

Learning and Variability in Spiking Neural Networks

Jeffrey J. Rodny (jrodny@ucmerced.edu)
Christopher T. Kello (ckello@ucmerced.edu)

University of California, Merced,
Department of Cognitive and Information Sciences,
5200 North Lake Road, Merced, CA 95343 USA

Abstract

Neural networks exhibit ongoing, spatiotemporal patterns of spiking activity. Evidence shows that these patterns are metastable, i.e. temporary, transient, and non-stationary. Metastability is theorized to be adaptive for neural and cognitive function, but learning must somehow remain stable in the context of highly variable spike dynamics. In the present study, a neural network learning algorithm is developed to co-exist with intrinsic variability that arises from regulating spike propagation to stay near its critical branching point. The learning algorithm is based on *reinforcement traces* stored at synapses that change much more slowly than synaptic switches triggered to maintain critical branching. As a result, learning establishes a stable synaptic space within which variability and metastability can arise from critical branching. Model efficacy is demonstrated using time-delayed XOR learning, and spike dynamics are compared with evidence of metastability in hippocampal recordings.

Keywords: Spiking neural network, critical branching, metastability, learning mechanism.

Introduction

All neural and behavioral activities are characterized by intrinsic variability—variations that are not attributable to forces outside the system in question (Kello, Beltz, Holden, & Van Orden, 2007). These variations are found to follow power laws of various kinds (Kello, Brown, Ferrer-i-Cancho, Holden, Linkenkaer-Hansen, Rhodes, & Van Orden, 2010), which means they cannot be discounted as Gaussian noise. Kello (2013) recently reported a spiking neural network model that exhibits pervasive power laws in its intrinsic dynamics. It does so by virtue of a simple, local mechanism designed to regulate spike propagation. Such homeostatic regulation is a basic, necessary function of any network with propagating activity. Too much or too little activity leads to sub-optimal communication and transport over networks (Beggs & Plenz, 2003). For cognitive science, this means diminished memory and computational capacity (Bertschinger & Natschläger, 2004).

A long-standing puzzle in neuroscience is how homeostatic regulation and its concomitant variability in spike dynamics are coordinated with learning (Turrigiano & Nelson, 2000). Both regulation and learning are expressed through potentiation and de-potentiation of synapses, which suggests these functions should be prone to interfering with each other. For instance, long-term potentiation of excitatory synapses may instantiate learning while also increasing overall spike rates. Learning may then be undone when de-potentiation is triggered by regulatory

mechanisms to lower spike rates. Some learning mechanisms have been proposed that inherently regulate spike activity, most notably spike timing dependent plasticity (Markram, Lübke, Frotscher, & Sakmann, 1997). However, STDP and similar learning mechanisms have not been shown to produce the pervasive power laws that characterize intrinsic variations of neural and behavioral activity.

Let us approach the puzzle from the starting point of homeostatic regulation instead. Regulatory mechanisms have been shown to produce power laws, particularly mechanisms based on critical branching (Beggs & Plenz, 2003). Critical branching is a general dynamic of event propagation over networks, whereby each event on a given node triggers exactly one subsequent event on a downstream node, on average. This general dynamic leads a very specific power law—the sizes of contiguous event cascades (i.e. number of propagated events till extinguishing) should follow an inverse power law distribution with an exponent of two, $P(S) \sim 1/S^{a=2}$. Beggs and Plenz referred to such cascades as *neuronal avalanches* and found evidence for them in multi-cell recordings of rat somatosensory cortex. Many studies have since found the same power law in many kinds of neural activity, including human electrophysiology (Poil, van Ooyen, & Linkenkaer-Hansen, 2008).

