ARTICLE IN PRESS

Neural Networks ▮ (▮▮▮▮) ▮▮▮–▮▮▮

Contents lists available at SciVerse ScienceDirect

Neural Networks

journal homepage: www.elsevier.com/locate/neunet

2012 Special Issue

Approximating distributions in stochastic learning

Todd K. Leen [a,*], Robert Friel [b], David Nielsen [a]

[a] Department of Biomedical Engineering, Oregon Health & Science University, Portland, OR, United States

[b] Courant Institute of Mathematical Sciences, New York, NY, United States

ARTICLE INFO

Keywords:
STDP
Online learning
Fokker–Planck equation
Master equation
State-space distributions
Perturbation theory

ABSTRACT

On-line machine learning algorithms, many biological spike-timing-dependent plasticity (STDP) learning rules, and stochastic neural dynamics evolve by Markov processes. A complete description of such systems gives the probability densities for the variables. The evolution and equilibrium state of these densities are given by a Chapman–Kolmogorov equation in discrete time, or a master equation in continuous time. These formulations are analytically intractable for most cases of interest, and to make progress a nonlinear Fokker–Planck equation (FPE) is often used in their place. The FPE is limited, and some argue that its application to describe jump processes (such as in these problems) is fundamentally flawed.

We develop a well-grounded perturbation expansion that provides approximations for both the density and its moments. The approach is based on the system size expansion in statistical physics (which does *not* give approximations for the density), but our simple development makes the methods accessible and invites application to diverse problems. We apply the method to calculate the equilibrium distributions for two biologically-observed STDP learning rules and for a simple nonlinear machine-learning problem. In all three examples, we show that our perturbation series provides good agreement with Monte-Carlo simulations in regimes where the FPE breaks down.

1. Introduction

Both on-line machine learning algorithms, and adaptation in many neural networks mediated by spike-timing-dependent plasticity (STDP) are stochastic processes described by a Markov process. While average dynamics may describe much phenomena, the fluctuations inherent in noisy learning create scatter around the average. This variability affects the accuracy with which a circuit can carry out its task, as is well-appreciated in on-line learning applied to adaptive filter theory (Haykin, 1991). Similarly, the dynamics and function of neural *activity* can be described in mean terms by firing rates. However this treatment neglects the role of correlated firing (Buice, Cowan, & Chow, 2010). A complete account of the effects of a noisy learning or dynamical system specifies the time evolution of the *probability density* for the dynamical variables.

The master or Chapman–Kolmogorov equations (Gardiner, 2009) for the evolution of the probability density are usually intractable. The most common approach is to use, in their place, a Fokker–Planck equation (FPE) (Gardiner, 2009; Risken, 1989) to describe the dynamics. The FPE has been successfully applied to both basin-hopping in machine learning (Der & Villmannn, 1993; Hansen, 1993; Orr & Leen, 1993; Radons, 1993) and to equilibrium distributions in biological learning (Babadi & Abbott, 2010; Câteau & Fukai, 2003; Kepecs, van Rossum, Song, & Tegney, 2002; Masuda & Aihara, 2004; Rubin, Lee, & Sompolinsky, 2001; van Rossum, Bi, & Turrigiano, 2000; Zhu, Lai, Hoppensteadt, & He, 2006). Despite its utility, adopting the FPE to approximate jump processes is not well-justified. Van Kampen argues strongly against adopting a nonlinear FPE for jump processes (unless rigorously derived by a proper limiting procedure), saying of its solutions that "any features they contain beyond that [predicted by linear drift and constant diffusion] are spurious" (van Kampen, 2007, p. 262). Linear drift and constant diffusion produce Gaussian densities, so van Kampen's claim is that non-Gaussian densities predicted from an FPE should be viewed with great suspicion. Nonetheless, an FPE with nonlinear drift or non-constant diffusion adopted ad hoc often gives good results, as shown in both the machine-learning and theoretical biology literature cited.

In our view, the issue is not whether the FPE is a useful approximation, but rather that there is no immediately-apparent way to use it as the lowest-order approximation to a complete perturbation expansion. Such a construction would enable one to make corrections to its predictions, and hence more precisely discuss its domain of applicability.

van Kampen (1961) and van Kampen (2007) developed his system-size expansion as a rigorous alternative to the FPE to describe jump processes. However it is widely-neglected and

* Corresponding author. Tel.: +1 503 332 5897.
E-mail address: leent@ohsu.edu (T.K. Leen).

0893-6080/$ – see front matter © 2012 Elsevier Ltd. All rights reserved.
doi:10.1016/j.neunet.2012.02.006

poorly understood, in part because his development of the higher order terms is complicated. Furthermore the method does not give higher-order corrections to the *probability density*, which would provide useful visualization.

We give here a perturbation approach appropriate for machine and STDP learning, and for neurodynamics. Our development is clear and straightforward, and supplies the *density* as well as its moments. Breaking with the bulk of the physics literature, we use a discrete-time approach, which is more appropriate than continuous-time dynamics for comparison with transients in learning experiments and simulations, and eliminates any possible confusion arising from continuous-time analysis applied to manifestly discrete-time processes. It is not our intent to discourage use of the FPE, but rather to provide tools for computation in regimes for which the FPE is inadequate.

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, Sugawara, and Grant (1997), and first modeled by Roberts and Bell (2000). We also apply the expansions to an example machine learning problem, logistic regression fit by on-line maximum likelihood estimation.

For the van Rossum Hebbian rule, we compare equilibrium densities predicted by the perturbation expansion with those from the FPE, and with histograms from Monte-Carlo simulations. Contrary to van Kampen's pronouncements, the solution of the FPE is an excellent fit to the highly-skewed densities revealed by the Monte-Carlo. To break the FPE solution, we must significantly raise the learning rate. When we do so (using 100× the physiological learning rate), the FPE does not agree with the simulations, but our perturbation expansion shows good agreement.

In contrast to van Rossum's model, for the electric fish anti-Hebbian rule the FPE is inadequate at the physiologically-observed learning rates. We compare moments of the equilibrium densities predicted by our perturbation expansion with those from the FPE and from a linearized theory. The results show large differences between the theoretical methods. The FPE fails to predict the moments, producing the *wrong sign* for the skew. The perturbation expansion is in excellent agreement with simulations.

For the on-line logistic regression problem, we show the equilibrium distributions and the moments from both the FPE and the perturbation expansion, comparing them with Monte-Carlo simulations. The FPE over-predicts the moments examined, and again, predicts the *wrong sign* for the skew.

An abbreviated version of this article was published in the Proceedings of the IJCNN (Leen & Friel, 2011). The formal development is here extended in several places, we provide new examples, and discuss the approach and results in more depth.

In the following Section 2.1 we develop the ensemble dynamics and equations of motion for the probability density. We develop our perturbation solution to those equations in Section 2.2. In Section 3 we apply the methods as discussed above.

2. Noisy learning

2.1. Ensemble dynamics

We consider learning rules of the form

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

where w_t is the synaptic weight at time-step t, η is the learning rate (constant for this discussion), $h(\cdot)$ is the learning rule, and ψ is a random variable giving rise to the randomness of the dynamics.

