

Local self-adaptation mechanisms for large-scale neural system building

M. García Ortiz and A. Gepperth

Honda Research Institute Europe GmbH
Carl-Legien-Str.30, 63073 Offenbach, Germany
garciaom@minatec.inpg.fr, alexander.gepperth@honda-ri.de,
<http://www.honda-ri.de>

Abstract. For integrating neural networks into large systems, dynamical stability and parameter settings are key issues, especially for popular recurrent network models such as dynamic neural fields. In neural circuits, homeostatic plasticity seems to counter these problems. Here we present a set of gradient adaptation rules that autonomously regulate the strength of synaptic input and the parameters of the transfer function for each neuron individually. By doing this, we actively maintain the average membrane potentials and firing rates as well as the variances of the firing rate at specified levels. A focus of this contribution lies on clarifying at which time scales these mechanisms should work. The benefit of such self-adaptation is a significant reduction of free parameters as well as the possibility to connect a neural field to almost arbitrary inputs since dynamical stability is actively maintained. We consider these two properties to be crucial since they will facilitate the construction of large neural systems significantly.

1 Introduction

It is well known that even single neurons are complex, nonlinear dynamical systems (see, e.g., [1]). Furthermore, neurons are massively interconnected with other neurons by (possibly recurrent) synaptic connections, with their own nonlinear behavior. To maintain dynamical stability under such circumstances, there exist a multitude of activity control mechanisms [2,1,3], which autonomously adapt the processing parameters of each neuron according to local rules. These mechanisms are collectively denoted as "homeostatic plasticity". Popular neuron models are facing similar stability problems, especially when researchers construct large-scale neural systems [4,5,6,7]. In addition, most neuron models contain a multitude of free parameters which cannot always be related to experimental findings.

For both reasons and based on our previous experiences with large-scale neural systems [4], this contribution focuses on local activity control mechanisms for a well-known network model, the dynamic neural field model [8,9,10]. As a first step, we show how the membrane potentials (related to synaptic scaling [3]) and the firing rates (related to intrinsic plasticity [11,12]) can be autonomously

regulated in face of variable afferent input, thereby maintaining the temporal mean and variance of individual membrane potentials or firing rates at specified target values. Due to the nonlinear dynamic character of the neural field model, it is important to determine time scales such as to minimize interference between adaptation mechanisms.

Several mechanisms of homeostatic plasticity have been previously modeled: on the one hand, intrinsic plasticity, the adaptation of the intrinsic excitability of a neuron's membrane, has been modeled as an adaptation of neural transfer functions in [13,14] and applied to a number of problems such as reservoir computing [15] or self-organization of sensory representations [16]. On the other hand, authors have modeled synaptic scaling [17], the activity-dependent modification of synaptic strengths [2].

Overall, our work differs from related work in two respects. Firstly, in addition to modeling synaptic scaling and intrinsic plasticity, we demonstrate the operation of these mechanisms concurrently with each other and, secondly, we present a strategy of decoupled time scales to prevent interference.

More precisely, the main difference between our work and [13,16,14] is the use of recurrent neural networks with a dynamic state instead of input-output encoding neurons. Not needing to address stability problems, these articles strongly focus on achieving a certain output distribution for each neuron. In contrast, we emphasize the reduction in the number of free parameters as well as the dynamic stability. In this respect, our work is related to [17] which also employs the dynamic neural field model, although the focus of our work is on the effects of activity control rather than on self-organization processes.

2 Dynamic Neural Fields

The dynamic neural field model [8] has been proposed to describe pattern formation in the visual cortex. Essentially, dynamic neural fields are a class of recurrent neural network models that have been extensively used for modeling cognitive phenomena like decision making [10], motor planning [9], spatial cognition [18], eye movement preparation [19,20] and object recognition [21,22]. Basic elements are simple dynamic-state neurons, a fixed lateral connectivity, and a (usually sigmoid) nonlinearity.

