$$\begin{aligned} P(w) &= \frac{K}{\alpha_2(w)} \exp\left(2 \int^w \frac{\alpha_1(w)}{\alpha_2(w)} dw \right) \\ &= K \frac{\exp(\sqrt{2} \arctan(\sqrt{2}\sigma' w / c_p)/\sigma')}{\left(2\sigma'^2 w^2 + c_p^2\right)^{\frac{2\sigma'^2 + c_d}{2\sigma'^2}}}, \end{aligned} \tag{3.7}$$

where $2\sigma'^2 \equiv 2\sigma_v^2 + c_d^2$.[7]

Figure 3A shows the FPE equilibrium density, equation 3.7, together with a histogram of weights (shaded area) from a Monte Carlo simulation of the learning rule with 2×10^4 ensemble members sampled at 10^3 time steps after convergence to equilibrium (as judged by stabilization of the moments). (The probability density has been rescaled so its area matches the number of points in the histogram.) The parameters are the same as those

[7] Van Rossum et al. (2000) further use $c_d^2 \ll \sigma_v^2$, as is appropriate for the biologically observed parameters. Their solution follows from ours with the replacement $\sigma'^2 \to \sigma_v^2$.

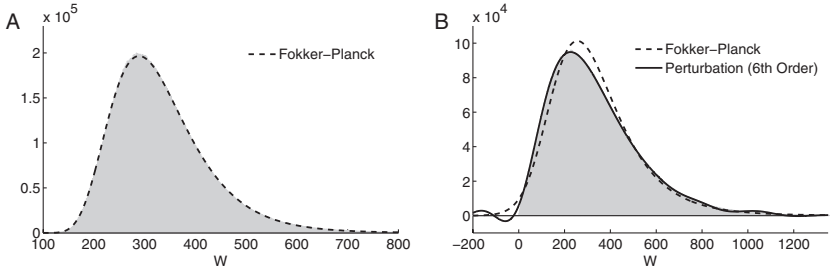

Figure 3: Equilibrium distribution of weights from Monte Carlo (shaded histogram) and as predicted by the FPE (dashed curve) using (A) parameters from van Rossum et al. (2000), $\eta = c_p = 1, c_d = 0.003, \sigma_v = 0.015$, and (B) parameters $100\times$ larger ($\eta = 1, c_p = 100, c_d = 0.3, \sigma_v = 0.015$) for which the FPE begins to break down. Also shown in panel B is the perturbation solution using terms through sixth order (solid curve).

used by van Rossum et al. (2000). Specifically, $\eta = 1, c_p = 1, c_d = 0.003$, and $\sigma_v = 0.015$. The FPE solution matches the Monte Carlo simulation extremely well. The quadratic dependence (in w) of the diffusion coefficient $\alpha_2(w)$, equation 3.6, skews the theoretical distribution toward larger weights. The good fit of the FPE equilibrium for such a strongly skewed distribution contradicts van Kampen's argument that such features arising from a nonlinear FPE are "spurious" (see section 1).

In fact, the FPE remains a remarkably good fit to the Monte Carlo densities until we reach values of c_p and c_d far larger than those observed physiologically. At high enough values of the learning rate constants, the FPE ceases to model the distribution accurately as the contributions of the higher jump moments ($\alpha_3, \alpha_4, \ldots$) become significant. We show an example in Figure 3B for which $c_p = 100$, $c_d = 0.3$, and $\sigma_v = 0.015$ (i.e., c_p and c_d are 100 times the values used by van Rossum et al., 2000, while the multiplicative noise variance is the same). The plot shows both the FPE (dashed curve) and the perturbation solution through the sixth order (solid curve) for the equilibrium density. The density plotted in Figure 3B is that given in equation 2.49 with $P(\xi)$ computed to the sixth order in perturbation.

The FPE solution fails to capture the position of the peak properly. The perturbation solution does a better job since, through operators L_1 through L_6 that occur in the solution (see equations 2.34 and 2.35), it takes into account contributions from the higher jump moments (through α_8). Hence, the perturbation expansion makes effective use of information dropped by the Fokker-Planck truncation.

The unphysical negative lobe in the density at the far left (a feature of the oscillations, which damp out by $w \approx -600$) and the small oscillation in the tail of the distribution at the far right are inevitable consequences of

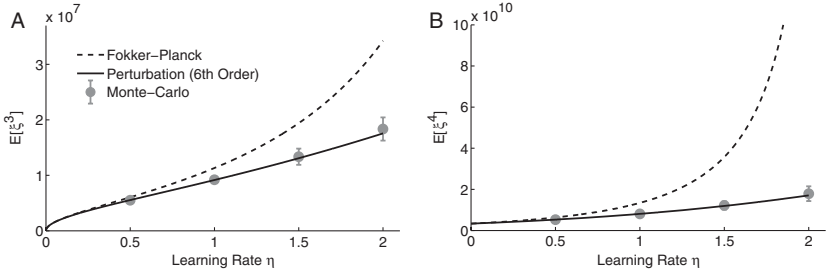

Figure 4: Equilibrium third (A) and fourth (B) moments of the fluctuations ξ for the van Rossum model with $c_p = 100, c_d = 0.3, \sigma_v = 0.015$ as a function of the learning rate η. Curves show the prediction of the FPE (dashed) and the perturbation expansion (solid) through the sixth order. The point $\eta = 1$ corresponds to the density in Figure 3B. Error bars show the $\pm 2\sigma$ confidence intervals on the moments estimated from the Monte Carlo simulation.

the expansion of the perturbation correction in orthogonal functions. All such functions oscillate and can thus lead to negative lobes in the density. The appearance of negative lobes in the perturbation density recalls the same phenomena in the Edgworth expansions of densities from classical statistics (Abramowitz & Stegun, 1972; Blinnikov & Moessner, 1998, for example). It also recalls the Pawula theorem (Risken, 1989), which states that finite truncations of the KM expansion that include terms beyond the first two are not guaranteed to yield positive densities.