The learning rule (1) generates a Markov process on w; the weight w_{t+1} depends only on the weight w_t and the probability of the transition $w_t \to w_{t+1}$. The probability density for the weights thus evolves according to the Chapman–Kolmogorov equation (Gardiner, 2009)

$$P(w, t+1) = \int P(w', t) W(w' \to w)\, dw' \tag{2}$$

where $W(w' \to w)$ is the probability of making the transition from w' to w in one time step. The transition probability for *fixed* ψ is a delta function satisfying the learning rule (1), and we obtain the unconditional transition probability by taking its expectation with respect to ψ

$$\begin{aligned} W(w' \to w) &= E_\psi[W(w' \to w|\psi)] \\ &= E_\psi[\delta(w - w' - \eta h(w', \psi))]. \end{aligned} \tag{3}$$

The Kramers–Moyal (KM) expansion of the Chapman–Kolmogorov equation (2) serves as the starting point for our perturbation treatment, and for comparison with the commonly-adopted Fokker–Planck equation (FPE). The expansion, derived in standard texts (Gardiner, 2009, for example), proceeds by substituting the transition probability (3) into the master equation (2) and expanding the result in powers of η. The result is

$$P(w, t+1) - P(w, t) = \sum_{j=1}^{\infty} \frac{(-\eta)^j}{j!} \partial_w^j (\alpha_j(w) P(w, t)) \tag{4}$$

where ∂_w^j denotes the jth partial derivative with respect to w, and the *jump moments* are

$$\alpha_j(w) = E_\psi[h^j(w, \psi)]. \tag{5}$$

The first jump moment $\alpha_1(w)$, called the *drift coefficient*, gives the average change in one time step. The second jump moment $\alpha_2(w)$, called the *diffusion coefficient*, describes the propensity of the system to spread out.

The KM expansion (4) is usually intractable. The common procedure is to retain only the first two terms, leaving the FPE

$$\begin{aligned} P(w, t+1) - P(w, t) &= -\eta \partial_w (\alpha_1(w) P(w, t)) \\ &\quad + \frac{\eta^2}{2} \partial_w^2 (\alpha_2(w), P(w, t)). \end{aligned} \tag{6}$$

The truncation of the KM expansion to the FPE is convenient and widely-used. However it is only strictly justified when arrived at by a limiting procedure in which the higher order jump moments $\alpha_3, \alpha_4, \ldots$ necessarily vanish. In machine and biological learning systems (and in the treatment of fluctuations in neural activity), the truncation is generally *not* justified.

Instead of adopting the FPE equation, we develop a perturbation expansion for the dynamics of $P(w, t)$. Our treatment follows the roots of van Kampen's system size expansion, and so gives the same results for the moments. However we give a much clearer development than previously available, and our approach provides approximations for the *density* as well as the moments. Finally, we retain a discrete-time approach, rather than passing to continuous time as does most of the literature (Gardiner, 2009, for example). This eliminates possible ambiguity arising from adopting a continuous-time approach, and may be more convenient for studying transients in learning experiments. The equilibrium densities are, of course, identical for the discrete and continuous-time treatments.[1]

[1] Bedeaux, Lakatos-Lindenberg, and Shuler (1971) show that after many time-steps, the solutions to the continuous time master equation and the associated discrete-time random walk agree.

Please cite this article in press as: Leen, T. K., et al. Approximating distributions in stochastic learning. Neural Networks (2012), doi:10.1016/j.neunet.2012.02.006

Following van Kampen, we decompose the weight trajectory w_t into the sum of a deterministic and a fluctuating piece

$$w_t = \phi_t + \sqrt{\eta}\xi_t \tag{7}$$

where ϕ_t is the deterministic piece and ξ_t are the fluctuations. In analogy with a simple random walk (and in the spirit of the central limit theorem) we have assumed that the *variance* of w scales *linearly* with the step size η, hence the $\sqrt{\eta}$ multiplying ξ. This scaling assumption is ultimately verified by confirming that the lowest order contribution to the variance of ξ is indeed independent of η.

Next we extract the dynamics of the deterministic piece, by substituting the decomposition of Eq. (7) into the learning rule equation (1)

$$\phi_{t+1} + \sqrt{\eta}\xi_{t+1} = \phi_t + \sqrt{\eta}\xi_t + \eta h(\phi_t + \sqrt{\eta}\xi_t, \psi_t). \tag{8}$$

Next, expand the last term on the right-hand side of Eq. (8) in powers of $\sqrt{\eta}$ and take the expectation of both sides with respect to ψ. Extracting the pieces of the result that are *independent of* ξ_t yields the equation for the deterministic trajectory

$$\phi_{t+1} = \phi_t + \eta\alpha_1(\phi_t). \tag{9}$$

To obtain the equation of motion for the fluctuations, substitute Eq. (9) into Eq. (8) yielding

$$\begin{aligned}\xi_{t+1} &= \xi_t + \sqrt{\eta}[h(\phi_t + \sqrt{\eta}\xi_t, \psi) - \alpha_1(\phi_t)] \\ &\equiv \xi_t + \sqrt{\eta}H(\phi_t + \sqrt{\eta}\phi_t, \phi_t, \psi_t),\end{aligned} \tag{10}$$

where we have defined the *fluctuation learning function*

$$H(w, \phi, \psi) \equiv h(w, \psi) - \alpha_1(\phi). \tag{11}$$

We can obtain a KM expansion for the evolution of the probability density on ξ analogous to Eq. (4). The result is

$$\begin{aligned}&P(\xi, t+1) - P(\xi, t) \\ &= \sum_{j=1}^{\infty} \frac{(-\eta)^{j/2}}{j!}\partial_\xi^j (A_j(\phi_t + \sqrt{\eta}\xi)P(\xi, t))\end{aligned} \tag{12}$$

where the *fluctuation jump moments* are

$$A_j(w, \phi) = E_\psi[H^j(w, \phi, \psi)]. \tag{13}$$

(Compare Eqs. (5) and (13).) We complete the derivation by expanding the jump moment on the right-hand side of Eq. (12) in powers of $\sqrt{\eta}$ and re-arranging terms to find

$$\begin{aligned}&P(\xi, t+1) - P(\xi, t) \\ &= \sum_{k=2}^{\infty}\sum_{j=1}^{k} \frac{(-1)^j}{j!(k-j)!}\eta^{k/2}A_j^{(k-j)}\partial_\xi^j(\xi^{k-j}P(\xi, t))\end{aligned} \tag{14}$$

where we have used $A_1(\phi) \equiv 0$, and defined

$$A_j^{(m)} \equiv \partial_w^m A_j(w, \phi)|_{w=\phi_t}, \tag{15}$$

the coefficients of the expansion of the jth jump moment about ϕ_t.

Eqs. (9) and (14) together provide a complete description of the dynamics. The right-hand side of Eq. (14) is identical to that in the continuous-time system size expansion (Gardiner, 2009; van Kampen, 2007) apart from the appearance of the $A_j^{(m)}$ in place of the $\alpha_j^{(m)}$.

