

A model for the interaction of oscillations and pattern generation with real-time computing in generic neural microcircuit models[☆]

Alexander Kaske, Wolfgang Maass*

Institute for Theoretical Computer Science, Technische Universitaet Graz, A-8010 Graz, Austria

Abstract

It is shown that real-time computations on spike patterns and temporal integration of information in neural microcircuit models are compatible with potentially descriptive additional inputs such as oscillations. A minor change in the connection statistics of such circuits (making synaptic connections to more distal target neurons more likely for excitatory than for inhibitory neurons) endows such generic neural microcircuit model with the ability to generate periodic patterns autonomously. We show that such pattern generation can also be multiplexed with pattern classification and temporal integration of information in the same neural circuit. These results can be interpreted as showing that periodic activity provides a second channel for communication in neural systems which can be used to synchronize or coordinate spatially separated processes, without encumbering local real-time computations on spike trains in diverse neural circuits.

© 2005 Elsevier Ltd. All rights reserved.

Keywords: Neural microcircuit; Oscillations; Synchronization of neural activity; Real-time computing; Pattern classification; Learning; Computational models; Dynamic synapses

1. Introduction

Theoretical results and computer simulations have shown that generic models for neural microcircuits are capable of integrating information from spatio-temporal spike patterns, thereby facilitating real-time computing on complex input streams (Maass, Natschläger, & Markram, 2002). Of course sensory input is not the only source of neuronal firing activity in real biological neural systems. In particular, firing activity in neural circuits is frequently entrained by additional oscillatory input. This oscillatory signal may be generated in a distributed way by the brain, or it can be generated by specific neural circuits. For example circuits in the spinal cord of vertebrates (Orlovsky, Deliagina, & Grillner, 1999) generate periodic firing patterns that are used to drive rhythmic movements such

as walking. We show in this article that both types of periodic activity are in principle compatible with temporal integration of information and fast computations on spike patterns, that need not even be phase-locked or otherwise coordinated with the phase or frequency of an externally or internally generated periodic activation pattern. In particular neural microcircuits appear to have the capability to multiplex the transmission of information contained in a periodic pattern, such as its current phase, with the transmission of information contained in nonperiodic input streams in such a way that both types of information can easily be separated by simple linear readouts that have been trained for such task.

However, our computer simulations also show that there is ‘no-free-lunch’ in the sense that simultaneous processing of two types of inputs in the same neural circuits reduces their capacity for each type of information processing. But this imposes no a-priori limitation on their computational power, since we also show that a somewhat larger neural circuit has even in the presence of additional oscillatory input the same capability for fast computations on spike patterns as a smaller circuit without oscillatory input. On the other hand, an additional oscillatory input may facilitate the synchronization of distributed processes since its phase can be read out with high fidelity from small ensembles of neurons in different spatial regions of the circuit.

[☆] The work was partially supported by the Austrian Science Fund FWF, project # P15386, and PASCAL project # IST2002-506778 of the European Union.

* Corresponding author. Tel.: +43 316 873 5822; fax: +43 316 873 5805.

E-mail addresses: alexander.kaske@mtc.ki.se (A. Kaske), maass@igi.tugraz.at (W. Maass).

We have tested the multiplexing of information processing on periodic and nonperiodic input streams on standard models for generic neural microcircuits (see Maass, Natschläger, & Markram, 2004). These circuits consisted of 900 leaky integrate-and-fire neurons¹ with biologically realistic models for dynamic synapses according to (Markram, Wang, & Tsodyks, 1998). Connectivity was assumed to be primarily local. More precisely, we assumed that the neurons were located on the integer points of a three-dimensional grid in space, where $D(a, b)$ is the Euclidean distance between neurons a and b . Synaptic connections between neurons were chosen randomly according to a probability law that favored synaptic connections between neurons a, b whose distance $D(a, b)$ was not much larger than a parameter λ : the probability of a synaptic connection from a to b was $C \exp(-D^2(a, b)/\lambda^2)$, with a suitable scaling constant $C > 0$. The short term dynamics of each synapse was chosen to correspond to the types of pre- and post-synaptic neurons (excitatory or inhibitory) according to data from (Gupta, Wang, & Markram, 2000) and (Markram et al., 1998). Further details of the circuits are described in the appended section on methods.