Critical branching is relevant to cognitive science because it serves to not only regulate spike propagation, but also maximize memory and computational capacities of spike dynamics. Beggs and Plenz (2003) showed this consequence with an abstracted probabilistic model, and Kello (2013) showed it in a spiking neural network model using reservoir computing techniques. Kello formulated a general, biologically plausible mechanism that maintains network-wide critical branching using only information local to a given neuronal unit and its synapses. The mechanism was shown to generate pervasive power laws, i.e. not only neuronal avalanches, but also power law inter-spike intervals (ISIs), and 1/f noise in neural activity as well as simulated behavioral response dynamics.

In summary, Kello's (2013) model simulated homeostatic regulation and its concomitant variability, and it related critical branching to metastability and cognitive function. However, the model did not address learning. In the present study, a reinforcement learning mechanism is formulated to work in conjunction with critical branching in a spiking neural network. The network is shown to learn a temporal nonlinear function of spike inputs in the face of ongoing variability due to ongoing actions of the critical branching

mechanism. This variability is shown to be metastable by comparison with multi-cell hippocampal recordings (Sasaki, Matsuki, & Ikegaya, 2007), yet learning is stable even after reinforcement signals are removed and synapses continue to change as a function of critical branching. Thus the model provides one solution to the puzzle of homeostatic regulation, variability, and learning in spiking neural networks.

Methods

The model is composed of excitatory and inhibitory leaky integrate-and-fire (LIF) neurons, with both feed forward and recurrent synaptic connections. As in Kello (2013), connectivity is arranged so that input spikes impinge on a set of source units, propagate to a set of reservoir units, and then circulate and propagate to a set of sink units. Model architecture is shown in Fig 1. Model parameters were chosen as defaults, based on preliminary simulations. Further work is needed to investigate the impacts of certain parameter choices on performance, but in general, small variations have little or no effect on results.




Figure 1: Model diagram with arrows showing connectivity, with each unit in the sending group connected to each in the receiving group with a probability of 0.1. Source units represented input bit values 0 or 1 (half and half), and sink units represented the XOR function on consecutive input bits. All source and sink units were excitatory.

The critical branching algorithm relies on temporal ordering of spikes, and so the model was simulated asynchronously, unlike most clock-based models. Here we describe the model algorithmically, because its temporal discontinuities would complicate an analytic description. Pseudo-code for the algorithm is diagrammed in Figure 2.




Figure 2: Pseudo-code for spike propagation and regulation guided by learning, overlaid on a diagram showing one neuronal unit i (circle) with its input j and output k synapses (triangles). Synapses are enabled or disabled (filled or unfilled), and an example chain of events is shown for one synaptic potential δ_j at time t . Other variables and functions are explained in the text.

LIF dynamics Each neuronal unit has a membrane potential (“activation”) that is updated instantaneously according to $V(t)$ update equation in Figure 2, where t and t' are the current and previously updated times, i and j are indices for units and incoming synapses, λ is the decay rate, and $\delta(t)$ is the instantaneous input. Inputs come from synaptic potentials for reservoir and sink units, and from external sources for source units. A spike (action potential) occurs whenever membrane potential crosses threshold, $V(t) > \theta$ with $\theta = 1$. $V(t)$ is immediately reset to its resting state $\zeta = 0.5$ and remains in its refractory period for 1 time interval (denoted by Γ), during which no inputs are applied. Synaptic potentials i.e. weights were held constant at $v = \pm 0.75$ for excitatory/inhibitory synapses. This absolute magnitude was chosen to be somewhat under threshold so that spikes would often but not always be triggered or inhibited. External inputs were set to always trigger a spike.

Each spike also triggers events over its unit’s synapses that serve to propagate spikes, and also regulate propagation according to critical branching, and guide propagation according to reinforcement signals. With synaptic strengths held constant, additional variables are needed to support homeostatic regulation and learning. We introduced three into our simplified LIF model, described next.

Critical Branching First, each synapse can be either disabled or enabled, $A = 0$ or 1 . When a spike occurs, synaptic potentials are triggered over a unit’s output (axonal) synapses at a delay of $t+\tau$ (τ sampled uniformly

between 1 and 2), but *only* for enabled synapses $A = 1$. The model is initiated with all synapses disabled, and the critical branching algorithm enables (and eventually disables) synapses as spikes propagate through the network.