Although the tail oscillations and negative lobe in the perturbation density are troubling artifacts of the expansion, a look at the third and fourth moments of the distribution confirms that the perturbation expansion is a large improvement over the FPE solution. Figure 4 shows the equilibrium third and fourth moments[8] of the fluctuations (that is, $E[\xi^3]$ and $E[\xi^4]$) as a function of learning rate η from the FPE, from the sixth-order perturbation results, and as estimated from Monte Carlo. To produce the Monte Carlo estimates, we first computed averages of ξ^k over the 2×10^4 ensemble members at each time step after burn-in. Next, this time series of moment estimates was subsampled at the first zero of its autocorrelation function. The subsamples were used to calculate the reported moment and its variance. The density in Figure 3 B corresponds to $\eta = 1$ on these curves.

The FPE overestimates both the third and fourth moments at all but the lowest learning rates, while the perturbation expansion gives results in

[8]We calculated the moments under the FPE by multiplying equation 2.10 by $w^k, k = 1, 2, 3, 4$, integrating by parts, solving the resulting algebraic equations, and transforming from w to ξ. Since the jump moments $\alpha_i(w)$ in this model are polynomials of order i or less, these moments equations can be solved in closed form. The first two moments thus calculated from the FPE are exactly correct.

excellent agreement with the Monte Carlo estimates. (We give the moments $E[\xi^k]$ since the perturbation order for these is directly comparable to that for the density, whereas calculations of the central moments $E[(\xi - E[\xi])^k]$ to a particular order use contributions of density calculations from several different orders.)

3.2 Anti-Hebbian Learning. In contrast to the antisymmetric Hebbian learning rule discussed in the previous section, for the anti-Hebbian learning rule in mormyrid electric fish, the FPE breaks down at the physiologically observed learning rates. Hence a perturbation method is required.

Mormyrid weakly electric fish have several electrosensory modalities that they use for hunting and navigating. The fish emit a short electric organ discharge (EOD) (of duration $\sim$0.25 msec). Mormyromast electroreceptors on the surface of the skin sense the EOD field, and their responses are processed centrally to detect distortions induced by nearby objects. A second system, that of interest here, detects weak low-frequency external signals (e.g., from prey) using sensitive ampullary receptor organs.

The ampullary organs respond to the fish's own discharge with spike trains whose instantaneous rate resembles a damped (roughly) sinusoidal oscillation (Bell & Szabo, 1986), persisting for roughly 100 milliseconds. This response resembles ringing initiated by the EOD. The ringing response would grossly interfere with the fish's ability to detect weak signals from other organisms. To overcome this, processing in the electrosensory lateral line lobe (ELL) of the fish's brain builds a memory of the time course of the ampullary ringing response. The memory, modifiable on timescales of a few minutes (Bell, Bodznick et al., 1997), is actually an inverted copy, or negative image, of the ampullary response to the EOD. The negative image is added to, and nulls out, the ampullary response to the EOD, allowing the brain to attend to weak, externally generated signals. Thus, the system adapts to ignore the habitual signal from the fish's own EOD, allowing the organism to attend to potentially important, novel signals.

The negative image is formed and combined with the ampullary input in medium ganglion (MG) cells in ELL. The MG cells have apical dendrites that receive inputs from parallel fibers originating outside ELL. The parallel fibers carry corollary discharge signals following the motor command that elicits the EOD. The corollary discharge signals on the parallel fibers arrive at MG cells with delays from 0 to about 100 msec following the EOD (Bell, 1982; Bell & Szabo, 1986; Bell, Bodznick et al., 1997; Roberts & Bell, 2000). (Presumably one such delayed signal is carried on each parallel fiber.)

The mormyrid ELL was one of the first systems in which STDP was observed (Bell, Han et al., 1997; Han et al., 2000). Bell and his colleagues discovered that plasticity at the parallel fiber synapses is responsible for forming and adapting the negative image. The plasticity is mediated by broad dendritic spikes in the MG cells with thresholds several times larger than the threshold for action potentials. This plasticity is sensitive to the

time between the onset of the excitatory postsynaptic potentials (EPSP) resulting from the incoming parallel fiber spikes and the broad dendritic spikes (Bell, Han et al., 1997).

3.2.1 Average Dynamics and Negative Image. Roberts gave a model of ELL that predicts the formation and stability of the negative image using average learning dynamics (Roberts & Bell, 2000). Later work relaxed restrictions on the early models (Williams, Roberts, & Leen, 2003) and began to explore the stochastic dynamics of the model (Williams, Leen, & Roberts, 2004).