2.2. Perturbation solution for the density and moments

Eq. (14) for the fluctuation density is of the form

$$P(\xi, t+1) - P(\xi, t) = \eta(L_0 + \eta^{1/2}L_1 + \eta L_2 + \cdots)P(\xi, t) \tag{16}$$

where

$$L_i f(\xi) = \sum_{j=1}^{i+2} \frac{(-1)^j}{j!(i+2-j)!}A_j^{(i+2-j)}\partial_\xi^j(\xi^{i+2-j}f(\xi)). \tag{17}$$

For reference, the first two operators are

$$L_0 f(\xi) = -A_1^{(1)}\partial_\xi(\xi f) + \frac{1}{2}A_2^{(0)}\partial_\xi^2 f \tag{18}$$

$$L_1 f(\xi) = -\frac{1}{2}A_1^{(2)}\partial_\xi(\xi^2 f) + \frac{1}{2}A_2^{(1)}\partial_\xi^2(\xi f) - \frac{1}{3!}A_3^{(0)}\partial_\xi^3 f. \tag{19}$$

For the remainder of the paper, we confine attention to *equilibrium* solutions for which $\phi_{t+1} = \phi_t$, and $P(\xi, t+1) = P(\xi, t)$. We assume that there are no constraints bounding the values of the weights and restrict attention to equilibrium solutions that are peaked up around asymptotically-stable fixed points ϕ_* of the deterministic trajectory (9)

$$\alpha_1(\phi_*) = 0 \tag{20}$$

$$-2 < \eta\alpha_1^{(1)}(\phi_*) < 0, \tag{21}$$

where the first condition establishes the fixed point, and the second its asymptotic stability[2] (Guckenheimer & Holmes, 1983). In what follows, we assume that $\phi = \phi_*$. Note that from Eqs. (11) and (20), at $\phi = \phi_*$, $H = h$. It follows from Eqs. (5) and (13) that $A_j(w) = \alpha_j(w)$ at $w = \phi_* + \sqrt{\eta}\xi$. This insures that the equilibrium of the discrete-time and continuous-time systems are identical, as one would expect. Finally, at equilibrium, the left side of Eq. (16) is zero. Henceforth we drop the time-dependence.

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

$$P(\xi) = P^{(0)}(\xi) + \eta^{1/2}P^{(1)}(\xi) + \eta P^{(2)}(\xi) + \cdots \tag{22}$$

where the $P^{(k)}(\xi)$ are unknown functions (independent of η) that we will solve for. This decomposition, along with the operator expansion in Eq. (16), are the *two critical elements* that provide a clear development for the approximations to the density and its moments.

To solve for the $P^{(k)}$ we substitute Eq. (22) into Eq. (16) (with the latter's left-hand side set to zero at equilibrium). Since the resulting equation must hold regardless of the value of η, it must be true that the terms of equal order in η separately vanish. Hence we obtain the succession of equations

$$L_0 P^{(0)}(\xi) = 0 \tag{23}$$

$$L_0 P^{(1)}(\xi) = -L_1 P^{(0)} \tag{24}$$

$$L_0 P^{(2)}(\xi) = -L_1 P^{(1)}(\xi) - L_2 P^{(0)}(\xi)$$

$$\vdots \tag{25}$$

and so on.

The zeroth-order density $P^{(0)}(\xi)$ satisfies Eq. (23), which carries the *linear* part of the drift and the *constant* part of the diffusion coefficients — see Eq. (18). This describes the equilibrium of an Ornstein–Uhlenbeck process (Gardiner, 2009). The solution is a zero-mean Gaussian

$$P^{(0)}(\xi) = \frac{1}{\sqrt{2\pi\sigma_0^2}}\exp\left(-\frac{\xi^2}{2\sigma_0^2}\right) \tag{26}$$

[2] For the reader familiar with stability of fixed points of differential equations, the lower bound in Eq. (21) may be surprising, but it is the correct result for discrete-time maps (Guckenheimer & Holmes, 1983).

Please cite this article in press as: Leen, T. K., et al. Approximating distributions in stochastic learning. Neural Networks (2012), doi:10.1016/j.neunet.2012.02.006

with variance

$$\sigma_0^2 = \frac{\alpha_2(\phi_*)}{2|\alpha_1^{(1)}(\phi_*)|}. \tag{27}$$

This is the lowest order (so-called *linear noise approximation*) piece of van Kampen's system size expansion. By Eq. (22), in the limit $\eta \to 0$ the density is dominated by $P^{(0)}(\xi)$. Thus the equilibrium solution (peaked up 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 (27) is *independent* of η, the scaling of fluctuations in the original decomposition is justified, as per the paragraph immediately following Eq. (7).

2.2.1. 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[3] $L_0^\dagger$

$$L_0 f_k(\xi) = \lambda_k f_k(\xi) \tag{28}$$

$$L_0^\dagger g_k(\xi) = \lambda_k g_k(\xi), \quad k = 0, 1, 2, \ldots. \tag{29}$$

The $g_k(\xi)$ are Hermite polynomials, and the $f_k(\xi)$ are Hermite polynomials times a Gaussian (harmonic oscillator wave functions). The first of these is the zeroth-order density $P^{(0)}(\xi) = f_0(\xi)$. The eigenvalues are $\lambda_k = -k|\alpha_1^{(1)}|$. The two sets of functions form a bi-orthogonal set $(g_j, f_k) = \delta_{jk}$ under the standard $\mathcal{L}_2$ inner product (Gardiner, 2009, Section 5.2).

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

$$P^{(1)}(\xi) = \sum_{j=1}^{\infty} a_j f_j(\xi). \tag{30}$$

To find the coefficients a_j, substitute this expansion into Eq. (24), use Eq. (28), take the inner product of both sides with g_k and use the bi-orthogonality. The result is

$$P^{(1)}(\xi) = -\sum_{k=1}^{\infty} \frac{1}{\lambda_k}(g_k, L_1 P^{(0)}) f_k(\xi). \tag{31}$$

With this correction in hand, the next is found following the same procedure on Eq. (25). The result is

$$P^{(2)}(\xi) = -\sum_{k=1}^{\infty} \frac{1}{\lambda_k}(g_k, L_1 P^{(1)} + L_2 P^{(0)}) f_k(\xi). \tag{32}$$

Higher order (in $\eta^{1/2}$) corrections are similarly found.

This procedure involves solving at each order i an inhomogeneous equation of the form

$$L_0 P^{(i)}(\xi) = q^{(i)}(\xi) \tag{33}$$

where the $q^{(i)}$ are known once $P^{(i-1)}, P^{(i-2)}, \ldots, P^{(0)}$ are obtained. Eq. (33) is solved by

$$P^{(i)}(\xi) = \sum_{k=1}^{\infty} \frac{1}{\lambda_k} f_k(\xi)(g_k, q^{(i)}) \equiv \int \left(\sum_{k=1}^{\infty} \frac{1}{\lambda_k} f_k(\xi) g_k(\xi') \right) q^{(i)}(\xi') d\xi' \tag{34}$$

which identifies

$$\gamma(\xi, \xi') = \sum_{k=1}^{\infty} \frac{1}{\lambda_k} f_k(\xi) g_k(\xi') \tag{35}$$

as the Green function (kernel for the formal inverse) for L_0.

Calculating the corrections to the density requires evaluating inner products of the form $(g_k, L_p f_j) = (L_p^\dagger g_k, f_j)$. Using the well-known recursion relations for the Hermite polynomials (Abramowitz & Stegun, 1972, for example) we can write $L_p^\dagger g_k$ as a finite sum of several different Hermite polynomials $g_l(\xi)$. The bi-orthogonality condition $(g_n, f_l) = \delta_{n,l}$ allows one to complete the calculations easily. The inner products $(L_p^\dagger g_k, f_j)$ are nonzero over only a finite set of the eigenfunctions f_j; hence the sums over k in Eqs. (31), (32), and the higher-order analogs (more generally, Eq. (34)) terminate after a finite number of terms.