The objective function of the algorithm is for each unit to be blamed for exactly one spike across its output synapses, during each of its ISIs. This objective embodies critical branching, i.e. exactly one spike propagated for every spike generated. The objective function is achieved by adding a unit variable B that counts the number of times a unit is blamed during each of its ISIs. When a spike occurs on unit i , one of its enabled input (dendritic) synapses is chosen, and blame B_j for corresponding unit j is incremented. Also unit i 's blame is checked. If unit i has not been blamed since its last spike ($B_i < 1$), then one of its disabled axonal synapses is enabled with probability 0.05. If unit j has been blamed more than once ($B_j > 1$), then input synapse j is disabled with probability 0.05. Finally, B_i is reset to 0.

Over time, synapses will be enabled when units are propagating less than one spike per spike on average, and disabled when units are propagating more than one spike per spike on average. The algorithm as just described is a canonical type of critical branching. It leaves unspecified *which* synapses are chosen to switch between enabled/disabled states. Kello (2013) showed that this choice can be random or based on a rule biased towards choosing recently switched synapses. Thus our algorithm has a natural free parameter for learning—instead of choosing synapses randomly or based on recency, they can be chosen on the basis of learning signals, as described next.

Reinforcement Learning To integrate a simple learning algorithm with critical branching, each synapse is provided with a reinforcement signal R that directly rewards or punishes sets of synaptic potentials associated with individual units (analogous to extracellular dopamine released to signal reward or stimuli predictive of reward; Schultz, Dayan, & Montague, 1997). Excitatory potentials are rewarded when downstream spikes are signaled, and punished when not. Inhibitory potentials are conversely punished or rewarded. Reinforcement signals do not have direct effects on enabling/disabling synapses. Instead, a running average is stored as a reinforcement trace C on each synapse, with the weighting of its current value set to 0.9 (i.e. traces change relatively slowly over time). Traces are updated only when a reinforcement signal is present (indicated by Ω). Such “synaptic tags” or traces have been suggested to be vital in the long term potentiation of neurons (Frey & Morris, 1997). Similar reinforcement traces have been hypothesized previously as a biologically plausible and effective learning mechanism for spiking networks with STDP (Izhikevich, 2007). However, previous studies did not integrate traces with critical branching or other homeostatic regulation mechanisms, to our knowledge.

The critical branching algorithm determines *when* a synapse should be chosen for enabling or disabling, and reinforcement traces determine *which* synapses are chosen

for enabling and disabling. Very simply, the disabled output synapse with the highest trace value is chosen for potential enabling, and the enabled input synapse with the lowest trace value is chosen for potential disabling. A small amount of noise (ε sampled uniformly with ± 0.1) is included with the trace value at each choice point. The function of this noise is to implement a random choice when reinforcement signals are very weak or unavailable (i.e. when $C \sim 0$).

Temporal XOR Classification

The spiking network model presented herein is very general in terms of network architecture and learning task. Any pattern of connectivity can be specified, as long as there is an external source to drive activity, and a sink where activity can (eventually) exit the network. Any synapse can receive reinforcement signals, and any schedule of direct reinforcement can be applied. Currently the model handles delayed reinforcement only to the extent that spike dynamics have memory. That is, learning will be contingent on the past to the extent that effects of past events are reflected in current synaptic potentials.

Kello (2013) showed that critical branching maximizes the “fading memory” property of spike dynamics as tested using the paradigm of reservoir computing (Maass, Natschläger, & Markram, 2002). The model had the same three-group architecture as herein, and external inputs drove recurrent spike dynamics in the reservoir. Memory was tested by examining whether the reservoir spike pattern during a given time interval T held information about past input patterns $T-\tau$. The model’s computational capacity was tested by examining whether this information could be used to compute a nonlinear function of past inputs. In particular, successive input patterns were treated as successive input bits (i.e. categorized as input pattern 0 or input pattern 1), and least squares regression was used to compute the XOR function on past pairs of input bits using only the current reservoir spike pattern. XOR accuracy fell off slowly as function of τ , most slowly when spike dynamics were near their critical branching point (rather than sub- or super-critical).