In the model formulated by Roberts and adapted by Williams et al., the MG cells receive signals originating at the ampullary receptors and from the parallel fiber synapses. The resulting MG membrane potential is

$$U(x, w) = \gamma(x) + \sum_i w_i \mathcal{E}(x - x_i), \tag{3.8}$$

where x is the time since the last EOD; $\gamma(x)$ carries the effects of the ampullary inputs; w_i is the weight of the synapse from the ith parallel fiber, carrying a corollary discharge spike delayed from the EOD by time x_i; and $\mathcal{E}(x - x_i)$ is the EPSP initiated at time x_i arising from the parallel fiber input to the ith synapse. The second term on the right of equation 3.8 carries the effect of the cascade of corollary discharge spikes, and through the synaptic weights, w_i builds up the negative image. The EPSP time course is modeled as

$$\mathcal{E}(x) = \begin{cases} (x/\tau^2)\exp(-x/\tau), & x \geq 0 \\ 0, & x < 0 \end{cases},$$

where τ is the time constant and the factor $1/\tau^2$ normalizes the EPSP to unit area.

In the absence of a broad dendritic spike, EPSPs evoke a synaptic potentiation α. If a broad dendritic spike occurs at time x, the ith synapse undergoes a change,

$$\Delta w_i = \eta(\alpha - \beta L(x - x_i)), \tag{3.9}$$

where α and β are positive constants, $L(x - x_i) = L(\delta t_i)$ is the learning function, and we have inserted the factor η to facilitate the perturbation expansion. The learning function is normalized to have unit area. The two components, the nonassociative potentiation and the associative depression, are depicted in Figure 5.

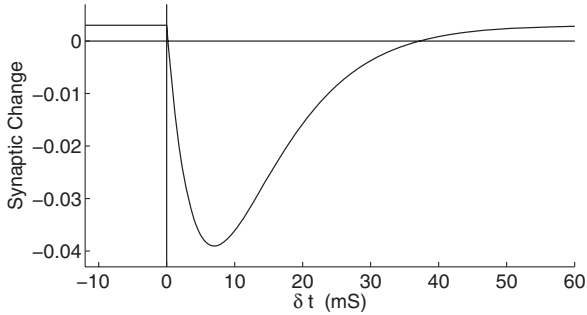

Figure 5: Modeled electric fish learning rule from equation 3.9. The EPSP begins at 0.0 ms, coinciding with the beginning of the learning window for associative depression. The parameters are $\alpha = 0.003$, $\beta = 0.0008$s, $\tau = 7$ ms and fall in the physiologically observed range.

Following Roberts (2000) and Williams et al. (2003, 2004), we model the broad spike productions as a Poisson process with rate at time x after the EOD given by

$$f(x) = \frac{f_{\max}}{1 + \exp\left[-\mu(U(x) - \Theta)\right]}, \tag{3.10}$$

where Θ is the broad spike threshold and μ quantifies the effects of noise on spike generation. The learning rule, equation 3.9, implicitly assumes at most a single broad spike (at time x) per EOD cycle. Within this restriction, equation 3.10 is sensible only if $f(x)$ is small enough to ensure

$$\text{Prob(one spike in EOD cycle)} = \int_0^T f(x)\,dx \leq 1, \tag{3.11}$$

where T is the EOD cycle length. As discussed below in section 3.2.2, the physiologically observed parameters allow this constraint to be satisfied for the single-synapse model treated here.

The average weight change in an EOD cycle is, from the expectation of equation 3.9,

$$\eta\alpha_{1i}(w) = E[\Delta w_i] = \eta\left(\alpha - \beta\int_0^T f(x)L(x - x_i)dx\right). \tag{3.12}$$

Roberts (2000) and Williams et al. (2003) show that with $L(x) \approx \mathcal{E}(x)$ (as for the real system), there is a stable fixed point at which this average weight

change is zero ($\alpha_{1i}(w) = 0$), and at this fixed point, the values of w are such that the MG membrane potential is approximately constant:

$$U(x) = \gamma(x) + \sum_i w_i \mathcal{E}(x - x_i) \approx U_0 = \text{constant}. \tag{3.13}$$

That is, there is a stable fixed point for which an approximate negative image of the ampullary signal $\gamma(x)$ is formed by the weighted sum of EPSPs initiated by the parallel fiber corollary discharges.

3.2.2 Beyond the Average Motion. Williams et al. (2004) move beyond the average dynamics to explore the statistics of the equilibrium weight distribution. They proceed by linearizing the firing rate, equation 3.10, about the fixed point. This allows an exact solution for the moments of the equilibrium distribution valid when the distribution is narrow enough not to feel the curvature in the logistic function, equation 3.10. We refer to this treatment as the linearized theory. They address both a simplified model with a single synapse and a full model with multiple synapses.

Here we apply our perturbation techniques from section 2.3 to the single-synapse model and compare our results with those from the linearized theory, the FPE, and Monte Carlo simulations. The single parallel fiber EPSP occurs at $x_1 = 0$, and we take the membrane potential to be

$$U(x; w) = U_0 + (w - w_*)\mathcal{E}(x). \tag{3.14}$$

At $w = w_*$, $U(x; w_*) = U_0$. This will be a fixed point $\alpha_1(w_*) = 0$ if (from equation 3.12 with $x_1 = 0$)

$$f_0 \equiv f(U_0) = \alpha/\beta. \tag{3.15}$$

We set U_0 and Θ to satisfy this for all that follows. The jump moments follow from equation 3.9 and the broad spike rate in equation 3.10. We find

$$\alpha_n(w) = \alpha^n + \sum_{j=1}^{n} \alpha^{n-j}(-\beta)^j \binom{n}{j} \int_0^T f(x) L^j(x). \tag{3.16}$$

Our perturbation methods require the derivatives of the jump moments at the fixed point $\partial_w^j \alpha_n(w_*)$, and we evaluate the required integrals,

$$\int \frac{\partial^j f(U(x, w_*))}{\partial w^j} L^i(x) dx,$$

numerically.