2.2.2. Perturbation corrections to the moments

The lowest order solution $P^{(0)}(\xi)$ is a Gaussian with mean zero and variance given by Eq. (27). 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 Eq. (22) by ξ^k and integrating to obtain

$$M_k(\eta) = M_k^{(0)} + \eta^{1/2} M_k^{(1)} + \eta M_k^{(2)} + \cdots \tag{36}$$

where the nth order correction to the kth moment is

$$M_k^{(n)} \equiv \int \xi^k P^{(n)}(\xi)\, d\xi. \tag{37}$$

Now return to the succession of perturbation Eqs. (23)–(25). Multiply both sides of each by ξ^k and use the definition of the adjoint of L_n to move their action off the Ps and onto the ξ^k. This leaves the set of equations

$$0 = \int P^{(0)} L_0^\dagger \xi^k \, d\xi \tag{38}$$

$$0 = \int P^{(1)} L_0^\dagger \xi^k + P^{(0)} L_1^\dagger \xi^k \, d\xi \tag{39}$$

$$0 = \int P^{(2)} L_0^\dagger \xi^k + P^{(1)} L_1^\dagger \xi^k + P^{(0)} L_2^\dagger \xi^k \, d\xi \tag{40}$$

$$\vdots$$

and so forth. These yield a set of algebraic equations for the corrections to the moments as follows. From the adjoint of Eq. (17) and the definition in Eq. (37) we write the terms in (Eqs. (38)–(40),…) directly

$$\int P^{(n)}(\xi) L_p^\dagger \xi^k \, d\xi = \sum_{j=1}^{\min(k,p+2)} \binom{k}{j} \frac{1}{(p+2-j)!} \alpha_j^{(p+2-j)} M_{p+2+k-2j}^{(n)}. \tag{41}$$

One can verify that to obtain $M_k^{(n)}$ one needs only the moments $M_{k+i}^{(n-i)}, M_{k+i-2}^{(n-i)}, \ldots, M_{i-i-2}^{(n-i)}$, for $i = 0, 1, 2, \ldots, n$. Thus the perturbation framework yields a set of algebraic equations that provide the equilibrium moments in increasing order of $\eta^{1/2}$ without having to directly evaluate the corrections to the density.

In *all* orders, the normalization of the perturbation density (22) is unity, reflected by the result $M_0^{(n)} = 0$, $n > 0$. The first few corrections are

[3] Given an inner product between functions in a Hilbert space (a, b) and an operator $\mathcal{O}$, the adjoint $\mathcal{O}^\dagger$ is defined by $(a, \mathcal{O}b) = (\mathcal{O}^\dagger a, b)$, for all a, b for which the inner products exist.

$$M_2^{(1)} = M_1^{(2)} = M_3^{(2)} = 0$$
$$M_1^{(1)} = -\frac{\alpha_1^{(2)}}{2\alpha_1^{(1)}} M_2^{(0)},$$
$$M_3^{(1)} = -\frac{1}{3\alpha_1^{(1)}}\left(\frac{3}{2}\alpha_1^{(2)} M_4^{(0)} + 3\alpha_2^{(0)} M_1^{(1)} + 3\alpha_2^{(1)} M_2^{(0)} + \alpha_3^{(0)}\right)$$
$$M_2^{(2)} = -\frac{1}{2\alpha_1^{(1)}}\left(\alpha_1^{(2)} M_3^{(1)} + \frac{1}{3}\alpha_1^{(3)} M_4^{(0)} + \alpha_2^{(1)} M_1^{(1)} + \frac{1}{2}\alpha_2^{(2)} M_2^{(0)}\right). \quad (42)$$

One assembles the terms $M_k^{(0\ldots n)}$ into the series in Eq. (36) and lastly forms the moments of the weight w from the original decomposition in Eq. (7) $E[w^k] = E[(\phi_* + \eta^{1/2}\xi)^k]$. Note that calculating $E[\xi^k]$ through terms $\eta^{n/2}$ leads to an expression for $E[w^k]$ that contains terms up to $\eta^{(n+k)/2}$.

2.3. Limitations and implementation

It is important to note that perturbation expansions are typically asymptotic but not convergent (Bender & Orszag, 1999). Hence, the accuracy of our expansions improves as $\eta \to 0$, but retaining more terms in the perturbation expansion (22) may not improve its accuracy at a fixed, non-zero value of η.

We implemented the perturbation expansions for the density functions and for the moments in the Mathematica[4] computer algebra language. This provides fast and accurate evaluation of the other-wise tedious and error-prone expressions. In *Mathematic 8.0* running under Windows XP on an Intel Xeon 3.2 GHz processor with 3.0 GB RAM, computing corrections to the density at fifth and tenth orders required 8.5 s and 236 s respectively.[5]

3. Biological and machine plasticity

We apply the methods from the last section to two observed biological STDP learning rules and a machine learning problem. 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 post-synaptic spike, and depressed if the presynaptic spike follows the post-synaptic spike (Bi & Poo, 1998; Markram, Lubke, Fortscher, & Sakmann, 1997; Zhang, Tao, Holt, Harris, & ming Poo, 1998). Van Rossum et al. applied a nonlinear Fokker–Planck equation to model the equilibrium synaptic weight distribution in one such learning rule (van Rossum et al., 2000). Their analysis draws on Bi and Poo's experimental results (1998) for the particular learning rule form and modeling parameters.

The simplified learning rule analyzed by van Rossum et al. (2000) is

$$w_{t+1} = w_t + \eta(c_p + v w_t), \quad 0 \le \delta t < t_w$$
$$w_{t+1} = w_t + \eta(-c_d w + v w_t), \quad -t_w \le \delta t < 0. \quad (43)$$

The first expression describes potentiation and the second depression. The constants c_p and c_d determine the amount of potentiation and depression respectively, δt is the time of the post-synaptic spike minus the time of the synaptic event, and t_w is the window width for potentiation and depression. The model includes multiplicative noise v that is zero mean Gaussian with variance σ_v^2. (The scaling factor η facilitates power counting for the perturbation analysis. It is set to unity to recover the physiological learning rate.)

Van Rosum et al. consider an integrate-and-fire neuron receiving Poisson-distributed spikes on the single synapse w under study, plus background input. This synapse undergoes depression with probability p_d, the probability that the incoming spike arrives in the depression window $-t_w \le \delta t < 0$ and potentiation with probability $p_p(w)$. An input spike cannot influence the generation of an *earlier* postsynaptic spike, so p_d is independent of w. However an incoming spike elevates the probability of a following post-synaptic spike, so $p_p(w)$ is an increasing function of w.

The jump moments are simple to evaluate, combining the definition equation (5) with the learning rule equation (43) yields (Adrian, 2008)

$$\alpha_n(w) = p_p E_v[(c_p + vw)^n] + p_d E_v[(-c_d w + vw)^n]$$
$$= \sum_{k=0}^{\lfloor n/2 \rfloor} \binom{n}{2k} \sigma_v^{2k} (2k-1)!!$$
$$\times (p_p(w) c_p^{n-2k} w^{2k} + p_d (-c_d)^{n-2k} w^n) \quad (44)$$

where $\lfloor q \rfloor$ is the largest integer equal to or less than q. When the total synaptic input to the cell (from the background inputs) is much larger than the input through the synapse under study $w \ll W_{Tot}$, the input through w has little effect on p_p and one has $p_p = p_d \equiv p$. This is the limit treated by van Rossum et al. and we follow that here.

