

Homeostatic Development of Dynamic Neural Fields

Claudius Gläser, Frank Joublin, Christian Goerick
Honda Research Institute Europe

Carl-Legien-Str. 30, D-63073 Offenbach/Main, Germany
{*claudius.glaeser, frank.joublin, christian.goerick*}@honda-ri.de

Abstract—Dynamic neural field theory has become a popular technique for modeling the spatio-temporal evolution of activity within the cortex. When using neural fields the right balance between excitation and inhibition within the field is crucial for a stable operation. Finding this balance is a severe problem, particularly in face of experience-driven changes of synaptic strengths. Homeostatic plasticity, where the objective function for each unit is to reach some target firing rate, seems to counteract this problem. Here we present a recurrent neural network model composed of excitatory and inhibitory units which can self-organize via a learning regime incorporating Hebbian plasticity, homeostatic synaptic scaling, and self-regulatory changes in the intrinsic excitability of neurons. Furthermore, we do not define a neural field topology by a fixed lateral connectivity; rather we learn lateral connections as well.

Index Terms—Dynamic Neural Field, Development, Homeostasis

I. INTRODUCTION

The beginnings of dynamical neural field theory (DNFT) reach back to the 1950's [1]. Since that and particularly after seminal contributions have been done by Wilson and Cowan [2] as well as Amari [3], DNFT has become a popular technique for modeling the spatio-temporal evolution of activity in the brain. Due to its variety in exhibited dynamic behavior ranging from periodic activity patterns over activity bumps up to traveling waves [4], neural fields have been used to model and understand cortical computation in various domains.

For robotic applications localized regions of excitation (activity bubbles) following the presentation of a stimulus are of particular interest. However, those solutions are difficult to achieve, since the dynamic behavior of neural fields is very sensitive to the right balance between excitation and inhibition within the field. Small changes to this balance will result in runaway excitation or quiescence. Furthermore, the influence of different parameter settings is only roughly estimated [3], [5], [6], such that parameters are often chosen heuristically or optimized via evolutionary strategies.

In contrast, the development of dynamic neural fields is only rarely explored even though taking a developmental approach might ultimately overcome the above mentioned problems. Learning most often focuses on the synaptic weights of projections from the input space to the neural field, thereby adapting the input-driven dynamics, but leaving the self-driven dynamics unchanged (see [7] for an exception of this). For this reason, extensive normalization strategies, often using network level knowledge, have to be applied in order to keep the neural field in a stable regime.

Our approach differs insofar as we do not make any assumption on the connectivity of the field. In other words,

synaptic weights of both, afferent projections to the field as well as lateral connections within the field, are learned. As a direct consequence, neural fields have to self-regulate in order to maintain a stable operation mode even in face of these experience-driven changes. Our model incorporates recent advances in the understanding of homeostatic processes regulating neuronal activity, namely homeostatic synaptic scaling and changes in the intrinsic excitability of neurons. We will experimentally show how homeostasis in form of locally operating processes contributes to the global stability of the neural field. Due to the self-regulatory nature of our model, the number of free parameters reduces to a minimum which eases its use for applications in various domains. It is particularly suited for modeling cortical development, since the process of learning the mapping is self-organizing, intrinsically regulated, and only depends on the statistics of the input patterns.

The rest of the paper is organized as follows. In Section II we will present the structure of our network model and discuss the differences to existing approaches. Section III highlights recent advances in our understanding of homeostatic processes regulating neuronal activity and shows how these mechanisms can be incorporated into a learning regime for the self-regulatory development of dynamic neural fields. In Section IV we experimentally evaluate our model in the domain of reference frame transformation. Finally, we will give a conclusion.