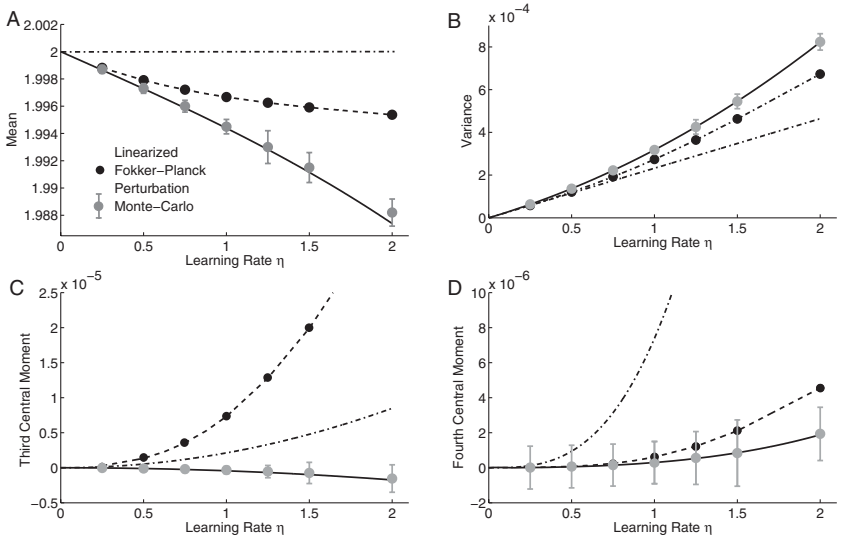

Figure 6: Plots of the (A) mean, (B) second, (C) third, and (D) fourth central moments of w at equilibrium for ELL synaptic weights. Shown are the moments calculated with the linearized theory (dot-dashed line), with the Fokker-Planck equilibrium density (dashed line with circle markers), with the perturbation analysis (solid line), and from Monte Carlo simulations of the learning rule. Error bars show the $\pm 2\sigma$ confidence intervals on the moments from the simulations. Learning rate $\eta = 1$ corresponds to the physiological values. For the perturbation theory curves, the mean was calculated to second order, the variance to sixth order, the third moment to fourth order, and the fourth moment to eighth order.

Figure 6 shows the first four (central) moments of w for the equilibrium distribution as a function of the learning rate η. The plots show the moments from the linearized theory, the Fokker-Planck equation, and our perturbation expansion. (We obtained the moments from the FPE equilibrium density by numerical integration.) Also shown are the moments estimated from Monte Carlo simulations and their $\pm 2\sigma$ confidence intervals. For the perturbation results, the first moment was calculated to second order, the second moment to sixth order, the third moment to fourth order, and the fourth moment to eighth order. Thus, for example, all of the terms through order $\eta^{8/2}$ were retained in the calculation of $E[(w - E[w])^4]$.

The parameters were taken from Roberts (2000) and fall in the physiologically observed range: $\alpha = 0.003$, $\beta = 0.0008$s, $\tau = 7$ ms, $\mu = 2$, and $f_{\max} = 15$/s (P. D. Roberts, personal communication, 2011). We integrate over the effective support of the $\mathcal{E}$ and L functions, $T = 60$ ms. Thus, $f_{\max} < 1/T$ and the constraint in equation 3.11 is guaranteed to hold. The

fixed point in the weights is free, and we arbitrarily set it to $w_* = 2.0$. Our Monte Carlo has 4800 ensemble members. The simulation was burned in for 1000 EOD cycles, and an additional 5000 EOD cycles were used to estimate the moments and the variance (subsampled at the first zero of the autocorrelation function, as for the moment estimates in Figure 4).

Under both the FPE and the perturbation treatment, the mean in Figure 6A is shifted downward with respect to the linearized theory (in which the mean is fixed at w_*). The fixed point $f_0 = \alpha/\beta$ is near the bottom of the sigmoid, equation 3.10, where the curvature is positive, so the sigmoid is locally above its linearization. Small fluctuations in the weight away from w_* in either direction lead to a higher spike rate than for the linearized theory, and thus more associative depression. Hence, the mean is shifted downward relative to the fixed point.

As shown by the Monte Carlo simulation (gray dots with error bars), the perturbation treatment captures this shift in mean more accurately than the FPE. The prediction of the perturbation theory departs slightly from the simulation at the highest learning rate. Perturbation expansions such as these are typically asymptotic rather than convergent, and higher orders in perturbation did not give improved results. The perturbation treatment captures the variance (see Figure 6B) more faithfully than either the FPE or the linearized theory. Saliently, both the FPE and linearized theory predict the wrong sign for the central third moment, while the perturbation theory agrees well with the simulation (see Figure 6C). The predictions for the fourth moment (see Figure 6D) by the FPE and perturbation analyses are quite close to each other and to the Monte Carlo results. Larger Monte Carlo simulations are required to show statistically significant disagreement with the FPE results for the fourth moment at all but the highest learning rate shown.