3.1.1. Equilibrium density

As is typically done, van Rossum et al. truncate the Kramer–Moyal expansion (4) at the second term and use the resulting FPE for the ensemble dynamics. For their model, the drift and diffusion coefficients are (using Eq. (44))

$$\alpha_1(w) = p(c_p - c_d w)$$
$$\alpha_2(w) = p(c_p^2 + (c_d^2 + 2\sigma_v^2) w^2). \quad (45)$$

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

$$P(w) = \frac{K}{\alpha_2(w)} \exp\left(2\int^w \frac{\alpha_1(w)}{\alpha_2(w)} dw\right)$$
$$= K \frac{\exp(2\arctan(\sqrt{2}\sigma' w/c_p)/\sigma')}{(2\sigma'^2 w^2 + c_p^2)^{\frac{2\sigma'^2 + c_d}{2\sigma'^2}}} \quad (46)$$

where[6] $2\sigma'^2 \equiv 2\sigma_v^2 + c_d^2$, and we have set $\eta = 1$.

Fig. 1(A) shows the FPE density (46) together with a histogram of weights from a Monte-Carlo simulation for the parameter values $\eta = 1$, $c_p = 1$, $c_d = 0.003$, and $\sigma_v = 0.015$ used by van Rossum et al. The FPE solution is indistinguishable from the histogram derived from the Monte-Carlo. The quadratic (in w)

[4] *Mathematica 8.0*, Wolfram Research, Inc. Champaign, Illinois, 2010.

[5] These estimates are for a system whose jump moments are evaluated analytically; when numerical evaluation of the jump moments is required, the calculation is slower.

[6] Van Rossum et al. 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 $\sigma'^2 \to \sigma_v^2$.

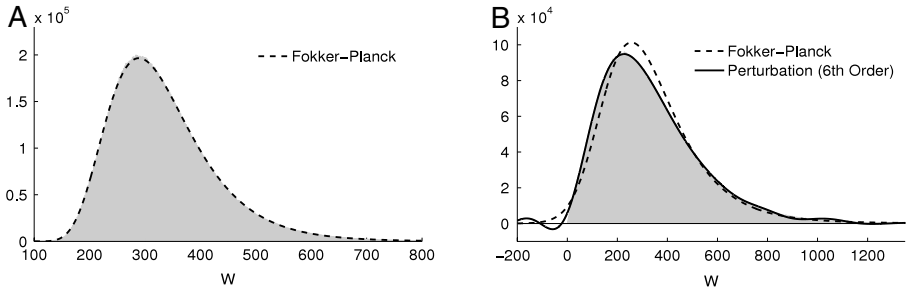

Fig. 1. Equilibrium distribution of weights from Monte-Carlo (shaded histogram) and as predicted by the FPE using (A) parameters from van Rossum et al., and (B) parameters for which the FPE fails. Also shown is the perturbation solution using terms through sixth order.

dependence of the diffusion coefficient α_2 (45) increases spreading at larger w, skewing the distribution to the right. The good fit of the FPE equilibrium for such a strongly non-Gaussian distribution contradicts van Kampen's arguments that such features arising from a non-linear FPE are "spurious" (van Kampen, 2007, p. 262).

At much higher values of c_p and c_d, the FPE ceases to model the distribution accurately since the contributions of the higher jump moments $(\alpha_3, \alpha_4, \ldots)$ become significant. Fig. 1(B) shows an example with $\eta = 1$, $c_p = 100$, $c_d = 0.3$ (that is $100\times$ the physiologically-observed values), and $\sigma_v = 0.015$. The plot shows the FPE (dashed curve) and the perturbation solution through sixth order (solid curve). The perturbation solution does a better job than the FPE since it takes into account contributions from higher jump moments as required at higher step sizes. For both plots the Monte-Carlo had 2×10^4 ensemble members sampled for 1000 time steps following burn-in.

Fig. 1(B) also shows a defect of the perturbation expansion. The oscillations in the tails of the distribution (and the consequent unphysical negative lobe in the density just left of $w = 0$) are unavoidable consequences of an expansion in orthogonal functions. All such functions oscillate, so any expansion in an orthogonal complete set will exhibit oscillations. This recalls the oscillations in density estimations from the classical Edgeworth expansion (Blinnikov & Moessner, 1998), an asymptotic expansion developed to improve on the central limit theorem. Despite these oscillations, the perturbation theory provides very accurate prediction of the *moments* of the equilibrium density for this problem (not given here, but see Leen and Friel (2012)).

3.2. Anti-Hebbian learning in mormyrid fish

This section treats the distribution of weights following the anti-Hebbian STDP learning rule exhibited in mormyrid weakly electric fish. In contrast to the anti-symmetric Hebbian learning just discussed, in this case the density derived from the FPE is inadequate even at the physiological learning rate.

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 ∼0.25 ms) that is used for active location and navigation. A second system detects weak, low-frequency signals from animals (e.g. 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 oscillations (Bell & Szabo, 1986), persisting up to a hundred milliseconds. This 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 an inverted copy, or *negative image*, of the ampullary response to the EOD that is combined with and nulls out the ampullary response to the EOD. Thus the systems adapts to ignore the habitual signal arising from the fish's own EOD, allowing the organism to attend to potentially-important, novel signals. The negative image is modifiable on a time scale of minutes, so the fish can adjust to slow changes in the environment (for example changes in water conductivity) that would affect the ampullary response to the EOD (Bell, 1982; Bell, Bodznick, Montgomery, & Bastian, 1997; Bell & Szabo, 1986).

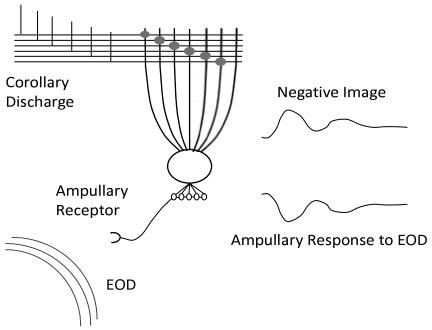

Fig. 2. Schematic showing ampullary receptor and corollary discharge inputs onto an MG cell, along with the ampullary response to the EOD, and its negative image.

Fig. 2 shows a schematic of the system. The negative image is formed and combined with the ampullary input in *medium ganglion* (MG) cells in ELL. Apical dendrites on MG cells receive inputs from parallel fibers that 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 ms following the EOD (Bell, 1982; Bell, Bodznick et al., 1997; Bell & Szabo, 1986; Roberts & Bell, 2000).

The mormyrid ELL was one of the first systems in which STDP was observed (Bell, Han et al., 1997). 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. This plasticity is sensitive to the time between the onset of the excitatory post-synaptic potentials (EPSP) resulting from the incoming parallel fiber spikes, and the broad dendritic spikes (Bell, Han et al., 1997; Roberts & Bell, 2000, and references therein).

3.2.1. Simplified model

Roberts gave a model of ELL that predicts the formation and stability of the negative image using average learning dynamics based on the observed STDP learning rule (Roberts, 2000; Roberts & Bell, 2000). Later work by Williams et al. relaxed restrictions on the early models (2003), and began to explore the stochastic dynamics of the model (Williams, Leen, & Roberts, 2004).