II. DYNAMIC NEURAL FIELDS

Dynamic neural fields are composed of model units distributed on a plane mimicking the neural tissue. Populations of neurons interact through extensive lateral connections which results in dynamic patterns of activity following stimulus presentation. Amari [3] formulated a field equation by which most of the present models can be described:

$$\tau \frac{\partial u(x, t)}{\partial t} = -u(x, t) + \int w(x, x') \cdot f(u(x', t)) dx' + S(x, t) + h \quad (1)$$

Here $u(x, t)$ represents the local activity of a population of neurons at position x of the cortical plane at time t . The stimuli to the neural field at position x and time t is denoted by $S(x, t)$, h is some resting potential which is approached in the absence of other inputs, and f is a monotonically increasing transfer function describing the relation between the activation and the firing rate of neurons. The lateral connectivity from neurons located at position x' to neurons located at position x of the neural tissue is defined by $w(x, x')$. For most present




Fig. 1. The structure of the recurrent neural network.

network models $w(x, x') = w(x - x')$ holds, that is the synaptic strength of lateral connections linking neurons of the neural field only depends on their distance on the cortical plane. There, a Mexican Hat connectivity realizing local excitation and distal inhibition is typically chosen. A distance-dependent lateral connectivity, however, may not be suitable, i.e. for high-dimensional data which cannot be adequately mapped on a 2-dimensional surface. For this reason, afferent projections to the field as well as connections within the field should underly experience-driven changes of synaptic strength.

The structure of the recurrent neural network we present here is based on the work of Wilson and Cowan [2]. It is composed of excitatory units E and inhibitory units I, both being arranged on a 2-dimensional grid mimicking the neural tissue. The wiring of the network is shown in Fig. 1. Afferent projections to excitatory units provide the input to the neural field. Furthermore, the units are laterally connected such that excitatory cells (E-cells) excite other E-cells as well as inhibitory cells (I-cells). In turn, E-cells receive inhibitory projections originating from I-cells.

The membrane potentials of excitatory and inhibitory units are denoted by the variables u and v , respectively. In the following we will use i for specifying the unit located at position x_i of the cortical plane. Therewith the spatio-temporal evolution of the activity in the neural field can be described by the following differential equations:

$$\begin{aligned} \tau_E \frac{du_i}{dt} = & -u_i + \sum_j g(d_{ij}) \cdot w_{ij}^{EE} \cdot f(u_j) \\ & - \sum_j w_{ij}^{EI} \cdot f(v_j) \\ & + \sum_j w_{ij}^{EXT} \cdot s_j + h_i^E \end{aligned} \quad (2)$$

$$\tau_I \frac{dv_i}{dt} = -v_i + \sum_j g(d_{ij}) \cdot w_{ij}^{IE} \cdot f(u_j) + h^I \quad (3)$$

Here, τ_E and τ_I are time constants, h_i^E and h^I are the resting potentials, and w_{ij}^* denotes the synaptic weight of a connection from unit j to unit i where $* \in \{EE, EI, IE, EXT\}$ specifies the type of connection. The monotonically increasing transfer function defining the relation between the membrane potential and the firing rate of a unit is denoted by f . We used a sigmoidal function of the form:

$$f(z) = \frac{1}{1 + \exp(-\gamma(z - \theta))} \quad (4)$$

with θ and γ denoting the threshold value and the gain factor, respectively. Additionally, we introduced a function g

which modulates the efficacy of excitatory lateral connections depending on the distance $d_{ij} = x_i - x_j$ between the pre- and postsynaptic unit positions. It was chosen to follow the normal distribution with a mean of 0 and a standard deviation of σ :

$$g(d) = \frac{1}{\sigma\sqrt{2\pi}} \cdot \exp\left(-\frac{d^2}{2\sigma^2}\right) \quad (5)$$

Thus, we define that excitatory lateral connections between units within a local neighborhood are more efficient than those between far-distant units. It is, however, important to note that this is fundamentally different compared to existing network models. There the synaptic weight values of lateral connections were chosen as a function of the distance between the pre- and postsynaptic units by which a topology on the neural tissue as well as within the feature space is defined. In contrast, we introduce a distance-dependent modulation of synaptic efficiency but do not make any assumption on the synaptic weight values themselves. Following this argumentation, g could be understood as a physical constraint. The integration of excitatory synaptic input as it is carried out by passive dendrites is one possible interpretation. There, it has been shown that the efficacy of a synaptic input significantly varies depending on synapse location [8].