In the neural field model described in [8], natural candidates for self-adaptation are the strengths of afferent inputs and the transfer function parameters, which need to be made position and time dependent for this purpose. We thus formulate a generalized version of the original model suitable for local adaptation mechanisms:

$$\begin{aligned} \tau \dot{u}(\mathbf{x}, t) = & -u(\mathbf{x}, t) + \alpha(\mathbf{x}, t)S(\mathbf{x}, t) \\ & + \beta \int w(\mathbf{x} - \mathbf{x}') f[u(\mathbf{x}', t)] d\mathbf{x}' + \gamma \sigma(\mathbf{x}, t) + h \end{aligned} \quad (1)$$

$$\text{where } f[u(\mathbf{x}, t)] = \frac{1}{1 + \exp\left(\frac{-2(u(\mathbf{x}, t) - \theta(\mathbf{x}, t))}{\nu(\mathbf{x}, t)}\right)}$$

Here, the quantity $u(\mathbf{x}, t)$ represents the membrane potential of the field at time t and position \mathbf{x} , $S(\mathbf{x}, t)$ the afferent input, $w(\mathbf{x} - \mathbf{x}')$ the fixed lateral *interaction kernel*, $f[u]$ the nonlinearity or *transfer function*, and $\sigma(\mathbf{x}, t)$ the noise. τ determines the time scale of field evolution, and h is the *resting potential*, i.e., the equilibrium potential in case of no input. We choose a sigmoid transfer function, parameterized for each neuron by a threshold and a gain value $\theta(\mathbf{x}, t), \nu(\mathbf{x}, t)$. In addition to the original model equation [8], we introduce time and position dependent coefficients $\alpha(\mathbf{x}, t), \theta(\mathbf{x}, t), \nu(\mathbf{x}, t)$, as well as the coefficients β, γ which will not be subject to adaption for now. The coefficients $\alpha(\mathbf{x}, t), \beta$ and γ respectively determine the contribution of the afferent input, the lateral recurrent interactions and the noise. The interaction kernel $w(\mathbf{x} - \mathbf{x}')$ is usually chosen to be symmetric: $w(\mathbf{x} - \mathbf{x}') = a_0 G_{\mu=0, \sigma_{\text{on}}}(\mathbf{x} - \mathbf{x}') - b_0 G_{\mu=0, \sigma_{\text{off}}}(\mathbf{x} - \mathbf{x}') - c_0$, where $G_{\mu=0, \sigma}(\mathbf{x})$ denotes a Gaussian with mean μ and standard deviation σ , and $\sigma_{\text{on}} < \sigma_{\text{off}}$. The constants a_0, b_0, c_0 are chosen suitably to achieve the desired level of local excitation/inhibition (a_0, b_0) as well as global inhibition (c_0).

3 Experimental Setup

The configuration used for simulation experiments consists of a single neural field discretized to 128x128 neurons. Constant parameters are chosen to $\tau = 12, a_0 = 0.3, b_0 = 1.5, \sigma_{\text{on}} = 10, \sigma_{\text{off}} = 20, \beta = 1, \gamma = 0, h = -0.15$. Input patterns stay constant for one *pattern cycle* consisting of 800 iteration steps. For clearing activity from previous pattern cycles, a value of $h = -20.0$ is used in the first 150 steps of a pattern cycle. Afferent input $S(\mathbf{x}, t)$ is additively composed of uniform noise with amplitudes between 0.14 and 0.16 and Gaussians (random peak values between 0.5 and 0.7, random variances between 4.0 and 6.0) that appear, at the start of each pattern cycle, at random positions in two distinct 30x30 areas (see Fig. 1). One area always contains two Gaussians whereas an other area contains only one. Thus, we can distinguish three activity levels in the afferent input. Using fixed values of $\alpha(\mathbf{x}, t) = 1.0, \nu(\mathbf{x}, t) = 0.3$ and $\theta(\mathbf{x}, t) = 0.5$, the mean and variance distribution resulting from this input can be seen in Fig. 1.