2. Periodic and nonperiodic inputs can be processed in parallel by generic neural circuits

In previous investigations of temporal integration and computations on input streams in neural microcircuit models (Maass et al., 2002, 2004; Legenstein et al., 2003) the firing activity in the neural circuits was caused exclusively by the input stream, and all firing in the circuit ceased within 50–150 ms after the input stopped. In this section we demonstrate that temporal integration and real-time computing on complex input streams is also possible if the neural circuit is activated by some additional oscillatory input that keeps the circuit active independently of the information-carrying input stream. Obviously, such additional oscillatory input just represents ‘noise’ from the point of view of computing on the primary (nonperiodic) input stream. In fact, this additional oscillatory input may be seen as a large source of noise from the point of view of readout neurons that aim at extracting temporally integrated information from the current firing activity (‘liquid state’) of the circuit. The liquid state $x(t)$ of a circuit is defined in (Maass et al., 2002) as that part of the current state of the circuit that is in principle accessible to a readout neuron that receives synaptic input from all neurons in the circuit. Hence the number of components of $x(t)$ is equal to the number of neurons in the circuit, and the component of $x(t)$ that corresponds to a neuron v in the circuit is a low pass filtered version of the spike train emitted by neuron v . More precisely: a spike of

neuron v at time $t' \leq t$ adds $\exp(-(t-t')/30 \text{ ms})$ to the corresponding component of $x(t)$, where the time constant of 30 ms reflects the assumed temporal dynamics of receptors and active channels as well as the membrane time constant of the readout neuron. Examples of liquid states of circuit models considered in this article are shown in Fig. 5.

The periodic input may activate substantially more neurons in the circuit than the information-carrying online input streams (see Fig. 1a,b). Furthermore the phase of such periodic firing activity may vary in relation to the primary input stream whose timing is typically determined by events in the external world. Consequently the equivalence classes of liquid states that all represent equivalent information with regard to the online input stream (e.g. information about the second to last spike pattern that had previously been injected into the circuit) become substantially larger and more diverse in the presence of an unrelated oscillatory input: the ‘up-phases’ of the oscillatory activity tend to overwrite the trace of activity left over from a preceding input segment, and these ‘up-phases’ affect this memory trace in different ways, depending on the accidental relationship between the phase of the oscillatory input relative to the ‘phase’ of the preceding temporal spike pattern. We show that nevertheless simple linear readouts can be trained to become invariant to all these different appearances of the same information caused by unrelated oscillatory activity. It should however be stressed that these readouts need to be trained to become invariant to specific types of background noise. At least in our model we are not aware of a-priori reasons, why the inherent bias of linear separators should facilitate automatic invariance to new types of high-amplitude background noise.

Since neural microcircuits in the nervous system often receive salient input in the form of spatio-temporal firing patterns (e.g. from sensory neurons, or from other brain areas), we have concentrated on online input streams of this type. Such firing patterns could for example represent visual information received during a saccade, or the neural representation of a phoneme or syllable in auditory cortex. The online input streams in our computer simulations consisted of 10 simultaneous spike trains over 1000 ms. Each of the 4 segments of length 250 ms of such input stream was a noisy version of one of two templates that had been generated by 10 Poisson spike trains over 250 ms each, with a firing rate of 20 Hz; see Fig. 2. The computational task was to classify in real-time at the moment when this input stream of length 1000 ms ended for each of the preceding 4 segments of length 250 ms which of the corresponding two templates had been used to generate that segment.²

It is clear that this task requires substantial capability for temporal integration of information over several hundreds

¹ Only for the simulation whose results are reported in Fig. 3c) a circuit with 1800 neurons was used.

² Similar tests can also be carried out with more than two templates for each segment, but larger circuits are needed to achieve the same level of performance for that.