We examine a simplified model that considers a *single* parallel fiber synapse, but uses the STDP learning rule and broad spike firing probability used in the full model (Roberts, 2000). We take the MG membrane potential to be

$$U(x, w) = U_0 + \gamma(x) + w\mathcal{E}(x) = U_0 + (w - w*)\mathcal{E}(x) \qquad (47)$$

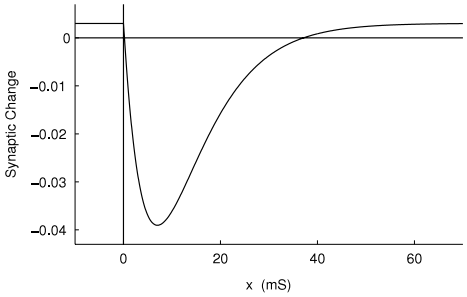

Fig. 3. Modeled time-dependence of electric fish anti-Hebbian learning rule (after Roberts (2000)) with $\tau = 7$ ms, $\alpha = 0.003$, $\beta = 0.8$ ms.

where x is the time since the last EOD, w is the weight of the synapse from the parallel fiber, $\mathcal{E}(x)$ is the EPSP initiated by the parallel fiber input The function $\gamma(x)$ carries the effects of the ampullary inputs (here vastly simplified to provide a convenient fixed point for the system). The EPSP time course is modeled as

$$\mathcal{E}(x) = \begin{cases} \dfrac{y}{\tau^2}\exp(-y/\tau) & y \geq 0 \\ 0 & y < 0 \end{cases} \tag{48}$$

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

The learning rule is as follows. In the absence of a broad dendritic spike, the EPSP evokes a synaptic potentiation α. If a broad dendritic spike occurs at time x after the EPSP, the synapse undergoes a change

$$\Delta w = \eta(\alpha - \beta L(x)) \tag{49}$$

where α and β are positive constants, $-L(x)$ is the learning function, which takes the same form as the EPSP (closely followed in the real system), and we have inserted the factor η to facilitate the perturbation expansion. Fig. 3 shows the two components of the learning rule, the non-associative potentiation, and the associative depression using the physiologically-motivated parameters that we adopt for our simulations (Roberts, 2000).

The probability density for a broad dendritic spike occurring at time x after the EOD (equivalently the instantaneous broad spike rate) is modeled as a logistic sigmoid

$$f(x) = \frac{f_{\max}}{1 + \exp -\mu(U(x) - \Theta)} \tag{50}$$

where Θ is the broad spike threshold, $U(x)$ is the membrane potential, and $1/\mu$ quantifies noise level in spike generation. We view Eq. (50) as the instantaneous spike rate for a Poisson process in the limit that $f(x)$ is so small that no more than one spike is likely in an EOD cycle. The corresponding probability that *no* dendritic spike occurs in an EOD cycle is

$$P(\text{no spike}) = 1 - \int_0^T f(x)dx. \tag{51}$$

The average weight change in an EOD cycle follows from the expectation of Eq. (49)

$$\langle \Delta w \rangle = \eta\alpha_1(w) = \eta\left(\alpha - \beta\int_0^T f(x)L(x)dx\right) \tag{52}$$

where T is the EOD cycle duration (up to several hundred milliseconds in the full system). Analogous to the full system (Roberts, 2000; Roberts & Bell, 2000), our single-synapse model has a stable fixed point with

$$w = w_*, \qquad U(x) = U_0, \quad \text{and} \quad \langle \Delta w \rangle = 0$$

provided

$$f(U_0) \equiv f_o = \alpha/\beta$$

which establishes a condition on $U_0 - \Theta$.

3.2.2. Beyond the average motion

Williams et al. (2004) moved beyond the average dynamics to explore the statistics of the equilibrium weight distribution. They proceed by *linearizing* the firing rate $f(x)$ (50) about the fixed point $U(x) = U_0$. This allows an *exact solution* for the moments of the equilibrium distribution (but not for the probability density).

We apply our perturbation techniques from Section 2.2 to the single-synapse model, and compare our results with those from the linearized theory, from the FPE, and from Monte-Carlo simulations. The jump moments follow from Eq. (49) and the broad spike probability Eqs. (50) and (51). We find

$$\alpha_n(w) = \alpha^n + \sum_{j=1}^{n} \alpha^{n-j}(-\beta)^j \binom{n}{j} \int_0^T f(x)L^j(x). \tag{53}$$

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 $\int \partial_w^j f(U(x, w_*))L^i(x)dx$ numerically.

Fig. 4 shows the first three moments of the equilibrium distribution as a function of the learning rate η. Shown are the moments predicted by the linearized theory, by the FPE (from numerical integration), by our perturbation expansion, and from Monte-Carlo simulations (with $\pm 2\sigma$ confidence intervals). The point $\eta = 1$ corresponds to the physiologically-observed learning rates. The perturbation results are clearly in better agreement with the Monte Carlo simulations than the linearized theory or the FPE for all but the smallest learning rates. Notice that both the FPE and the linearized theory give the *wrong sign* for the third central moment. The first moment was calculated to fourth order in perturbation, the second moment to sixth order, and the third moment to fourth order.

The shift in mean from the fixed point $w = w_* = 2$ is caused by the nonlinearity in the drift coefficient $\alpha_1(w)$, which couples the higher-order moments into the mean. The sigmoidal firing probability curve lies *above* its linearization at the fixed point w_*, so fluctuations away from the fixed point lead to *higher* broad spike firing rate than predicted by the linearization, and hence more associative depression. Thus the mean is shifted *downward* relative to the linearized theory.

The parameters are from Roberts (2000) and are in the range observed physiologically: $\alpha = 0.003$, $\beta = 0.0008$ s,[7] $\tau = 7$ ms, $f_{\max} = 16.667\ \text{s}^{-1}$, and $\mu = 2$. We use a time window $T = 60$ ms which covers the support of the EPSP, and so is adequate for this single-synapse model. The Monte Carlo has 4800 weights in the ensemble. The simulation was burned in for 1000 EOD cycles and an additional 5000 EOD cycles (sub-sampled at the first zero of the autocorrelation function) were used to calculate the moment estimates and their error bars.

Finally, in another investigation not shown here, we have relaxed the low-firing rate condition, allowing multiple broad spikes per EOD cycle. The agreement between the perturbation theory and the simulations is similar to that shown here, but with further depression of the mean (relative to the fixed point), and larger equilibrium variance.

3.3. Online logistic regression

As a final example, we examine the equilibrium distribution for a simple algorithm that fits a logistic regression model using stochastic gradient descent. This is a useful example because, unlike linear regression via the LMS algorithm, the moments of the equilibrium distribution cannot be obtained in closed form and a perturbation method is *required*.

[7] This corrects an error in the value of β used in Leen and Friel (2011).

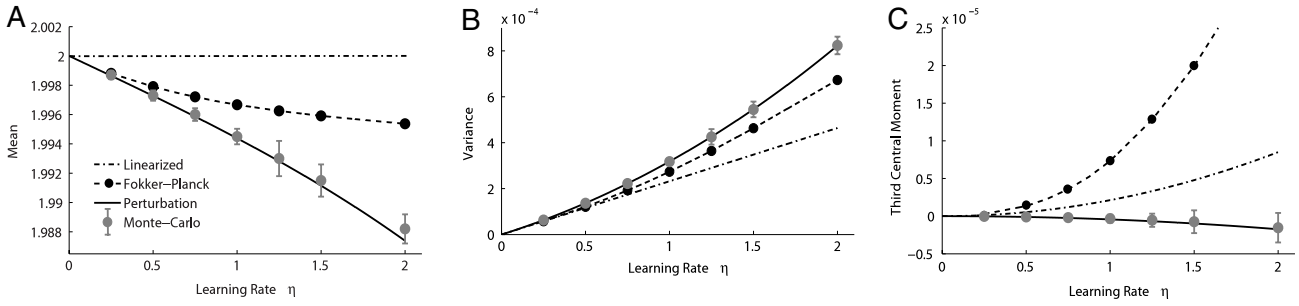

Fig. 4. Moments of the equilibrium distribution for ELL synaptic weights vs. learning rate. Shown are the moments calculated with the linearized theory, with the FPE, with the perturbation analysis and from Monte-Carlo simulations ($\pm 2\sigma$ confidence intervals). (A) Mean to 4th order in perturbation, (B) Variance to 6th order in perturbation, (C) Third central moment to 4th order in perturbation.