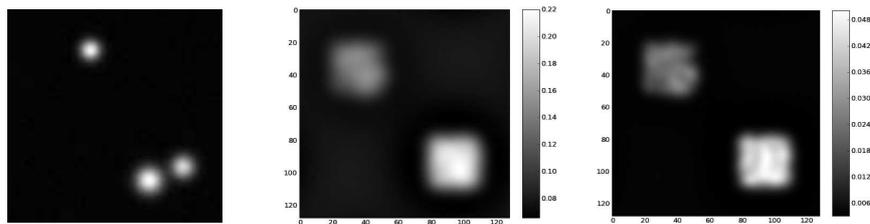


Fig. 1: Left: Example input stimulus. Middle, right: Temporal mean and variance of membrane potentials $u(\mathbf{x}, t)$ resulting from input stimuli, without adaptation.

4 Dynamic adaptation of membrane potential

In order to regulate the temporal mean value of the membrane potential $\bar{u}^\lambda(\mathbf{x}, t)$ to a target value \bar{u}_{target} , we use a gradient adaptation rule for the input strength $\alpha(\mathbf{x}, t)$ of each neuron:

$$\alpha(\mathbf{x}, t + 1) = \alpha(\mathbf{x}, t) - \epsilon_\alpha(\bar{u}^\lambda(\mathbf{x}, t) - \bar{u}_{target}). \quad (2)$$

$$\bar{u}^\lambda(\mathbf{x}, t + 1) = (1 - \lambda)\bar{u}^\lambda(\mathbf{x}, t) + \lambda u(\mathbf{x}, t). \quad (3)$$

Here, λ and ϵ_α denote the timescales at which mean calculation and adaptation take place. Importantly, λ and ϵ_α have to be properly chosen. If the adaptation is too fast, i.e. ϵ_α is too large, it has an immediate effect on the field potential, but $\alpha(\mathbf{x}, t)$ will vary constantly and not stabilize. If the adaptation is slow, $\alpha(\mathbf{x}, t)$ will converge slowly without oscillating. An aim of this section is to find a suitable combination $(\epsilon_\alpha, \lambda)$ that guarantees stable convergence of $\alpha(\mathbf{x}, t)$.

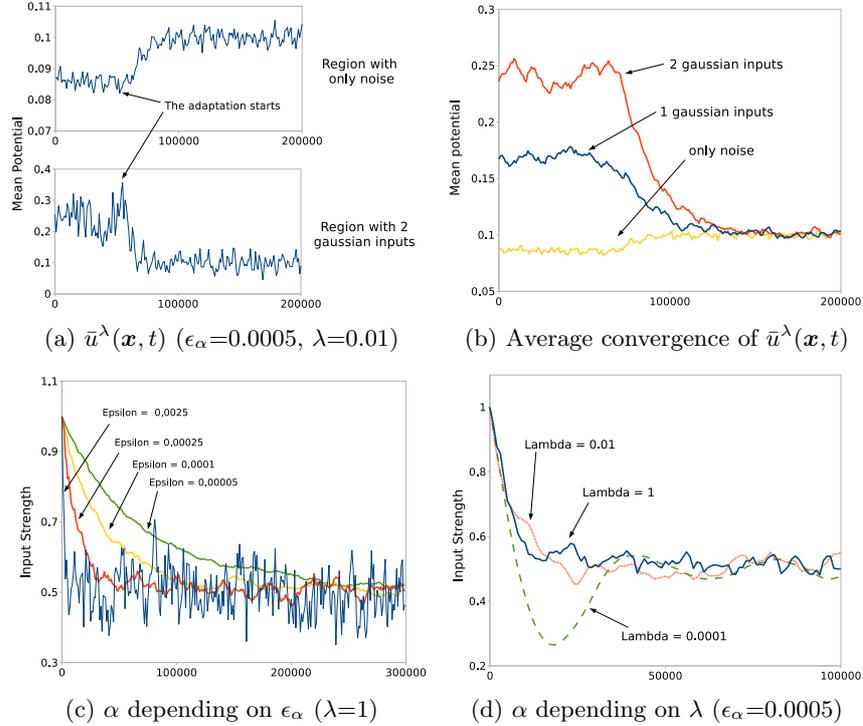


Fig. 2: Mean potential adaptation. (a) shows the convergence of the mean potential $\bar{u}^\lambda(\mathbf{x}, t)$. (b) shows the average convergence of $\bar{u}^\lambda(\mathbf{x}, t)$ for a set of 10 experiments, with the same parameters as in (a). Fig. (c) and (d) show the convergence of $\alpha(\mathbf{x}, t)$ depending on ϵ_α and λ , in region with 2 gaussian inputs.