A direct consequence of this is that large synaptic weight values could compensate for the distance-dependent modulation of connection efficiency. Thus, the mapping described by the neural field does not necessarily have to be topology preserving, that is nearby units having similar receptive fields. Direct support for this idea comes from a study recently carried out by van Hooser et al. [9]. They found orientation-sensitive cells in the primary visual cortex (V1) of a highly visual rodent, the gray squirrel, similar to those found in V1 of humans. But in contrast, orientation-selectivity did not smoothly vary across the cortical surface. Thus, topology preserving self-organization does not seem to be a fundamental principle of mammalian cortical development, rather it seems to depend on other mechanisms missing in rodents.

Up to now, we described the structure of the recurrent neural network, but in order to model neural field development, a learning regime has to be defined. This involves both, a mechanism for adapting the neural field such that it adequately represents the input pattern distribution as well as self-regulatory processes keeping the neural field in a stable state. The former can be achieved via Hebbian forms of synaptic plasticity, a learning principle usually summarized as *cells that fire together, wire together*. This means, if the postsynaptic cell repeatedly fires following a stimulation by the presynaptic cell, the synapse linking both cells is strengthened. In its simplest form Hebbian learning would lead to unconstrained weight growths. For this reason, saturation or normalization techniques have to be applied. A famous variant of Hebbian learning is Oja's rule [10] described by (6), where w_{ij} is the synaptic weight as well as η_i and ξ_j the pre- and postsynaptic activities, respectively.

$$\Delta w_{ij} \propto \eta_i \cdot \xi_j - w_{ij} \cdot \eta_i^2 \quad (6)$$

Here the first term represents the basic Hebbian learning rule, whereas the second term regulates the synaptic weight change

via an activity-dependent leakage. More importantly, it has been shown that a neuron using this learning rule tends to extract the statistically most significant factor from its inputs, the principle component. Therewith it is a suitable learning technique for adapting the synaptic weights of a neural field.

As mentioned earlier, a balanced excitatory and inhibitory connectivity is crucial in order to keep the dynamic neural field in a proper working state. Finding this right balance is a severe problem, which becomes even more significant when synaptic weights change through some learning mechanism. In other words, self-regulatory processes counteracting these changes have to be applied [11]. In the following we will highlight recent advances in the understanding of homeostatic processes regulating neuronal activity and show how they can be incorporated into a learning regime for the development of continuous neural fields.

III. HOMEOSTATIC PLASTICITY

Homeostatic mechanisms can be described as processes intrinsic to a system, which regulate the internal environment in face of a changing external environment in order to keep the system in a stable state. In the case of the nervous system numerous homeostatic processes acting at different parts of the neural circuit can be distinguished [12]. Here we will concentrate on those operating at the synaptic level, so called homeostatic synaptic plasticity [13], and those causing neuron-wide changes in intrinsic excitability [14].

Changes induced by homeostasis occur at a relatively slow time scale, possibly hours or days, which is important in order not to destroy moment-to-moment fluctuations in activity carrying the information to be learned. Conversely, they have to be fast enough in order to compensate for changes induced by (Hebbian) learning [13]. The need for a stably operating neural circuit is directly correlated to the need of regulating overall activity. Compelling evidence point to the fact that homeostatic processes aim at regulating the average activity of individual units towards some target firing rate, thereby implicitly regulating overall activity. For example it has been shown that blocking of excitatory or inhibitory input to a cell results in initially depressed or increased cell activity, respectively, but over time cell activity approaches control level again [15].

The need of regulating an individual neuron's average activity necessitates the ability of a neuron to monitor its own activity, a property typically attributed to changes in intracellular calcium concentrations [16]. Computationally, a mean firing rate \bar{A}_i of a neuron i can be easily computed using the following equation:

$$\bar{A}_i(k) = \left(1 - \frac{1}{\tau_H}\right) \cdot \bar{A}_i(k-1) + \frac{1}{\tau_H} \cdot A_i(k) \quad (7)$$

where $A_i(k) = f(u_i(k))$ is the instantaneous firing rate of the unit at discrete time k and τ_H a time constant defining the time scale on which integration takes place.