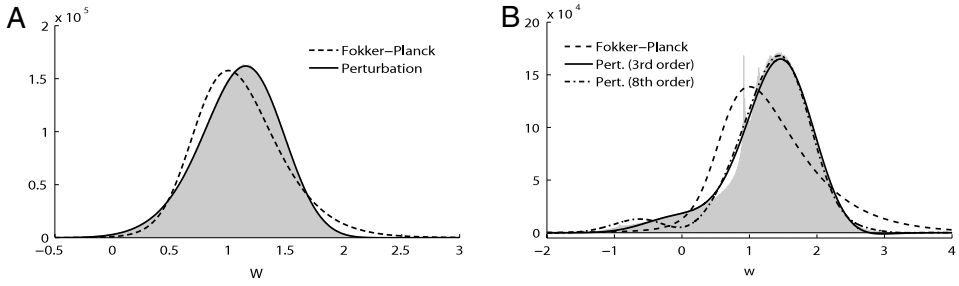

Fig. 5. Equilibrium densities for logistic regression computed from the Fokker–Planck equation, the perturbation expansion, and estimated from Monte-Carlo (shaded area). (A) Learning rate $\eta = 0.25$, perturbation to third order. (B) Learning rate $\eta = 0.75$, perturbation to third and eighth order.

The model is a simple one-parameter (w) logistic function of a scalar input feature $x \in \mathbb{R}$

$$p(b = 1|x) \equiv p(x; w) = \frac{1}{1 + \exp(-wx)}. \tag{54}$$

The parameter w is fit by maximum likelihood to the data input–output pairs $(b_i \in \{0, 1\}, x_i)$, $i = 1, \ldots$. The log-likelihood function is the familiar cross-entropy

$$L = \sum_i L_i(x_i, b_i; w)$$
$$\equiv \sum_i b_i \log(p(x_i)) + (1 - b_i) \log(1 - p(x_i)). \tag{55}$$

Under stochastic gradient ascent, the weight is re-estimated at each time step based on a single datum (x_i, b_i)

$$w_{t+1} = w_t + \eta \left. \frac{\partial L_i(x_i, b_i; w_t)}{\partial w} \right|_{w=w_t}. \tag{56}$$

For our simulations, we sample the inputs x from a uniform distribution on $[-3, 3]$ and the outputs b from a binomial distribution with $p(b = 1|x) = p(x; w_0)$ as in Eq. (54) with $w_0 = 1$. The jump moments and their derivatives $\alpha_i^{(j)}$ required for the perturbation analysis are computed by analytic averaging over b and numerical integration over the distribution on x (which motivated our use of the uniform distribution)

Fig. 5 shows the density of the equilibrium distribution at $\eta = 0.25$ and 0.75 computed from the FPE, from our perturbation expansion, and estimated from Monte Carlo. Higher order perturbation expansions (to eighth order) for the density are indistinguishable from the third order expansion for the lower learning rate in Fig. 5(A). As the learning rate is further dropped, all the densities approach a Gaussian (as discussed immediately after Eq. (27)).

As for the electric fish anti-Hebbian learning rule, the FPE gives the *incorrect sign* for the skew. The perturbation expansion gives a better estimate of the density than the FPE for both learning rates, and agrees very well with the Monte-Carlo for the lower learning rate. The agreement is not as good at the higher learning rate, and we show perturbation curves to third and to eighth order at $\eta = 0.75$. (The spikes in the histograms at $w \sim 1.0$ for the higher learning rate are real features, not sampling noise.) The oscillations in the tails of the predicted density in Fig. 5(B) grew with the increase from third to eighth order in the expansion. These recall the oscillations in the density in Fig. 1(B), and our caution that perturbation expansions are typically asymptotic but not convergent. Nonetheless, as we will see below, higher-order calculations can be quite effective for the moments. The simulation contained 50 k members in the ensemble and was burned in for 20 k time steps (moments were well-equilibrated), and the full ensemble for 400 additional time steps were used for the density histograms.

Fig. 6 shows the moments of the equilibrium distribution as calculated by the Fokker–Planck equation (through numerical integration of the density), by our perturbation expansion, and as estimated from the Monte-Carlo simulation. The perturbation expansion was carried through sixth order for the mean, and through eighth order for the variance and third moment. In all cases, the perturbation expansion produces better predictions for the moments than the FPE. At the highest learning rates the perturbation series predictions are in disagreement with the Monte-Carlo estimates, and the difference is statistically significant. (Like the results in Fig. 4, the Monte-Carlo estimates of the moments are accompanied by $\pm 2\sigma$ error bars, but they are hidden by the marker in the plots in this figure.) Increasing the order of the perturbation expansion did *not* substantively improve the theoretical predictions.

For the moment estimates, the 400 step time series of the moments estimated from the $50k$ ensemble members were subsampled at lag equal to the first zero of the auto-correlation function to produce estimates of the moments and their variances.

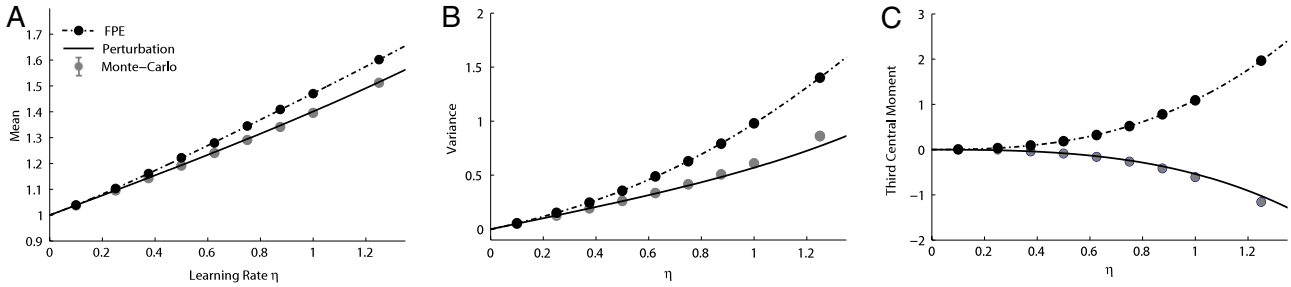

Fig. 6. Moments of equilibrium density for on-line logistic regression as functions of the learning rate η. (A) Mean through 6th order in perturbation), (B) Variance through 8th order in perturbation, (C) Third central moment through 8th order in perturbation.

4. Discussion

The perturbation methods given here provide systematic tools for approximating the density and its moments for on-line machine and STDP biological learning, and for stochastic neural dynamics, recently addressed at lowest order with the system size expansion (Buice et al., 2010). Our development gives a clear and simple order-by-order expansion that, unlike the system size expansion, gives estimates for the *density* as well as its moments.

Our implementation by computer algebra provides rapid and accurate evaluation. By hand, perturbation calculations are reasonable to undertake to low order — perhaps retaining up to the second order corrections. For the examples in this paper, we required higher order calculations to provide accurate predictions. Such calculations cannot be reliably and economically done by hand. To make our approach more accessible, we provide Mathematica code to carry out the calculations at http://www.bme.ogi.edu/~tleen/MEPerturb/.

The perturbation equations for the moments generalize in a straightforward way to multidimensional problems. Like the one-dimensional treatment given here, the moments can be evaluated *without* explicitly evaluating the corrections to the density. Evaluating the expansion for multi-dimensional *densities* requires the eigenfunctions of the corresponding multidimensional Ornstein–Uhlenbeck operators (L_0 and $L_0^\dagger$). These were previously not available in the literature, but we have recently derived them using raising and lowering operators (Leen, Friel, & Nielsen, in review).