4.1 Experiments and results

Considering the distribution of the mean potential without adaptation (Fig. 1), we set the target mean potential \bar{u}_{target} to 0.1 and study the convergence of the mean potential depending on ϵ_α and on λ .

The results in Fig. 2 show that we can successfully adapt the mean potential using the gradient adaptation rule. If the adaptation is on a faster time scale than the mean calculation, the input strength will oscillate, as can be observed in Fig. 2d with $\lambda=0.0001$. Therefore, a general rule for input strength adaptation is to set different time scales satisfying $\lambda \gg \epsilon_\alpha$.

5 Dynamic adaptation of firing rates

This section describes how to adapt the parameters $\theta(\mathbf{x}, t)$ and $\nu(\mathbf{x}, t)$ of the transfer function in order to control the mean and the standard deviation $\Sigma_f^\rho(\mathbf{x}, t)$ of each neuron's firing rate. Mathematically, we can express a running estimate on time scale ρ of the firing rate's mean and standard deviation as:

$$\bar{f}^\rho(\mathbf{x}, t+1) = (1 - \rho)\bar{f}^\rho(\mathbf{x}, t) + \rho f(\mathbf{x}, t) \quad (4)$$

$$\Sigma_f^\rho(\mathbf{x}, t+1) = (1 - \rho)\Sigma_f^\rho(\mathbf{x}, t) + \rho \sqrt{(f[u(\mathbf{x}, t)] - \bar{f}^\rho(\mathbf{x}, t))^2} \quad (5)$$

For the dynamic firing rate adaptation, we adapt the threshold $\theta(\mathbf{x}, t)$ to the mean value of the potential:

$$\theta(\mathbf{x}, t+1) = \theta(\mathbf{x}, t) - \epsilon_\theta(\theta(\mathbf{x}, t) - \bar{u}^\lambda(\mathbf{x}, t)) \quad (6)$$

By doing so, we center the transfer function on the mean value of the potential distribution. This is essential if we want to adapt the gain efficiently. Next, as illustrated in Fig. 3, we adjust the gain in order to adapt $\Sigma_f^\rho(\mathbf{x}, t)$:

$$\nu(\mathbf{x}, t+1) = \nu(\mathbf{x}, t) - \epsilon_\nu(\Sigma_f^\rho(\mathbf{x}, t) - \Sigma_{target}) \quad (7)$$

5.1 Experiments and results

For simulations, we use the following values: $\Sigma_{target} = 0.015$, $\rho = 0.01$, $\epsilon_\theta = 0.0001$ and $\epsilon_\nu = 0.00001$. As one can see in Fig.4b, the transfer function adaptation is efficient for areas with strong input (2 gaussian inputs). However, for the case of a region with a weak input (only noise), the activity may die out due to lateral interactions (see Fig. 4a) . This can be avoided by combining the firing rate adaptation with the mean potential adaptation (see section 6)