In order to reach some target firing rate, a neuron can adjust the synaptic weights of its afferent projections, an ability commonly referred to as synaptic scaling. Different forms of modifying synaptic strength are possible, namely

additive and multiplicative synaptic scaling. The different effects of both techniques have been highlighted in [17]. For the central nervous system it has been shown that neurons adjust their synaptic weights multiplicatively [18], which has the computationally attractive feature of leaving the relative difference in synaptic weights unchanged.

The previously mentioned Oja's rule (see (6)) as a special variant of Hebbian learning performs multiplicative scaling as well. Why should we then need additional mechanisms for the regulation of synaptic strengths? The key difference is the objective function controlling synaptic scaling. If a neuron aims at maintaining some target firing rate, it has to adjust the balance between excitatory and inhibitory currents as well, which is a property not targeted by Oja's rule. Even small changes in the excitatory-inhibitory balance can disrupt network performance, meaning that dynamic adjustments in the relative strength of excitatory and inhibitory feedback on excitatory neurons is an important component of firing rate homeostasis [13].

Neuroscientific studies indicate that this dynamic adjustment is mediated by the activity-dependent release of the neurotrophin BDNF (brain-derived neurotrophic factor) [19]. Interestingly, this means that the activity of inhibitory interneurons is regulated by adjusting the synaptic strengths of their afferent excitatory projections depending on the activity of the presynaptic pyramidal cell (releasing the BDNF). Furthermore, it has been shown that BDNF has opposing effects on the scaling of excitatory synapses on pyramidal neurons and interneurons, thereby changing the ratio between the firing rates of excitatory and inhibitory neurons [20]. A high BDNF level reduces the synaptic strengths on excitatory neurons, whereas it increases synaptic strengths on inhibitory neurons and vice versa.

We modeled the BDNF release of an excitatory unit i (E-cell) given its mean firing rate \bar{A}_i^E (see (7)) and a target firing rate \hat{A} as follows:

$$BDNF_i^E(k) = 1 + \beta_H \left(\frac{\bar{A}_i^E(k-1) - \hat{A}}{\hat{A}} \right) \quad (8)$$

where β_H is a homeostatic learning rate. According to (8) the BDNF release is greater than 1, if the mean firing rate exceeds its target level, whereas it is smaller than 1, if the mean firing rate lies below its target level. A similar normalization factor was recently proposed for the synaptic scaling within self-organizing maps [21]. However, the application of the normalization factor differs significantly. We modeled the evolution of the synaptic weights within our recurrent neural network by a combination of Oja's rule described by (6) and the BDNF normalization factor of (8). Additionally, different types of synapses are differently scaled according to the opposing effects of BDNF [20]. The resulting learning technique is described by the following formulas:

$$w_{ij}^{EXT}(k) = \frac{w_{ij}^{EXT}(k-1) + \alpha \cdot \Delta \tilde{w}_{ij}^{EXT}(k)}{BDNF_i^E(k) \cdot BDNF_j^{EXT}(k)} \quad (9)$$

$$w_{ij}^{EE}(k) = \frac{w_{ij}^{EE}(k-1) + \alpha \cdot \Delta \tilde{w}_{ij}^{EE}(k)}{BDNF_i^E(k) \cdot BDNF_j^E(k)} \quad (10)$$

$$w_{ij}^{EI}(k) = [w_{ij}^{EI}(k-1) + \alpha \cdot \Delta \tilde{w}_{ij}^{EI}(k)] \cdot BDNF_i^E(k) \quad (11)$$

$$w_{ij}^{IE}(k) = [w_{ij}^{IE}(k-1) + \alpha \cdot \Delta \tilde{w}_{ij}^{IE}(k)] \cdot BDNF_j^E(k) \quad (12)$$

where $\Delta \tilde{w}_{ij}^*(k)$ with $* \in \{EE, EI, IE, EXT\}$ denotes the weight change according to Oja's rule and α a learning rate.