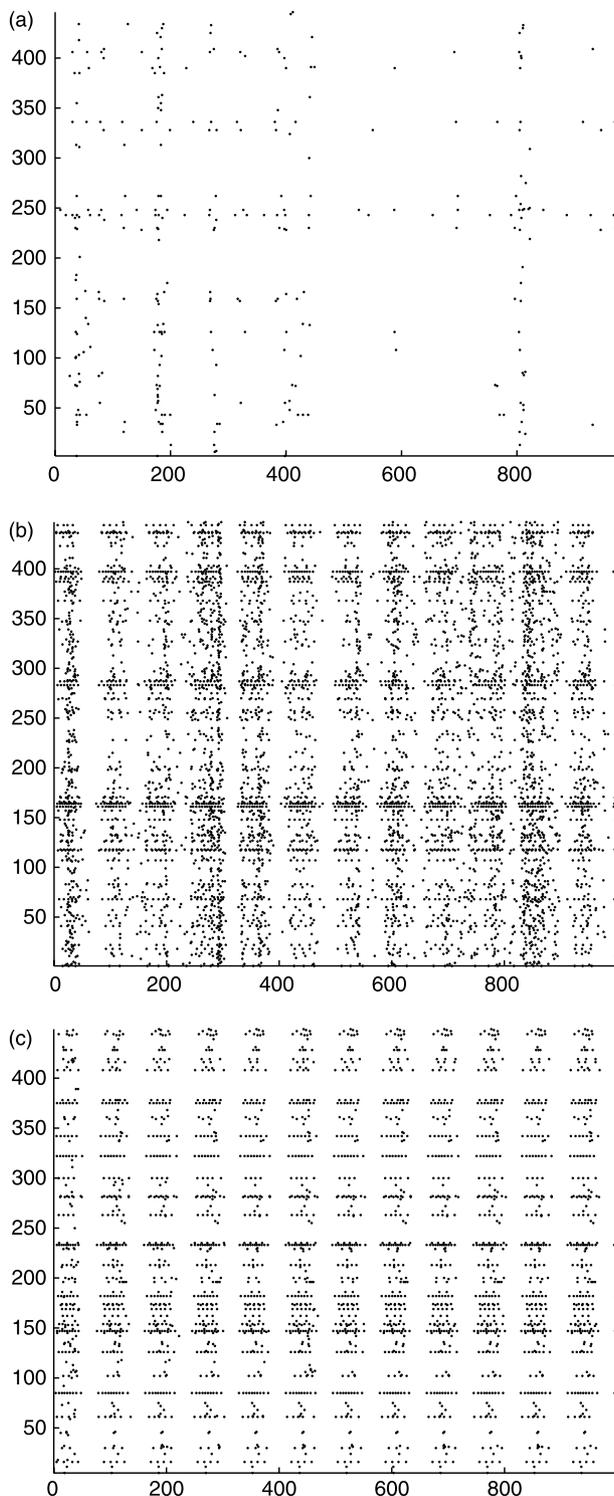


Fig. 1. Response of the generic microcircuit model to an input stream consisting of 10 Poisson spike trains. Spike times of the neurons 1–450 (y-axis) of the microcircuit are marked as dots. The x-axis denotes time (in ms). (b) The same representation of spike times in the same microcircuit for the case where in addition to the spike input each neuron in the circuit receives an additional analog sinusoidal input current at 12 Hz. Information about the online input stream is now encoded through modulations of the firing activity during up-phases of the oscillation. (c) For comparison the substantially more stereotypical firing response of the same circuit is shown for the case where it only receives the oscillatory input, and no multiplexing occurs.

of ms. We have demonstrated that this temporal integration was carried out by the circuits (in spite of the additional oscillatory input) by training 4 McCulloch-Pitts neurons to extract at a later point in time this temporally integrated information. They received as input the 450-dimensional liquid state $x(1000)$ of the circuit at time $t=1000$ ms (with time $t=0$ defined as the moment when the online input stream had begun). They were trained to output in real-time (in fact instantly) the labels (0 or 1) of the templates that had been used to generate the 4 segments of the preceding input stream. Obviously in the presence of additional oscillatory input whose phase can have any random value at time $t=1000$ ms, the liquid state $x(1000)$ is largely determined by this random phase of the oscillatory input (and is in addition influenced by the jitter used to generate the input streams from the underlying templates). Nevertheless Fig. 3b shows that the template that had been used to generate the last 250 ms segment of the input stream can be classified with the same precision as for the case without additional oscillatory input (see Fig. 3a). The recovery of information from the liquid state $x(t)$ at time $t=1000$ ms about the templates that had been used to generate the *earlier* segments of the input stream is substantially more difficult, since the spikes caused by these input segments were overwritten by independently chosen subsequent input segments and by several up-phases of the oscillatory input (with a phase that varies randomly from trial to trial, and is not communicated separately to the readouts that are trained to extract information from the liquid state $x(1000)$). Fig. 3b shows that the classification capability of these readouts on test data is reduced by additional oscillatory input, but is still above chance level (i.e. above an error fraction of 0.5). Furthermore a larger circuit (with 1800 instead of 900 neurons) has even in the presence of additional oscillatory input the same capability for temporal integration and real-time classification of information from complex spike patterns (see Fig. 3c) as the smaller control circuit without additional oscillatory input (for which results are given in Fig. 3a). Such automatic improvement of classification correctness for a larger circuit of the same type has previously already been reported in cases without oscillatory input (Maass et al., 2002). We refer to the methods section for further details of these computer simulations.