Reservoir computing methods were useful for demonstrating the memory and computational capacities of critical branching, but least squares regression is not a neural learning mechanism. More problematically, reservoir computing methods require stationary spike dynamics, which means that Kello (2013) had to disengage the critical branching algorithm at test, because the algorithm creates non-stationary spike dynamics. These dynamics reflect the metastability property of critical branching, which is hypothesized to benefit cognitive function and needed to account for neuroscience evidence (Tognoli & Kelso, 2014).

Here we applied the temporal XOR classification test as implemented by Kello (2013), but least squares regression was replaced with the learning mechanism described in the previous section. The stability of learning was tested by engaging the critical branching algorithm while

reinforcement signals were applied, and also while reinforcement signals were subsequently removed.

Each input bit was converted into a sequence of 20 spikes over source units (evenly spaced over one unit time interval), where half of the units represented bit 0 and half bit 1. Spikes were always sequenced in the same order for each bit, and random bit sequences were input to the model. Reinforcement signals were applied to sink units at a delay of 3-4 unit time intervals from the corresponding input bits. For example, reinforcement signals for the XOR output 1 were applied at time interval T when the two input bits were both 0 or both 1 at time intervals T-3 and T-4. Half the sink units represented XOR output 0 and the other half represented XOR output 1. A reinforcement signal of +1 was applied to input synapses of sink units representing the correct output for each time interval T, and -1 for incorrect outputs. As with other reservoir computing models, performance drops off for larger delays, with delays larger than 9-10 unit time intervals nearing chance performance.

Results

Sixty simulations were run using the model architecture shown in Figure 1. Each simulation was run for 200,000 time intervals, with different random initializations of parameters. Critical branching and reinforcement learning were always engaged for the first 40,000 time intervals. There were three different conditions for the remaining 160,000 time intervals (20 simulations per condition). For the “+CB+Rwd” condition, both critical branching and reinforcement learning continued to be engaged throughout. For the “+CB-Rwd” condition, critical branching continued to be engaged, but reinforcement learned was stopped. For the “-CB-Rwd” condition, both mechanisms were disengaged which means that synapses were held constant. There was no “-CB+Rwd” condition because our reinforcement learning cannot be engaged without engagement of critical branching. Model performance was always tested after the initial 40,000 time intervals.

Average XOR accuracies are shown in Figure 3, where chance performance is 0.5. Performance can be seen to ramp up to ~95% in all conditions, and remain there for the duration of the simulations. These results demonstrate the efficacy of the reinforcement learning algorithm, despite ongoing variability due to homeostatic regulation. Perhaps most impressive was steady performance near the 95% level even after reinforcement signals were stopped (+CB-Rwd). The reason for this steady performance is that the values of reinforcement traces remain constant so long as there are no reinforcement signals to trigger updates. Therefore, the critical branching algorithm continues to enable/disable synapses in accordance with reinforcement traces laid down for temporal XOR classification. It is this combination of stable traces and synaptic switching that allows stable learning to co-exist with homeostatic regulation and its concomitant variability.




Figure 3: Mean XOR accuracies as a function of time, shown for each of the three model conditions. Dashed line shows separation of the initial tuning/learning period, and the subsequent period during which conditions were distinguished.




Figure 4: Spectral analysis of fluctuations in reservoir spike counts per unit time, shown in log-log coordinates, with logarithmically spaced spectral bins.