So far we focused on homeostatic mechanisms operating at the synapses directly. Another category of self-regulatory processes dynamically adjust neuronal properties like the transfer function, thereby changing a neuron's intrinsic excitability. But what is the benefit of multiple simultaneously operating homeostatic processes and do we need to incorporate an additional mechanism at all? For sure, both above mentioned processes try to locally adapt neural circuits and even their objective functions seem to be the same (individual neurons try to reach a target firing rate). However, the way by which self-regulation is achieved differs significantly. Homeostatic synaptic scaling dynamically adjust the balance between excitation and inhibition, which is due to the opposing effects of BDNF. An activity-level dependent release of BDNF of course means that synaptic scaling operates with respect to the transfer function of neurons. Nevertheless, it does not ensure that the settled balance between excitation and inhibition results in membrane potentials which lie within the dynamic range a neuron is most sensible to. This drawback can be compensated by adjusting the transfer functions of neurons. For this reason, the incorporation of homeostatic processes changing synaptic strengths as well as the excitability of neurons may be necessary in order to reach optimal solutions.

Given a sigmoidal transfer function f according to (4), a neuron's intrinsic excitability can be changed by dynamically adjusting the gain and threshold parameter γ and θ , respectively. In a recent work, Triesch [22] derived an update mechanism for the parameters based on information theory.

Here, we restrict our work to dynamically adjusting the transfer function's threshold value of excitatory neurons. Since changing the resting potential h_i^E of an excitatory neuron i is equivalent to threshold adaption, we did not vary the threshold parameter θ , rather we adapted the resting potentials h_i^E . Due to the fact that changes in the intrinsic excitability of a neuron seems to be induced by changes in intracellular calcium concentrations, which is a correlate of a neuron's activity level, we follow the same argumentation as for homeostatic synaptic scaling. A neuron adjusts its excitability given its average activity level \bar{A}_i^E and some target firing rate \hat{A} . If a neuron's activity level exceeds its target level, the neuron's excitability has to be reduced which can be achieved by decreasing its resting potential. Conversely, an increase of a neuron's resting potential is suitable, if its excitability should be larger as it is the case when the neuron's activity level does not reach its target level. The following formula describes the evolution of the resting potentials we used within our network model:

$$h_i^E(k) = h_i^E(k-1) + \beta_T \cdot \left(\frac{\hat{A} - \bar{A}_i^E(k-1)}{\hat{A}} \right) \quad (13)$$

There β_T denotes a learning rate representing the time scale at which homeostatis via threshold adaption takes place.




Fig. 2. The receptive fields of the neurons used for coding the gaze position in body-centered coordinates (s^1), the hand position in body-centered coordinates (s^2), and the hand-position in an eye-centered reference frame (s^3).

IV. EXPERIMENTAL RESULTS

In order to evaluate the presented recurrent neural network model, we carried out an experiment in the domain of reference frame transformation. More precisely, we investigated the use of dynamic neural fields for 1-dimensional eye-hand coordination. In order to robustly perform eye-hand coordination, an animal has to be able to transform between the different reference frames [23]. This ability is usually attributed to an intermodal body-calibration obtained in the early stages of development [24]. There, the key aspect is that simultaneously present stimuli become linked together and can later be used for the transformation from one modality into another [25]. Here, we want to use our network for modeling the calibration process during early self-exploration. Therefore, we have chosen three stimuli s^1, s^2, s^3 with $s^1, s^2 \in [-1, 1]$ and $s^3 = s^1 - s^2$, where s^1 and s^2 mimic the gaze and hand position in a body-centered reference frame, respectively, as well as s^3 representing the hand position in eye-centered coordinates.