Finally, though not reported here, we have relaxed the restriction to a single broad spike per EOD cycle, developing a perturbation treatment for the case when $f(x)$ is large enough to allow multiple broad spikes per EOD. Following Roberts and Lafferriere (2005), the modified learning rule assumes that effects on plasticity of multiple broad spikes within the same learning window (support of $L(x)$) are additive.[9] The corresponding perturbation expansions for the equilibrium moments show agreement with Monte Carlo simulations qualitatively similar to that depicted in Figure 6. However, the distribution for the multispike case has a lower mean and larger variance than for the single-spike case. Those differences are expected from the differences in the jump moments between the two cases.

[9]There are no physiology data for this system on synaptic changes when several dendritic spikes occur within the depressive lobe of the learning rule in Figure 5.

4 Discussion

The perturbation methods we have given here provide systematic and well-grounded tools for approximating the moments and probability density for stochastic learning and neural dynamics. Our development gives a clear and simple order-by-order expansion that goes beyond the linear noise regime commonly retained from the system size expansion. Unlike van Kampen's seminal work, our approach gives the higher-order corrections for the density as well as for the moments. Application to several STDP learning rules shows that perturbation methods can provide accurate estimates of moments and densities in regimes where the Fokker-Planck equation breaks down. Our implementation by computer algebra provides rapid and accurate evaluation.

Our perturbation expansion of the density (see section 2.3.1) expresses the corrections to the zeroth order (linear noise regime) gaussian as a linear combination of basis functions (the $f_k(\xi)$), each of which is a Hermite polynomial times a gaussian. This construction recalls the Edgeworth expansion for probability densities (Abramowitz & Stegun, 1972, or Blinnikov & Moessner, 1998, for a review). In contrast with our perturbation development, the Edgeworth expansion provides successive approximations for the density when the cumulants are exactly-known, so we are not able to give a term-by-term correspondence between the two developments.

The pronouncements against the FPE made by van Kampen are clearly extreme. As Figure 3A and numerous examples in the machine learning and STDP literature show, it is clearly not the case that predictions of the FPE beyond the lowest-order gaussian approximation are spurious. One of our motivations for this article is to provide an even-handed critique of the FPE, as well as to give alternative tools for calculation when required. As we stated in section 1, the fundamental problem is not that the nonlinear FPE is never applicable, but rather that there is no apparent way to use it as the basis of a complete perturbation expansion. This is unfortunate; the FPE solution would be an attractive zeroth-order approximation as it already captures skew.

We compared solutions of the FPE with results from the perturbation expansion and with histograms and moments estimated from Monte Carlo experiments. In broad terms, one expects the FPE to break down when the contributions from the discarded jump moments $\alpha_3, \alpha_4, \ldots$ become appreciable. To track this precisely, one could take our perturbation expansion and set the jump moments $\alpha_3, \alpha_4, \ldots$ identically equal to zero and compare the full perturbation to the truncated perturbation results. This would give an order-by-order comparison between the solution of the FPE and the full solution and could more precisely illuminate the limits of applicability of the FPE.

In general, our perturbation expansions are asymptotic (becoming more accurate as $\eta \to 0$) but may not be convergent at arbitrary η. Thus, we expect them to be most useful when the per step changes in the learning

process are small. In our development, we introduced the scaling parameter η essentially as a power-counting device. The magnitude of step-wise changes is governed by the phenomenological constants in the observed learning rules: c_p, c_d, and σ_v in van Rossum's model and α, β in the electric fish rule. In the range presented, the perturbation calculations are useful. However, for larger values, they can certainly break down. Since our expansions are generally not convergent, going to higher order can degrade the approximation. In particular, one can find high-frequency oscillations in the tails of the distribution. (This also occurs in the classical Edgeworth expansions as explored by Blinnikov & Moessner, 1998.) However when integrated against smooth functions, such as low-order polynomials, these oscillations tend to cancel out. Consequently the perturbation expansion tends to give good estimates of the moments even if the density approximation carries large, spurious oscillations. However, even for the moments, the expansions are not typically convergent, and increasing the order can degrade the approximation. Given that perturbation expansions may not be convergent, one would like guidelines for truncating them, and heuristics for truncating asymptotic expansions do exist. Taking, for example, the series for the moments, equation 2.56, Boyd (1999) suggests retaining everything up to the smallest term $\eta^{i/2}M_k^{(i)}$, while Bender and Orszag (1999) suggest truncating at the term just before the smallest term. Neither of these heuristics produced good results for us.

Although we cannot make general statements about the radius of convergence for the perturbation expansion in general (indeed, it may be zero in some cases), one case deserves attention. When the jth jump moment $\alpha_j(w)$ is a polynomial of order j or less in w, the moments can be derived in closed form from the KM expansion (Gardiner, 2009). (This is the case, for example in van Rossum's model, Adrian, 2008.) In that case, the perturbation expansion must agree (term by term) with the Taylor series expansion (in $\eta^{1/2}$) of the exact moments, since the expansion asymptotic to the true moment function (in a given asymptotic scale) is unique (Bender & Orszag, 1999). We have confirmed the agreement through the third-order terms for the first four moments of the equilibrium distribution.