Whereas the additional oscillatory input reduces the information processing capability of the circuit for the unrelated input stream, it also introduces a second channel for communication and computing in the same neural microcircuit model. The phase of the oscillatory input can be read off from the liquid state $x(t)$ at any time t (see Fig. 4a), and could potentially be combined with the information about preceding spike patterns extracted from $x(t)$ by other readouts. Furthermore the relatively high amplitude of the oscillatory input (see Fig. 1) that we had chosen for demonstration purposes, has the advantage that information about its phase is broadcast quite clearly throughout the circuit, and hence can be recovered from the current firing

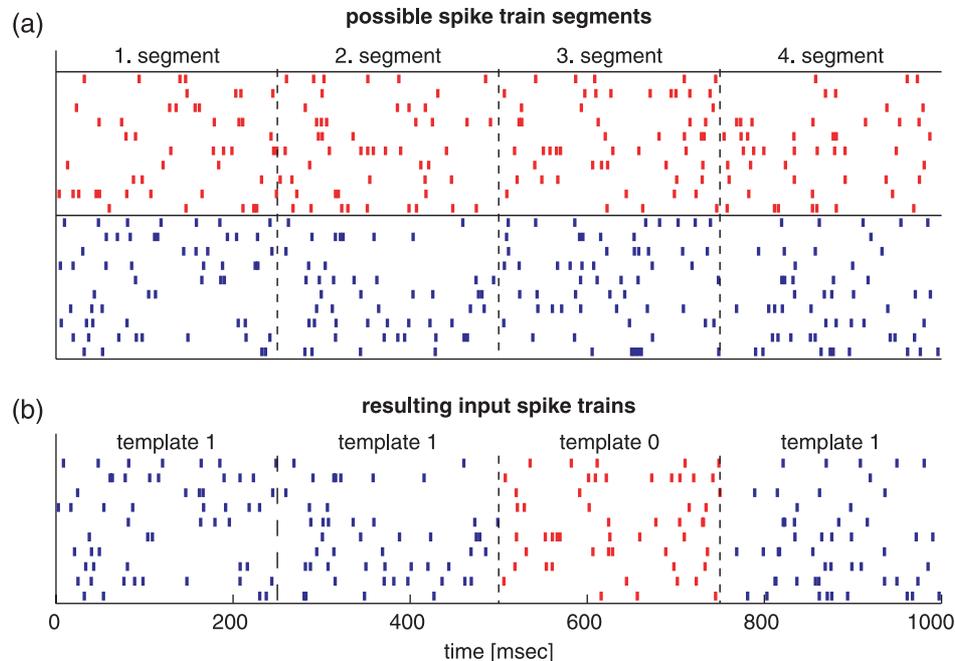


Fig. 2. Structure of online input streams used for all simulations. Each input consists of 10 spike trains of length 1000 ms generated from 4 segments of length 250 ms each. (a) For each segment two labeled templates 0 (upper row) and 1 (lower row) were generated randomly (more precisely, as Poisson spike trains with a frequency of 20 Hz). (b) The actual input streams for our simulations were generated by choosing randomly for each segment i , $i = 1, \dots, 4$, one of the two associated templates (labeled 0 or 1), and then generating a noisy version of that template by moving each spike by a different amount drawn from a Gaussian distribution with mean 0 and SD 4 ms.

pattern of a subpopulation of a few adjacent neurons anywhere in the circuit (see Fig. 4b), in spite of the unrelated input stream that may represent just a source of noise from this perspective (although in a biological circuit both sources might also contain complementary pieces of information that need to be combined by suitable readouts). Such information about the oscillatory input could be used for other computational tasks at different locations in a larger circuit or system, such as for example tasks that require a global synchronization of local processes.