The experimental setup is as follows. Each of the three stimuli s^1, s^2, s^3 is represented by a population code composed of 21 neuron responses, resulting in a total of 63 inputs to the neural field. The receptive fields of the neurons coding the stimuli are shown in Fig. 2. Target gaze- and hand-positions (s^1, s^2) were chosen randomly, but in order to obtain smooth transitions between target positions, head and hand movements exhibit simple linear dynamics. The recurrent neural network is composed of 100 excitatory units and 100 inhibitory units, both being arranged on a 10x10 grid. The synaptic weight values of afferent projections to the field w_{ij}^{EXT} were initialized with small random values. Weights of lateral connections were initialized uniformly. In contrast to the LISSOM model [7], which also features plastic lateral connections, we do not make any assumption on when learning takes place. More precisely, in [7] the synaptic weights were changed only after the network settled into a stable state. In our network model learning is performed at each timestep.

Over the course of development, the recurrent neural network self-organizes via balanced cooperation and competition. In an initial phase, the neural field grossly adapts to the input pattern distribution. After that, the receptive fields of the neurons become more and more distinct. Fig. 3 shows the developed receptive fields of the excitatory units in the s^1 - s^2 -plane after several stimuli have been presented. As can be




Fig. 3. The developed receptive fields of all excitatory units (arranged on a 2-dimensional grid) are shown in the s^1 - s^2 -plane.




Fig. 4. The position of the excitatory neurons in feature space (s^1 - s^2 - s^3 -plane) as obtained by calculating the receptive fields' center of masses.

seen, each neuron specializes to a particular combination of s^1 and s^2 . Furthermore, we calculated the center of masses of the receptive fields, thereby obtaining the neuron positions in the feature space. As illustrated in Fig. 4, the neurons are nicely distributed on the s^1 - s^2 - s^3 -plane which means that each input pattern is adequately represented by the neural field.

A quantitative measure of the functional quality of the mapping can be obtained by comparing the neural field responses following the presentation of different stimuli. Small changes in the input should result in small changes in the output, whereas large changes in the output should be obtained following the presentation of largely different inputs. Therefore, we treated inputs $s \in \mathbb{R}^3$ and outputs $A \in \mathbb{R}^{100}$ as vectors and calculated their normalized changes (Δs and ΔA) for all possible pairs of stimuli. In Fig. 5 (a) the means and the variances of the differences in output activity for all possible changes in input activity are plotted. The curve nicely follows the above argumentation, in that the difference between normalized input vectors of pair-wise presented stimuli is reflected in the difference between the resulting normalized output vectors. In other words, the more similar (different) inputs representing different stimuli are, the more similar (different) are the corresponding output vectors of the neural field.

Another interesting aspect is, whether the learned mapping is topology preserving, that is nearby units having similar receptive fields. As already discussed in Section II, the mapping learned by our recurrent neural network model does not necessarily have to be topology preserving, since we do not




Fig. 5. The means and variances of (a) the normalized changes in output activity and (b) the distance between the locations of peak activity given the normalized changes in input activity for pair-wise presented stimuli, respectively.

make any assumption on the synaptic weights of the lateral connections within the field. However, the development of topology preserving mappings does not seem to be a fundamental principle for cortical development. It might even be not necessarily possible since high-dimensional inputs cannot be easily mapped onto the 2-dimensional cortical surface. Nevertheless, in order to give a quantitative measure for the learned mapping, we calculated the euclidean distance between neurons on the grid (Δx), which exhibit the peak responses following the presentation of different stimuli. A pair-wise comparison results in the curve depicted in Fig. 5 (b). As it is shown, the larger the distances between two normalized input vectors are, the larger the mean distances between the locations of the corresponding peak responses become. However, the plot also indicates large variances in the measurements which demonstrates that the learned mapping is not strictly topology preserving. A typical example for this is shown in Fig. 6. In (a) the neural field activity following the presentation of a stimulus ($s^1 = 0.4, s^2 = -0.3, s^3 = 0.7$) is plotted according to the position of the neurons on the cortical plane, which results in multiple peaks of activity. In contrast, (b) shows the neural field activity following the presentation of the same stimulus according to the neurons' positions in the feature space. Here, a single activity bubble arises.