VI

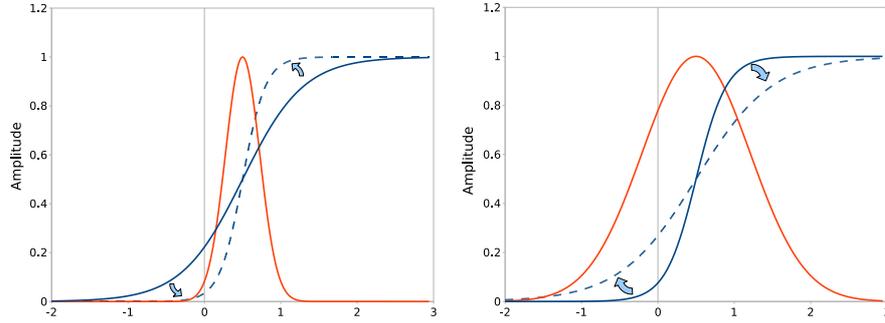


Fig. 3: Mechanisms of gain adaptation. Red solid curves indicate the distribution of membrane potential values, blue solid and dashed curves indicate the current and the optimal transfer function. As the gain is inversely proportional to the slope of the sigmoid function, increasing (left) or decreasing (right) the gain will both decrease or increase the firing rate variance and adapt the transfer function (blue arrows) so that it maps most potential values linearly.

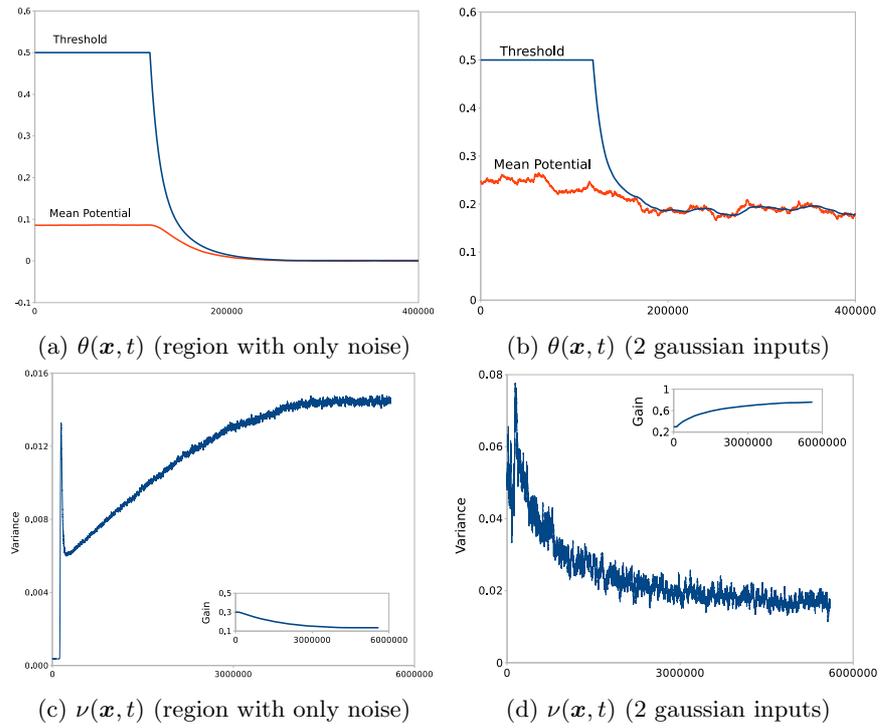


Fig. 4: Transfer function adaptation. The gain and the threshold both converge and lead to a stable system in the case of a strong input (b,d). The mean potential converge to 0 for a low input (a,c). Variance plots are obtained by averaging a set of 10 experiments.

6 Combination of dynamic adaptation methods

In this section, it will be described how the self-adaptation of the potentials and the firing rates is performed simultaneously. In order to avoid interferences, attention must be given to the time scales of the adaptation mechanisms ρ and λ , and to the adaptation constants ϵ_α , ϵ_ν and ϵ_θ .

We require that all statistical quantities should operate on a similar time scale ($\lambda \approx \rho$), and that adaptation mechanisms should be significantly slower. Threshold adaptation is coupled to input strength adaptation, then $\epsilon_\alpha \approx \epsilon_\theta$. Gain adaptation depends critically on centering of the threshold, therefore $\epsilon_\nu < \epsilon_\theta$:

$$1/\tau > 1/800 > \lambda \approx \rho > \epsilon_\alpha \approx \epsilon_\theta > \epsilon_\nu \quad (8)$$

Using target values of $\bar{u}_{target}=0.1$, $\Sigma_{target}=0.015$, time constants $\lambda=0.01$ and $\rho=0.01$, and adaptation coefficients $\epsilon_\alpha=0.0005$, $\epsilon_\theta=0.0001$, and $\epsilon_\nu=0.00001$, one can see in Fig. 5 that the joint adaptation is indeed successful. We do not observe the previous problem with weak input strength anymore.