3. Multiplexing pattern generation with real-time computing within the same circuit

With small modifications the same models for generic neural microcircuits that were previously used for real-time computing on complex input streams can also be used to generate periodic patterns. It suffices for example, to make sure that the probability of a synaptic connection from neuron a to neuron b is somewhat larger in case that a is an excitatory neuron. This assumption is justified both for the cortex (based on the wide axon trees of pyramid cells) and for the spinal cord in a variety of species (tadpole (Roberts et al., 1998), lamprey (Buchanan, 1982)). Then, if one just injects a constant input current into the neurons on layer 1 of the circuit, a periodic pattern is generated (see Fig. 5a). A more stable and spatially more clearly organized pattern (a traveling wave) is generated by such circuit if there exists

a spatial gradient for connection probabilities in the circuit, even if the strength of the tonic input is uniform throughout the circuit (see Fig. 5b); details can be found in the methods section. Another way to trigger the emergence of traveling waves in a generic neural microcircuit model is to choose the synaptic connection density uniform throughout the circuit (like for Fig. 5a), and to introduce a spatial gradient for the probability that a neuron receives a tonic depolarizing input current (see Section 6 for details). This option (see Fig. 5c), which may be viewed as the most realistic one for the generation of periodic patterns in the cortex, is used for the subsequent analysis discussed in Fig. 6.

The emergence of traveling waves in generic neural microcircuits is consistent with experimental data, in particular with the generation of traveling waves by slices of cerebral cortex (Golomb & Amitai, 1997). The generation of two-dimensional traveling waves in generic models for neural microcircuits could also be expected on the basis of preceding theoretical and simulation work for one-dimensional circuit models (Golomb & Ermentrout, 1999) and simpler models for two-dimensional circuits (Osan & Ermentrout, 2001; Kaske, Weinberg, & Coester, 2003). It is shown in (Kaske & Bertschinger, 2005) that the frequency, orientation and wavelength of the traveling waves generated by such generic neural microcircuit models can be influenced by various details of a biologically more realistic circuit model, such as the inclusion of commissural inhibitory neurons that impose inhibition on neurons on the other side of a spatially extended circuit, the presence of