We now move on to results concerning the variability and metastability of spike dynamics. First we examine whether the 1/f power law reported by Kello (2013) is replicated when reinforcement learning is integrated with critical branching (in the interest of space, we do not report results for the other power laws). As in the original study, the number of reservoir spikes was counted per unit time interval, creating a time series of spike counts for each simulation. Spectral analysis was conducted on each time

series, and spectra were averaged and plotted in Figure 4. As can be seen, the present model replicated the 1/f power law in lower frequencies that is observed while critical branching is engaged. Fluctuations become random and uncorrelated when critical branching is disengaged, as in the previous model without reinforcement learning.

1/f fluctuations and other power laws have been associated with metastability, but this hypothesized property of spike dynamics was not examined in Kello (2013). Here we test metastability using methods applied in a previous study of hippocampal spike dynamics (Sasaki et al., 2007). The authors used functional multi-neuron imaging to record spontaneous network activity in rat hippocampal slice cultures. Principal component analysis revealed that spike patterns varied over a diverse but organized set of broadly defined pattern spaces. Moreover, each space was visited only once or a small number of times. It is this transitioning between diverse sets of patterns, without settling into any one of them, that defines metastability.




Figure 5: Spike trains from spontaneous hippocampal recordings [top; (Sasaki et al., 2007)], and from reservoir neurons in our spiking neural network model (middle). At bottom are shown example series of bit inputs, S_0 & S_1 that drive reservoir spiking, with a sample of corresponding reservoir and output spikes, R and X_0 & X_1 .

To give the reader a visual sense of both empirical and model spike dynamics, Figure 5 shows spike trains from the study by Sasaki et al. (2007), along with spike trains from model in the +CB+Rwd condition. Heterogeneity in patterning can be seen for both model and empirical data, in terms of locally correlated patterns of clustering that change

over time. For the model, one can also see input bits as ordered sequences over two sets of units, as well as the relative stability of sink units compared with reservoir units. This comparison demonstrates that stable learning can co-exist with metastable dynamics.

Metastable dynamics are expressed more clearly through auto-correlation and principle components analysis (PCA). If spike patterns are organized locally in time, but transition through different pattern spaces over time, then auto-correlation of the spike pattern time series should reveal temporally local correlations but a lack thereof at more distant time delays. PCA is one way to visualize the hypothesized transitions through different pattern spaces. In particular, by projecting dynamics onto the first two principle components of a given spike pattern time series, metastability should be visualized as organized movement through different regions of the 2D space.




Figure 6: Auto-correlation (left) and PCA (right) analyses of spike pattern time series for +CB+Rwd model (top), hippocampal data from Sasaki et al. (middle), and the -CB-Rwd model (bottom). Temporal windows were 10 s for hippocampal data, and 10 time intervals for model data. For auto-correlation analyses of the model (top & bottom), one axis unit of time represents 20,000 time intervals.

Sasaki et al. (2007) quantified spike patterns as vectors of windowed spike counts over hippocampal neurons recorded during spontaneous activity. Successive windows created

time series of spike patterns, and the middle row of Figure 6 shows the results of auto-correlation and PCA. In the time period visualized, one can see three or four distinct pattern spaces. In auto-correlation, they appear as square regions of high correlation around the diagonal. In PCA, they appear as successive clusters of points in the 2D space.

We conducted the same analyses on reservoir units in our model. It is debatable whether activity driven by input bits is comparable to spontaneous hippocampal activity, but the same basic results hold for the spontaneous input conditions examined in Kello (2013; not reported here). In any case, the model shows the same basic earmarks of metastability, but only when critical branching is engaged. When critical branching is disengaged, spike dynamics fluctuate randomly within one region of the pattern space, as evidenced by high auto-correlations throughout time, and one cluster of random movements in PCA space.

Discussion

The model presented herein provides one solution to the puzzle of how learning can be stable in the face of ongoing variability and metastability. The model is cast as a biologically plausible spiking neural network, but the principles and mechanisms may be applied to complex adaptive networks in general. The model's intrinsic, power law variability derives from homeostatic regulation that supports and enhances memory and computation. The model is designed so that regulatory enabling and disabling of synapses does not interfere with changes to reinforcement traces driven by learning. Instead, these synaptic traces act as a stable memory for learning. This memory carves out a broad portion of synaptic space in which performance is maintained. Critical branching serves to drive synaptic changes within this space. As a result, spike patterns form and reform over time, but always within the space carved out by learning.