The perturbation treatment generalizes to multiple dimensions in a straightforward way, although the expressions become more cumbersome. The jump moments become tensor objects: $\alpha_1(w)$ is a vector, $\alpha_2(w)$ is a symmetric matrix, $\alpha_3(w)$ is a symmetric third-rank tensor, and so forth. As long as the jump moments and their derivatives can be computed from the learning rule, there is no fundamental impediment to developing the perturbation expansion.[10] As for the one-dimensional case treated here, the

[10] Burkitt et al. (2004) discusses challenges that can arise in computing the higher jump moments in the case where multiple excitatory postsynaptic potentials and postsynaptic spikes occur within the same learning window.

perturbation corrections to the equilibrium moments follow from simple algebraic expressions without explicit use of the perturbation corrections to the density. The Fokker-Planck equation could also be applied to such problems. Although solving for a multidimensional density for the FPE is generally not tractable, one can write a perturbation expansion similar to that developed here, but retaining only the first two jump moments α_1, α_2, as described three paragraphs back. Computing the perturbation corrections to the density requires the eigenfunctions of the multidimensional analog of L_0 and its adjoint. These are apparently not given in the literature, but we have derived them using a ladder operator formalism. Applications to multidimensional systems will be given in a later publication.

We leave for future work application to systems in which bounds on the weights determine the equilibrium (see the discussion of equilibria in section 2.3). Application to STDP rules with triplet and quadruplet interactions or rate dependence would also be useful. Mathematica code to carry out the density and moment expansions is available at http://www.bme.ogi.edu/~tleen/MEPerturb/.

Appendix: From the FPE to the Fluctuation Expansion

This appendix motivates the fluctuation expansion from the desire to use the Fokker-Planck equilibrium as the basis for a full perturbation expansion. In one dimension, the equilibrium solution of the Fokker-Planck equation reduces to quadrature (Risken, 1989):

$$P(w) = \frac{K}{\alpha_2(w)} \exp\left(\frac{2}{\eta}\int^w \frac{\alpha_1(w')}{\alpha_2(w')}dw'\right). \tag{A.1}$$

The appearance of η in this solution causes problems for a perturbation expansion. When equation A.1 is substituted into the K-M expansion, equation 2.9, the remainder pieces are all of order η^0 (unity) rather than of higher orders in η as required to form a perturbation series (Heskes & Kappen, 1993). The lowest-order piece of the perturbation series solution ought to be independent of η and should hold in the limit $\eta \to 0$. The coordinate transformation adopted in the fluctuation expansion, equation 2.11, accomplishes precisely that. We will arrive at the transformation by considering the limiting case of the FPE equilibrium.

As η is decreased, the step size decreases and (in analogy with the rudimentary random walk) the variance decreases. In the small η limit, the distribution will be peaked up sharply about some value w_*. Assume an isolated fixed point $\alpha(w_*) = 0$ and, for simplicity, $\alpha_1^{(1)}(w) < 0$ and $\alpha_2(w) > 0$ everywhere. Then the integral in the exponent of equation A.1 peaks at w_* and drops off to the left and right. As η decreases, the density drops off ever more quickly to the left and right of w_*, and the portion of the density

substantially above zero becomes well modeled by

$$P(w) \approx \frac{K_1}{\alpha_2(w_*)} \exp\left(\frac{2}{\eta}\int^{w} \frac{\alpha_1^{(1)}(w_*)}{\alpha_2(w_*)}(w' - w_*)dw'\right)$$
$$= K_2 \exp\left(-\frac{1}{\eta}\frac{(w - w_*)^2}{\left[\alpha_2(w_*)/|\alpha_1^{(1)}(w_*)|\right]}\right). \tag{A.2}$$

This equilibrium suggests that to eliminate the explicit dependence on η, we should make the change of coordinates

$$w = w_* + \sqrt{\eta}\xi \tag{A.3}$$

(compare with equation 2.11). Then the lowest-order approximation to the equilibrium is a gaussian peaked up at w_*, and the fluctuations ξ have variance

$$\sigma_\xi^2 = \frac{\alpha_2(w_*)}{2|\alpha_1(w_*)|}. \tag{A.4}$$

This gives a density for the fluctuations ξ that is independent of η, and one would expect that the coordinate transformation in equation A.3, together with this lowest-order equilibrium density establishes the basis for a proper perturbation expansion (in powers of $\eta^{1/2}$) of the density. The result is precisely the development given for the equilibrium density in section 2.2.

Acknowledgments

This work was supported by the NSF under grant IIS-0812687 and in part by IIS-0827722. We thank Joel Franklin, Michael Friedlander, and Ignacio Sáez for helpful conversation and the anonymous reviewers for helpful comments. Patrick D. Roberts was especially helpful in discussing the mormyrid electric fish model and our results. Frank Adrian prepared the initial Monte Carlo simulations for the van Rossum model. David Nielsen prepared the final Monte Carlo and perturbation results for the electric fish learning rule.

References

Abramowitz, M., & Stegun, L. (1972). *Handbook of mathematical functions*. Washington, DC: U.S. Department of Commerce, National Bureau of Standards.

Adrian, F. A. (2008). *Exact ensemble dynamics for spike-timing-dependent plasticity*. Master's thesis, Oregon Health & Science University.

Babadi, B., & Abbott, L. (2010). Intrinsic stability of temporally shifted spike-timing dependent plasticity. *PLoS Colmput. Biol.*, *6*(11):e1000961. doi:10,1371/journal.pcbi.1000961.

Bedeaux, D., Lakatos-Lindenberg, K., & Shuler, K. E. (1971). On the relation between master equations and random walks and their solutions. *Journal of Mathematical Physics*, *12*, 2116–2123.

Bell, C. (1982). Properties of a modifiable efference copy in an electric fish. *J. Neurophysiol.*, *47*, 1043–1056.

Bell, C., Bodznick, D., Montgomery, J., & Bastian, J. (1997). The generation and subtraction of sensory expectations within cerebellum-like structures. *Brain Behav. Evol.*, *50*, 17–31 (suppl. 1).