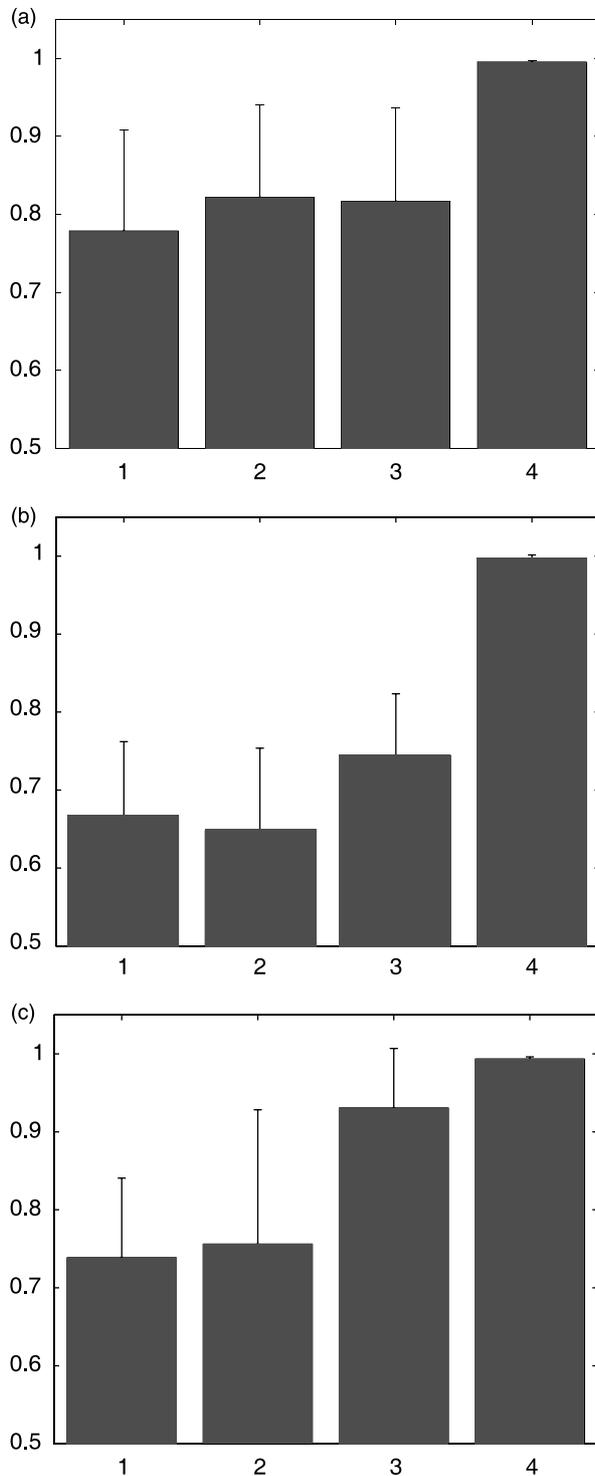


Fig. 3. Evaluation of the information about preceding spike inputs contained in the liquid state $x(t)$ of the neural microcircuit at time $t=1000$ ms. The y -axis gives the probability of a correct classification by a trained linear readout on test data (new noisy variation of spike templates), for each of the 4 segments of length 250 ms. Labels on the x -axis give the position of the segment within this spike input over 1000 ms for which such classification is carried out (there exist 4 readout neurons, one for each segment, which carry out their classification simultaneously and in parallel when the 4th input segment ends at time $t=1000$ ms. (a) Fraction of correct classifications (on test data) of the trained readout at time $t=1000$ ms of the templates used to generate each of the 250 ms segments of the preceding