Lastly, we investigated how locally operating homeostatic mechanisms with a neuron's objective function of reaching some average firing rate affect the overall activity of the neural field. Therefore, we calculated a running average of the overall field activity for different target firing rates \hat{A} of individual neurons. As Fig. 7 shows, the overall field activity quickly rises towards a mean activity level proportional to the target firing rates of individual neurons. After that, the neural field activity oscillates around the mean activity level where the period of oscillations is dependent on the time scale at which homeostatic processes operate.

V. CONCLUSION

We presented a recurrent neural network model for the self-regulatory development of dynamic neural fields. In contrast to previously published approaches, the synaptic weights of all connections within our network model are plastic. Particularly, we do not fix lateral connections, rather they undergo experience-driven changes via Hebbian forms of plasticity as well. As a direct consequence, a mapping described by the




Fig. 6. The neural field activity following the presentation of a stimulus ($s^1 = 0.4, s^2 = -0.3, s^3 = 0.7$) is plotted according to (a) the position of the neurons on the cortical plane and (b) the position of the neurons in feature space.




Fig. 7. Running averages of the overall field activity for different target firing rates of individual neurons.

neural field does not necessarily have to be topology preserving. We propose, that in order to develop topology preserving mappings, another mechanism has to be incorporated.

In order to keep the neural field in a stable state even in face of experience-driven changes of synaptic strengths, we applied biologically inspired forms of homeostatic plasticity, namely synaptic scaling and changes in a neuron’s intrinsic excitability. These locally operating self-regulatory mechanisms aim at changing the activity levels of individual neurons towards some target firing rate. Depending on its average activity level, a neuron is able to adjust its synaptic weights (thereby changing the ratio between excitation and inhibition within the field) and its intrinsic excitability via threshold adaptation.

By incorporating homeostatic processes, the number of free parameters that have to be controlled reduces to the time constants and the target firing rate of individual neurons. Here, particularly the latter is very interesting, since the target firing rate of individual units defines the overall activity in the field, the overlap between the receptive fields of single units and therewith the sparsity of the developed representation. However, here we did not experimentally evaluate the influence of different parameter settings.

Due to its low number in controllable parameters, the present work will ease the use of neural fields in applications of various domains. Our network is particularly interesting when taking a developmental perspective, because the specification of the input patterns (i.e. stimuli of different modalities that have to be correlated) is sufficient for obtaining the desired functional behavior, whereas the process of establishing the mapping is self-organizing and intrinsically regulated.

We experimentally tested our network model in the domain of reference frame transformation and showed that it self-organizes without the use of any external supervision. It adjusts the receptive fields of its units such that the input pattern distribution is adequately represented. Furthermore,

we illustrated that locally operating homeostatic processes are sufficient for keeping the overall network in a stable state.

The incorporation of a mechanism for enhancing topology preservation during neural field development as well as the use of our network model within architectures composed of several layers will be part of our future work.