Bell, C., Han, V., Sugawara, Y., & Grant, K. (1997). Synaptic plasticity in a cerebellum-like structure depends on temporal order. *Nature*, *387*, 278–281.

Bell, C., & Szabo, T. (1986). Electroreception in mormyrid fish: Central anatomy. In T. Bullock & W. Heiligenberg (Eds.), *Electroreception* (pp. 375–421). New York: Wiley.

Bender, C. M., & Orszag, S. A. (1999). *Advanced mathematical methods for scientists and engineers: Asymptotic methods and perturbation theory*. New York: Springer.

Bi, Q., & Poo, M. (1998). Precise spike timing determines the direction and extent of synaptic modifications in cultured hippocampal neurons. *Journal of Neuroscience*, *18*, 10464–10472.

Bienenstock, E. L., Cooper, L. N., & Munro, P. W. (1982). Theory for the development of neuron selectivity: Orientation specificity and binocular interaction in visual cortex. *Journal of Neuroscience*, *2*, 32–48.

Bishop, C. M. (2006). *Pattern recognition & machine learning*. New York: Springer.

Blinnikov, S., & Moessner, R. (1998). Expansions for nearly gaussian distribution. *Astron. Astrophys. Suppl. Ser.*, *130*, 193–205.

Boyd, J. P. (1999). The devil's invention: Asymptotic, superasymptotic and hyperasymptotic series. *Acta Applicandae Mathematicae*, *56*, 1–99.

Buice, M. A., Cowan, J. D., & Chow, C. C. (2010). Systematic fluctuation expansion for neural network activity equations. *Neural Computation*, *22*, 377–426.

Burkitt, A. N., Meffin, H., & Grayden, D. B. (2004). Spike-timing-dependent plasticity: The relationship to rate-based learning for models with weight dynamics determined by a stable fixed point. *Neural Computation*, *16*, 855–940.

Câteau, H., & Fukai, T. (2003). A stochastic method to predict the consequence of arbitrary forms of spike-timing-dependent plasticity. *Neural Computation*, *15*(3), 597–620.

Cowan, J. (1991). Stochastic neurodynamics. In R. Lipmann, J. Moody, & D. Touretzky (Eds.), *Advances in neural information processing systems, 3*. San Mateo, CA: Morgan Kaufmann.

Darken, C., & Moody, J. (1991). Note on learning rate schedules for stochastic optimization. In R. Lipmann, J. Moody, & D. Touretzky (Eds.), *Advances in neural information processing systems, 3*. San Mateo, CA: Morgan Kaufmann.

Darken, C., & Moody, J. (1992). Towards faster stochastic gradient search. In J. Moody, S. Hanson, & R. Lipmann (Eds.), *Advances in neural information processing systems, 4*. San Mateo, CA: Morgan Kaufmann.

Der, R., & Villmannn, T. (1993). *Dynamics of self organized feature mapping* (Tech. Rep.). Leipzig: Universitat Leipzig, Inst. für Informatik.

Fabian, V. (1968). On asymptotic normality in stochastic approximation. *Ann. Math. Statist.*, *39*, 1327–1332.

Fabian, V. (1973). Asymptotically efficient stochastic approximation: The RM case. *Annals of Statistics*, *1*, 486–495.

Gardiner, C. (2009). *Stochastic methods: A handbook for the natural and social sciences* (4th ed.). Berlin: Springer-Verlag.

Gelfand, I., & Shilov, G. (1964). *Generalized functions*. Orlando, FL: Academic Press.

Gray, C., & Singer, W. (1989). Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. *Proceedings of the National Academy of Sciences of the U.S.A.*, *86*(5), 1698–1702.

Guckenheimer, J., & Holmes, P. (1983). *Nonlinear oscillations, dynamical systems, and bifurcations of vector fields*. New York: Springer-Verlag.

Han, V., Grant, K., & Bell, C. (2000). Reversible associative depression and nonassociative potentiation at a parallel fiber synapse. *Neuron*, *27*, 611–622.

Hansen, L. K. (1993). Stochastic linear learning: Exact test and training error averages. *Neural Networks*, *6*, 393–396.

Haykin, S. (1991). *Adaptive filter theory* (2nd ed.). Englewood Cliffs, NJ: Prentice Hall.

Heskes, T. M., & Kappen, B. (1993). On-line learning processes in artificial neural networks. In J. Taylor (Ed.), *Mathematical foundations of neural networks* (pp. 199–233). Dordrecht: Elsevier.

Ismailov, I., Kalikulov, D., Inoue, T., & Friedlander, M. (2004). The kinetic profile of intracellular calcium predicts long-term potentiation and long-term depression. *J. Neuroscience*, *24*, 9847–9861.

Kepecs, A., van Rossum, M. C., Song, S., & Tegney, J. (2002). Spike-timing-dependent plasticity: Common themes and divergent vistas. *Biol. Cybern.*, *87*, 446–448.

Kuan, C.-M., & Hornik, K. (1991). Convergence of learning algorithms with constant learning rates. *IEEE Transactions on Neural Networks*, *2*, 484–489.

Kushner, H., & Clark, D. (1978). *Stochastic approximation methods for constrained and unconstrained systems*. New York: Springer-Verlag.

Leen, T. K. (1998). Exact and perturbation solutions to the ensemble dynamics. In D. Saad (Ed.), *Online learning in neural networks*. Cambridge: Cambridge University Press.