The present study shows how learning can be integrated with homeostatic regulation and metastability, but further work is needed to investigate how the latter might enhance the former. Based on the current results, it is safe to say that learning becomes highly redundant as a result of critical branching. That is, the same XOR function was accomplished using many different sets of neurons and spike patterns. Thus the variability from homeostatic regulation may enhance the robustness of learning (Kitano, 2004). Previous studies also have associated metastability with the flexibility and context-sensitivity of cognitive function (Kello & Van Orden, 2009). The demonstrated redundancy may provide a way for a given process or representation to emerge differently in different contexts, yet achieve the same underlying function. This is one of a number of different future directions.

Acknowledgments

This research was supported in part by awards from the National Academies Keck Futures Initiative, the National

Science Foundation (BCS 0842784 and 1031903), and DARPA under contract No. HR0011-09-C-0002.

References

Beggs, J. M., & Plenz, D. (2003). Neuronal avalanches in neocortical circuits. *The Journal of Neuroscience*, 23, 11167-11177.

Bertschinger, N., & Natschläger, T. (2004). Real-time computation at the edge of chaos in recurrent neural networks. *Neural Computation*, 16(7), 1413-1436.

Frey, U., & Morris, R. G. M. (1997). Synaptic tagging and long-term potentiation. *Nature*, 385(6616), 533-536.

Izhikevich, E. M. (2007). Solving the Distal Reward Problem through Linkage of STDP and Dopamine Signaling. *Cerebral Cortex*, 17(10), 2443-2452.

Kello, C. T. (2013). Critical branching neural networks. *Psychological Review*, 120(1), 230-254.

Kello, C. T., Beltz, B. C., Holden, J. G., & Van Orden, G. C. (2007). The emergent coordination of cognitive function. *Journal of experimental psychology. General*, 136(4), 551-568.

Kello, C. T., Brown, G. D. A., Ferrer-i-Cancho, R., Holden, J. G., Linkenkaer-Hansen, K., Rhodes, T., & Van Orden, G. C. (2010). Scaling laws in cognitive sciences. *Trends in Cognitive Sciences*, 14(5), 223-232.

Kello, C. T., & Van Orden, G. C. (2009). Soft-assembly of sensorimotor function. *Nonlinear Dynamics, Psychology, and Life Sciences*, 13(1), 57-78.

Kitano, H. (2004). Biological robustness. *Nat Rev Genet*, 5(11), 826-837.

Maass, W., Natschläger, T., & Markram, H. (2002). Real-Time Computing Without Stable States: A New Framework for Neural Computation Based on Perturbations. *Neural Computation*, 14(11), 2531-2560.

Markram, H., Lübke, J., Frotscher, M., & Sakmann, B. (1997). Regulation of Synaptic Efficacy by Coincidence of Postsynaptic APs and EPSPs. *Science*, 275(5297), 213-215.

Poil, S.-S., van Ooyen, A., & Linkenkaer-Hansen, K. (2008). Avalanche dynamics of human brain oscillations: Relation to critical branching processes and temporal correlations. *Human Brain Mapping*, 29(7), 770-777.

Sasaki, T., Matsuki, N., & Ikegaya, Y. (2007). Metastability of Active CA3 Networks. *The Journal of Neuroscience*, 27(3), 517-528.

Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, 275(5306), 1593-1599.

Tognoli, E., & Kelso, J. A. S. (2014). The Metastable Brain. *Neuron*, 81(1), 35-48.

Turrigiano, G. G., & Nelson, S. B. (2000). Hebb and homeostasis in neuronal plasticity. *Current Opinion in Neurobiology*, 10(3), 358-364.