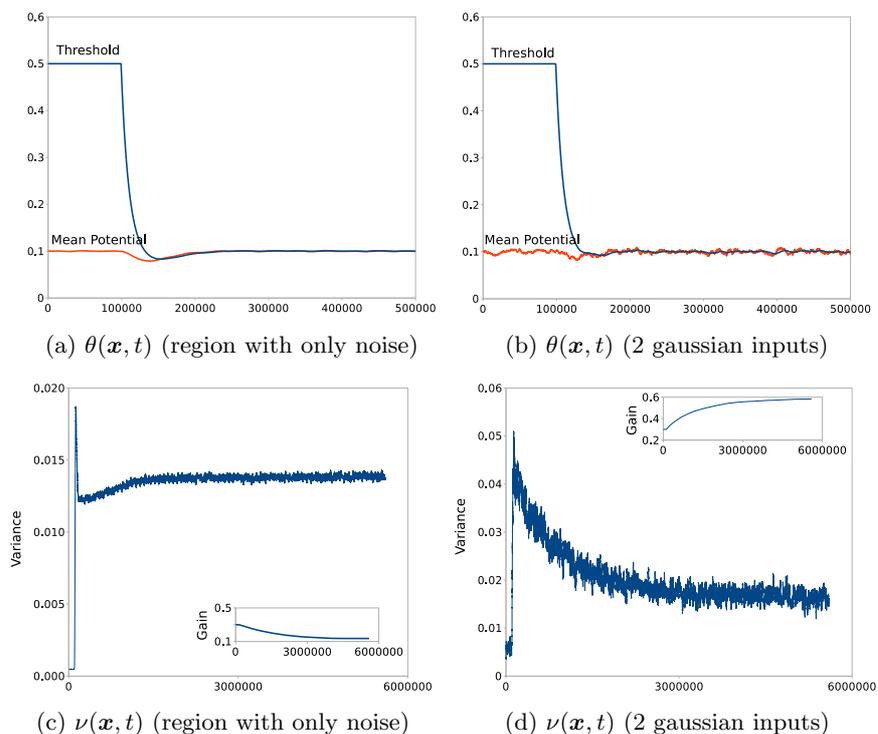


Fig. 5: Result with both mean potential adaptation and transfer function adaptation. Variance plots are obtained by averaging a set of 10 experiments.

7 Discussion

This contribution introduced several adaptation rules for actively maintaining neurons in a desired dynamic state by suitably modifying their internal processing parameters. An important outcome of our experiments is the fact that such adaptation mechanisms must work on appropriate time scales w.r.t. each other in order to avoid instabilities and divergences. The benefit of the work presented here can be summarized as follows: First of all, the number of free parameters that have to be set by a system designer is reduced. Secondly, neural fields can now be used in a "plug and play" manner, connecting them without having to explicitly consider parameter settings for dynamical stability. This will facilitate the construction of large neural systems as envisioned in [23,7]

However, there still remains a significant number of parameters that are not subject to adaptation. Many of these parameters may be eliminated by further adaptation mechanisms, which will be part of our future work. However, there must remain a set of parameters describing some high-level function of a neuron (e.g., accumulation, decision making, integration, discrimination) that are determined by a designer. Therefore it should be clear that self-adaptation can eliminate several but not all free parameters in a system.

8 Conclusion

As was mentioned before, there exists a number of parameters that is not considered for adaptation in this contribution. In the future, it should be investigated how to extend self-adaptation mechanisms to these parameters (which are treated as constants right now), most notably the lateral interaction strength β and the time constant τ .