Leen, T. K., & Orr, G. B. (1994). Optimal stochastic search and adaptive momentum. In J. Cowan, G. Tesauro, & J. Alspector (Eds.), *Advances in neural information processing systems, 6*. San Francisco: Morgan Kaufmann.

Leen, T. K., Schottky, B., & Saad, D. (1998). Two approaches to optimal annealing. In M. Jordan, M. Kearnes, & S. Solla (Eds.), *Advances in neural information processing systems, 10*. Cambridge, MA: MIT Press.

Leen, T., Schottky, B., & Saad, D. (1999). Optimal asymptotic learning rate: Macroscopic versus microscopic dynamics. *Physical Review, E*, *59*, 985–991.

Ljung, L. (1977). Analysis of recursive stochastic algorithms. *IEEE Trans. Automatic Control*, *22*, 551–575.

Markram, H., Lubke, J., Fortscher, M., & Sakmann, B. (1997). Regulation of synaptic efficacy by conicidence of postsynaptic AP's and EPSP's. *Science*, *275*, 213–215.

Masuda, N., & Aihara, K. (2004). Self-organizing dual coding based on spike-time-dependent plasticity. *Neural Computation*, *16*, 627–663.

Morrison, A., Diesmann, M., & Gerstner, W. (2008). Phenomenological models of synaptic plasticity based on spike timing. *Biological Cybernetics*, *98*, 459–478.

Ohira, T., & Cowan, J. D. (1993). Master-equation approach to stochastic neurodynamics. *Physical Review E*, *48*, 2259–2266.

Orr, G. B. (1996). *Dynamics and algorithms for stochastic search*. Unpublished doctoral dissertation, Oregon Graduate Institute.

Orr, G. B., & Leen, T. K. (1993). Weight space probability densities in stochastic learning: II. Transients and basin hopping times. In C. L. Giles, S. J. Hanson, & J. D. Cowan (Eds.), *Advances in neural information processing systems, 5*. San Mateo, CA: Morgan Kaufmann.

Orr, G. B., & Leen, T. K. (1997). Using curvature information for fast stochastic search. In M. Mozer, M. Jordan, & T. Petsche (Eds.), *Advances in neural information processing systems, 9*. Cambridge MA: MIT Press.

Pfister, J.-P., & Gerstner, W. (2006a). Beyond pair-based STDP: A phenomenological rule for spike triplet and frequency effects. In Y. Weiss, B. Schölkopf, & J. Platt (Eds.), *Advances in neural information processing systems, 18*. Cambridge, MA: MIT Press.

Pfister, J.-P., & Gerstner, W. (2006b). Triplets of spikes in a model of spike timing-dependent plasticity. *J. Neuroscience*, *26*, 9673–9682.

Radons, G. (1993). On stochastic dcynamics of supervised learning. *J. Phys. A*, *26*, 3445–3461.

Radons, G., Schuster, H., & Werner, D. (1990). Fokker-Planck description of learning in backpropagation networks. In *International Neural Network Conference—INNC 90* (pp. II 993–996). Norwell, MA: Kluwer.

Risken, H. (1989). *The Fokker-Planck equation*. Berlin: Springer-Verlag.

Robbins, H., & Monro, S. (1951). A stochastic approximation method. *Ann. Math. Stat.*, *22*, 400–407.

Roberts, P. D. (2000). Modeling inhibitory plasticity in the electrosensory system of mormyid electric fish. *Journal of Neurophysiology*, *84*, 2035–2047.

Roberts, P. D., & Bell, C. C. (2000). Computational consequences of temporally asymmetric learning rules: II. Sensory image cancellation. *Journal of Computational Neuroscience*, *9*, 67–83.

Roberts, P., & Lafferriere, G. (2005). *Notes on modeling plasticity in the ELL*. Unpublished research notes.

Rubin, J., Lee, D., & Sompolinsky, H. (2001). Equilibrium properties of temporally asymmetric Hebbian plasticity. *Phys. Rev. Lett.*, *86*, 364–367.

Sejnowski, T. (1977). Statistical constraints on synaptic plasticity. *J. Theor. Biol.*, *69*, 385–389.

van Kampen, N. (1976). The expansion of the master equation. In I. Prigogine & S. A. Rice (Eds.), *Advances in chemical physics*, *34*. New York: Wiley.

van Kampen, N. (2007). *Stochastic processes in physics and chemistry* (3rd ed.). Amsterdam: Elsevier North-Holland.

van Rossum, M., Bi, G., & Turrigiano, G. (2000). Stable Hebbian learning from spike timing-dependent plasticity. *Journal of Neuroscience*, *20*, 8812–8821.

Venter, J. H. (1967). An extension of the Robbins-Monro procedure. *Annals of Mathematical Statistics*, *38*, 117–127.

White, H. (1989). Learning in artificial neural networks: A statistical perspective. *Neural Computation*, *1*, 425–464.

Williams, A., Leen, T. K., & Roberts, P. D. (2004). Random walks for spike-timing-dependent plasticity. *Physical Review E*, *70*.

Williams, A., Roberts, P. D., & Leen, T. (2003). Stability of negative-image equilibrium in spike-timing-dependent plasticity. *Physical Review E*, *68*, 021923.

Zhang, L. I., Tao, H. W., Holt, C. E., Harris, W. A., & Poo, M. (1998). A critical window for cooperation and competition among developing retinotectal synapses. *Nature*, *395*, 37–42.

Received April 22, 2011; accepted October 26, 2011.