REFERENCES

- [1] R. L. Beurle, “Properties of a mass of cells capable of regenerating pulses,” *Phil. Trans. of the Roy. Soc. of London*, vol. 240, no. 669, pp. 55–94, 1956.
- [2] H. R. Wilson and J. D. Cowan, “A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue,” *Kybernetik*, vol. 13, no. 2, pp. 55–80, Sep 1973.
- [3] S. Amari, “Dynamics of pattern formation in lateral-inhibition type neural fields,” *Biol Cybern*, vol. 27, pp. 77–87, 1977.
- [4] S. Coombes, “Waves, bumps, and patterns in neural field theories,” *Biol Cybern*, vol. 93, no. 2, pp. 91–108, 2005.
- [5] J. G. Taylor, “Neural ‘bubble’ dynamics in two dimensions: foundations,” *Biol Cybern*, vol. 80, pp. 393–409, 1999.
- [6] I. Mikhailova and C. Goerick, “Conditions of activity bubble uniqueness in dynamic neural fields,” *Biol Cybern*, vol. 92, no. 2, pp. 82–91, 2005.
- [7] J. Sirosh and R. Miikkulainen, “Cooperative self-organization of afferent and lateral connections in cortical maps,” *Biol Cybern*, vol. 71, pp. 65–78, 1994.
- [8] J. C. Magee, “Dendritic integration of excitatory synaptic input,” *Nat Rev Neurosci*, vol. 1, no. 3, pp. 181–190, Dec 2000.
- [9] S. D. V. Hooser, J. A. F. Heimel, S. Chung, S. B. Nelson, and L. J. Toth, “Orientation selectivity without orientation maps in visual cortex of a highly visual mammal,” *J Neurosci*, vol. 25, pp. 19–28, 2005.
- [10] E. Oja, “Simplified neuron model as a principal component analyzer,” *J of Math Biol*, vol. 15, no. 3, pp. 267–273, 1982.
- [11] G. G. Turrigiano and S. B. Nelson, “Hebb and homeostasis in neuronal plasticity,” *Curr Opin Neurobiol*, vol. 10, no. 3, pp. 358–364, 2000.
- [12] E. Marder and J.-M. Goaillard, “Variability, compensation and homeostasis in neuron and network function,” *Nat Rev Neurosci*, vol. 7, no. 7, pp. 563–574, 2006.
- [13] G. G. Turrigiano and S. B. Nelson, “Homeostatic plasticity in the developing nervous system,” *Nat Rev Neurosci*, vol. 5, no. 2, pp. 97–107, 2004.
- [14] W. Zhang and D. J. Linden, “The other side of the engram: experience-driven changes in neuronal intrinsic excitability,” *Nat Rev Neurosci*, vol. 4, no. 11, pp. 885–900, 2003.
- [15] E. Marder and A. A. Prinz, “Modeling stability in neuron and network function: the role of activity in homeostasis,” *Bioessays*, vol. 24, no. 12, pp. 1145–1154, 2002.
- [16] M. J. Berridge, “Neuronal calcium signaling,” *Neuron*, vol. 21, no. 1, pp. 13–26, 1998.
- [17] G. J. Goodhill and H. G. Barrow, “The role of weight normalization in competitive learning,” *Neural Comput*, vol. 6, no. 2, pp. 255–269, 1994.
- [18] G. G. Turrigiano, K. R. Leslie, N. S. Desai, L. C. Rutherford, and S. B. Nelson, “Activity-dependent scaling of quantal amplitude in neocortical neurons,” *Nature*, vol. 391, no. 6670, pp. 892–896, 1998.
- [19] L. C. Rutherford, A. DeWan, H. M. Lauer, and G. G. Turrigiano, “Brain-derived neurotrophic factor mediates the activity-dependent regulation of inhibition in neocortical cultures,” *J Neurosci*, vol. 17, no. 12, pp. 4527–4535, 1997.
- [20] L. C. Rutherford, S. B. Nelson, and G. G. Turrigiano, “Bdnf has opposite effects on the quantal amplitude of pyramidal neuron and interneuron excitatory synapses,” *Neuron*, vol. 21, no. 3, pp. 521–530, 1998.
- [21] T. J. Sullivan and V. R. de Sa, “Homeostatic synaptic scaling in self-organizing maps,” *Neural Netw*, vol. 19, no. 6-7, pp. 734–743, 2006.
- [22] J. Triesch, “Synergies between intrinsic and synaptic plasticity in individual model neurons,” *Neural Comput*, vol. 19, no. 4, pp. 885–909, 2007.
- [23] Y. E. Cohen and R. A. Andersen, “A common reference frame for movement plans in the posterior parietal cortex,” *Nat Rev Neurosci*, vol. 3, no. 7, pp. 553–562, Jul 2002.
- [24] R. Morgan and P. Rochat, “Intermodal calibration of the body in early infancy,” *Ecol Psych*, vol. 9, no. 1, pp. 1–23, 1997.
- [25] L. E. Bahrick and J. S. Watson, “Detection of intermodal proprioceptive-visual contingency as a potential basis of self-perception in infancy,” *Dev Psych*, vol. 21, no. 6, pp. 963–973, 1985.