References

1. EM Izhikevich. Which model to use for cortical spiking neurons? *IEEE Transactions on Neural Networks*, 15(5), 2004.
2. GG Turrigiano and SB Nelson. Homeostatic plasticity in the developing nervous system. *Nat Rev Neurosci*, 5(2):97–107, Feb 2004.
3. G Turrigiano. The self-tuning neuron: Synaptic scaling of excitatory synapses. *Cell*, 135(3), 2008.
4. A Gepperth, J Fritsch, and C Goerick. Cross-module learning as a first step towards a cognitive system concept. In *Proceedings of the First International Conference On Cognitive Systems*, 2008.
5. R Fay, U Kaufmann, A Knoblauch, H Markert, and G Palm. Combining Visual Attention, Object Recognition and Associative Information Processing in a NeuroBotic System. In S Wermter, G Palm, and M Elshaw, editors, *Biomimetic Neural Learning for Intelligent Robots. Intelligent Systems, Cognitive Robotics, and Neuroscience.*, volume 3575 of *Lecture Notes in Computer Science LNAI*, pages 118–143. Springer Berlin Heidelberg, 2005.

6. T Wennekers, M Garagnani, and F Pulvermüller. Language models based on hebbian cell assemblies. *J Physiol Paris*, 100(1-3):16–30, 2006.
7. D Vernon, G Metta, and G Sandini. The iCub cognitive architecture: Interactive development in a humanoid robot. In *6th International Conference on Development and Learning*, 2007.
8. S-I Amari. Mathematical foundations of neurocomputing. *Proceedings of the IEEE*, 78(9):1441–1463, 1990.
9. Wolfram Erlhagen and Gregor Schöner. Dynamic field theory of movement preparation. *Psychol Rev*, 109(3):545–572, Jul 2002.
10. P Cisek. Integrated neural processes for defining potential actions and deciding between them: a computational model. *J Neurosci*, 26(38):9761–9770, Sep 2006.
11. W Zhang and DJ Linden. The other side of the engram: experience-driven changes in neuronal intrinsic excitability. *Nat Rev Neurosci*, 4, 2003.
12. M Stemmler and C Koch. How voltage-dependent conductances can adapt to maximize the information encoded by neuronal firing rate. *Nat Neurosci*, 2, 1999.
13. J Triesch. Synergies between intrinsic and synaptic plasticity mechanisms. *Neural Comput.*, 19(4):885–909, 2007.
14. T Elliott, X Kuang, NR Shadbolt, and KP Zauner. An invariance principle for maintaining the operating point of a neuron. *Network*, 19(3):213–35, 2008.
15. JJ Steil. Online reservoir adaptation by intrinsic plasticity for backpropagation-decorrelation and echo state learning. *Neural networks*, 20(3):353–364, April 2007.
16. N Butko and J Triesch. Learning sensory representations with intrinsic plasticity. In *Proceedings of the European Symposium on Artificial Neural Networks*, 2007.
17. C Gläser, F Joublin, and C Goerick. Homeostatic development of dynamic neural fields. In *International Conference on Development and Learning*, 2008.
18. JS Johnson, JP Spencer, and G Schöner. Moving to higher ground: The dynamic field theory and the dynamics of visual cognition. *New Ideas Psychol*, 26(2):227–251, Aug 2008.
19. C Wilmzig, S Schneider, and G Schöner. The time course of saccadic decision making: dynamic field theory. *Neural Netw*, 19(8):1059–1074, Oct 2006.
20. NP Rougier and J Vitay. Emergence of attention within a neural population. *Neural Netw*, 19(5):573–581, Jun 2006.
21. C Faubel and G Schöner. Fast learning to recognize objects: Dynamic fields in label-feature spaces. In *Proceedings of the 5th International Conference on Development and Learning*, 2006.
22. G Deco and ET Rolls. A neurodynamical cortical model of visual attention and invariant object recognition. *Vision Res*, 44(6):621–642, Mar 2004.
23. A Gepperth, J Fritsch, and C Goerick. Computationally efficient neural field dynamics. In *Proceedings of the 16th European Symposium on Artificial Neural Networks*, pages 179–185, 2008.