As pointed out, perturbation expansions are typically asymptotic (ours becomes more accurate as $\eta \to 0$), but may not be convergent at a particular nonzero η. We have observed oscillations in the tails of density estimates (e.g. Fig. 5(B)) that grow with increasing order, and growth in the correction terms for *moments* with increasing order — both suggesting divergence. Although this is a common problem with asymptotic expansions, they are nonetheless abundantly useful in mathematics (for function approximation), physics (in celestial mechanics and quantum theory), and engineering (in fluid mechanics). Heuristic guidelines do exist for truncating asymptotic expansions. Bender and Orszag (1999) suggest terminating the series at the term just preceding the smallest, while Boyd (1999) suggests truncating *at* the smallest term. Neither of these heuristics proved reliable for us.

There is one exception to this convergence story. When the jump moments $\alpha_i(w)$ are polynomials of degree i or less, the moments of the equilibrium distribution $M_k(\eta)$ can be calculated in *closed form* from the Kramers–Moyal expansion. When this is the case, the perturbation expansion for the moments will agree with the power series expansion (in η) derived from the closed-form solution, since any two series (in the same scale $\eta^{1/2}$) asymptotic to the same function are identical (Bender & Orszag, 1999). For such problems, at values of η where $M_k(\eta)$ exists and its power series expansion converges, the perturbation expansion will agree and also converge.

As stated in the introduction, it is not our intent to discourage use of the FPE; numerous successes of the FPE applied to jump processes in both machine and biological learning refute van Kampen's heavy-handed claims against its utility. Instead, our work offers theoretical tools that enable prediction in regimes where the FPE fails.

We have yet to explore perturbation methods for problems in which fixed points occur at boundaries of the weight space. Such problems are discussed in the STDP literature by Câteau and Fukai (2003) and Kepecs et al. (2002). Nor have we explored application to STDP rules with triplet and quadruplet interactions as discussed for example by Pfister and Gerstner (2006).

Acknowledgments

This work was supported by the NSF under grant IIS-0812687 and in part by grant IIS-0827722. The authors thank Joel Franklin, Michael Friedlander, and Ignacio Sáez for discussion. Patrick D. Roberts was helpful in discussions of the mormyrid learning rule and our results. Frank Adrian prepared the initial Monte-Carlo simulations for the van Rossum model. We thank the organizers of the IJCNN 2011 for inviting our contribution to this special issue of *Neural Networks* and an anonymous reviewer for comments.

References

Abramowitz, M., & Stegun, L. (1972). *Handbook of mathematical functions*. US Department of Commerce, National Bureau of Standards.
Adrian, F. A. (2008). *Exact ensemble dynamics for spike-timing-dependent plasticity. Master's thesis*. Department of Science & Engineering, OHSU.
Babadi, B., & Abbott, L. (2010). Intrinsic stability of temporally shifted spike-timing dependent plasticity. *PLoS Computational Biology*, *6*, 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. *Journal of Neurophysiology*, *47*, 1043–1056.
Bell, C., Bodznick, D., Montgomery, J., & Bastian, J. (1997a). The generation and subtraction of sensory expectations within cerebellum-like structures. *Brain Behavior and Evolution*, *50*(Suppl. 1), 17–31.
Bell, C., Han, V., Sugawara, Y., & Grant, K. (1997b). 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.
Blinnikov, S., & Moessner, R. (1998). Expansions for nearly Gaussian distribution. *Astronomy and Astrophysics Supplement Series*, *130*, 193–205.
Boyd, J. P. (1999). The devil's invention: asymptotic, superasymptotic and hyperasymptotic series. *Acta Applicandae Mathematicae*, *56*, 1–98.
Buice, M. A., Cowan, J. D., & Chow, C. C. (2010). Systematic fluctuation expansion for neural network activity equations. *Neural Computation*, *22*, 377–426.
Câteau, H., & Fukai, T. (2003). A stochastic method to predict the consequence of arbitrary forms of spike-timing-dependent plasticity. *Neural Computation*, *15*, 597–620.
Der, R., & Villmannn, T. (1993). *Dynamics of self organized feature mapping. Technical report*. Universitat Leipzig, Inst. f. Informatik.

Gardiner, C. (2009). *Stochastic methods, a handbook for the natural and social sciences* (4th ed.). Berlin: Springer-Verlag.
Guckenheimer, J., & Holmes, P. (1983). *Nonlinear oscillations, dynamical systems, and bifurcations of vector fields*. New York: Springer-Verlag.
Hansen, L. K. (1993). Stochastic linear learning: exact test and training error averages. *Neural Networks*, 393–396.
Haykin, S. (1991). *Adaptive filter theory* (2nd ed.). Englewood Cliffs, NJ: Prentice Hall.
Kepecs, A., van Rossum, M. C., Song, S., & Tegney, J. (2002). Spike-timing-dependent plasticity: common themes and divergent vistas. *Biological Cybernetics*.
Leen, T. K., & Friel, R. (2011). Perturbation theory for sotchastic learning dynamics. In *Proc. IJCNN 2011*. IEEE Press.
Leen, T. K., & Friel, R. (2012). Stochastic perturbation methods for spike-timing-dependent plasticity. *Neural Computation*, *24*(5).
Leen, T. K., Friel, R., & Nielsen, D. (2012). Eigenfunctions of the *N*-Dimensional Ornstein–Uhlenbeck operator (in review).
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.
Orr, G. B., & Leen, T. K. (1993). Weight space probability densities in stochastic learning: II. Transients and basin hopping times. In Giles, Hanson, & Cowan (Eds.), *Advances in neural information processing systems*: *Vol. 5*. San Mateo, CA: Morgan Kaufmann.
Pfister, J. P., & Gerstner, W. (2006). Triplets of spikes in a model of spike timing-dependent plasticity. *Journal of Neuroscience*, *26*, 9673–9682.
Radons, G. (1993). On stochastic dcynamics of supervised learning. *Journal of Physics. A.*, *26*, 3445–3461.
Risken, H. (1989). *The Fokker–Planck equation*. Berlin: Springer-Verlag.
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.
Rubin, J., Lee, D., & Sompolinsky, H. (2001). Equilibrium properties of temporally asymmetric Hebbian plasticity. *Physical Review Letters*, *86*, 364–367.
van Kampen, N. (2007). *Stochastic processes in physics and chemistry* (3rd ed.). Amsterdam: North-Holland.
van Kampen, N. (1961). A power series expansion of the master equation. *Canadian Journal of Physics*, *39*, 551–565.
van Rossum, M., Bi, G., & Turrigiano, G. (2000). Stable Hebbian learning from spike timing-dependent plasticity. *The Journal of Neuroscience*, *20*, 8812–8821.
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*.
Zhang, L. I., Tao, H. W., Holt, C. E., Harris, W. A., & ming Poo, M. (1998). A critical window for cooperation and competition among developing retinotectal synapses. *Nature*, *395*, 37–42.
Zhu, L., Lai, Y., Hoppensteadt, F., & He, J. (2006). Cooperation of spike timing-dependent and heterosynaptic plasticities in neural networks: a Fokker–Planck approach. *Chaos*, *16*, 023105.

Please cite this article in press as: Leen, T. K., et al. Approximating distributions in stochastic learning. Neural Networks (2012), doi:10.1016/j.neunet.2012.02.006