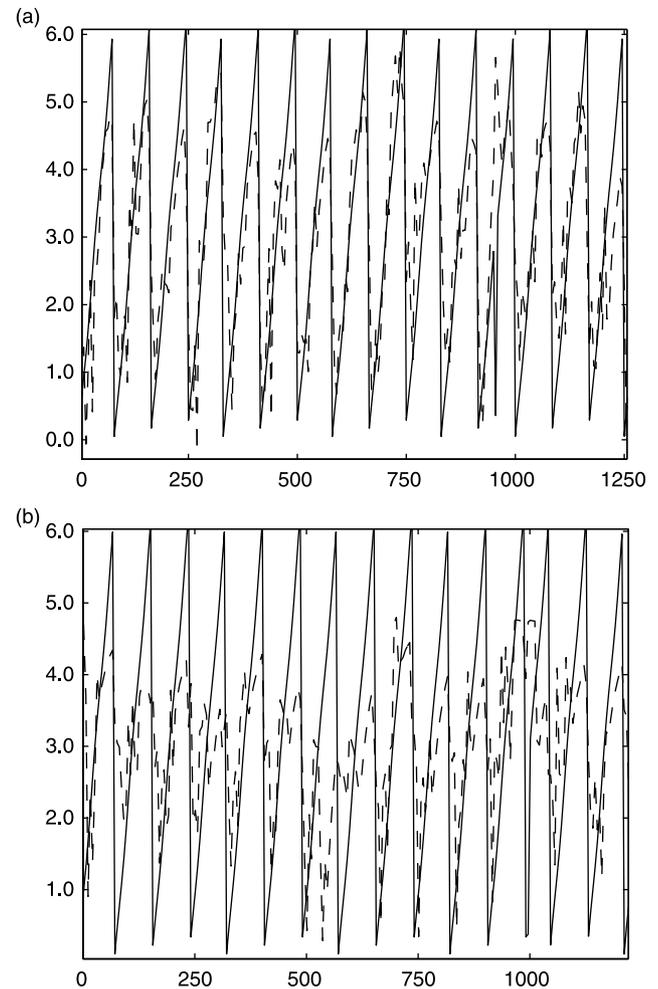
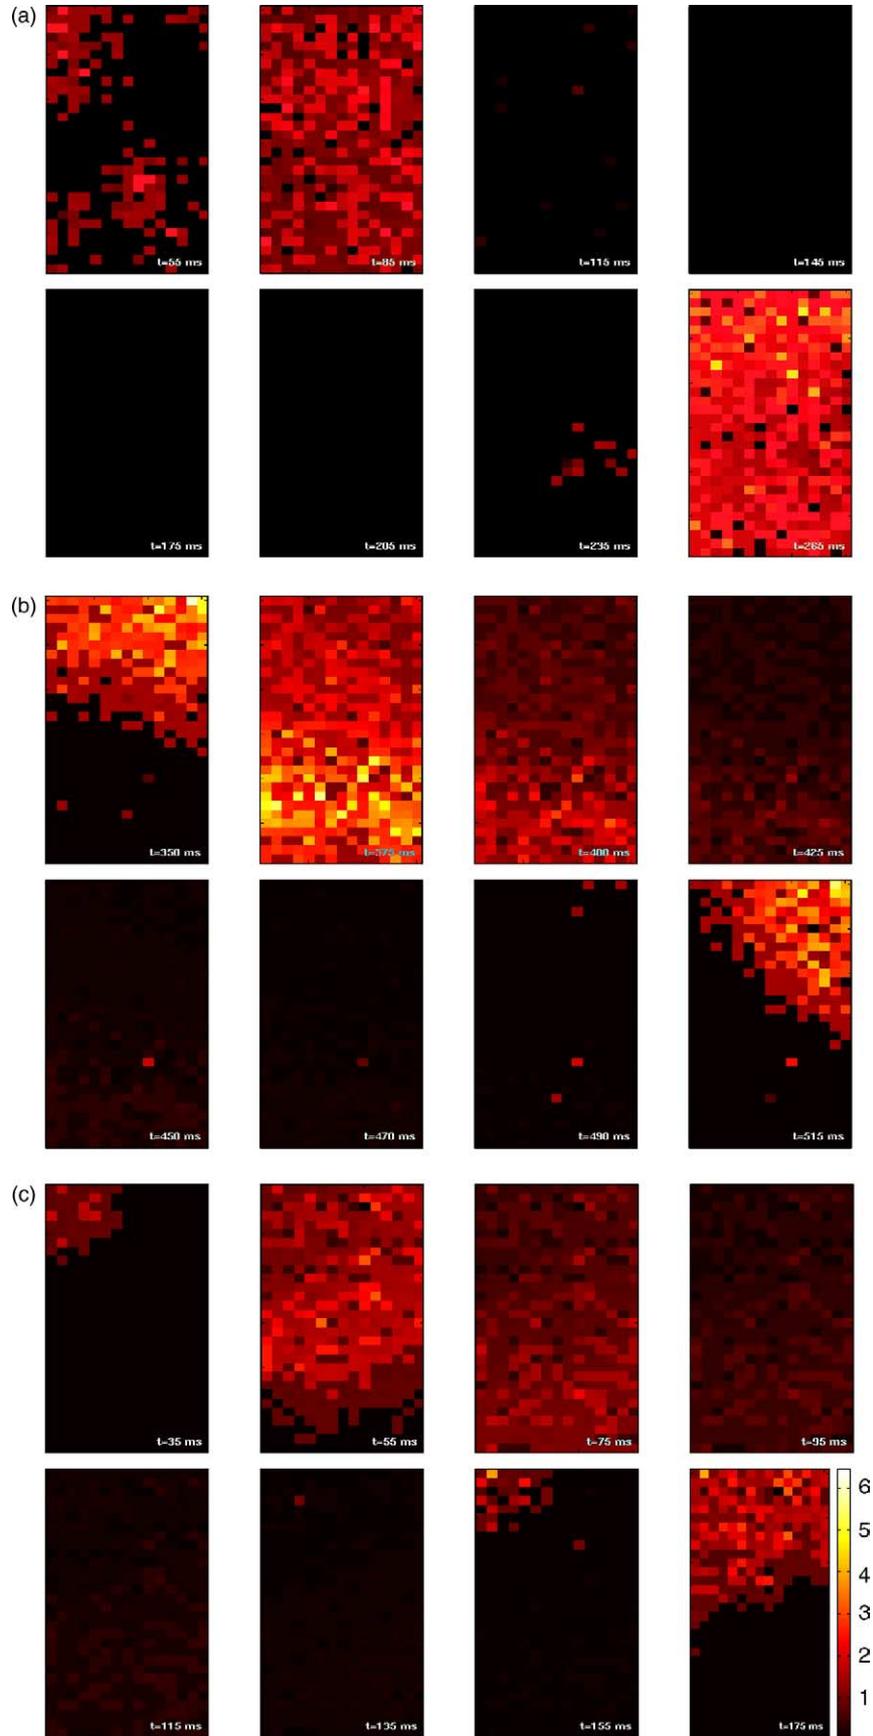


Fig. 4. Analysis of information transmission in generic neural microcircuits via oscillations. The solid line shows the phase of the oscillatory input as a function of time (x -axis in ms). (a) Dashed line shows the output of a linear readout that was trained to extract at any time t from the liquid state $x(t)$ of the 450 neurons on layer 1 the current phase of the oscillatory input. (b) Same as in (a), but for a linear readout that receives as input only 50 components of $x(t)$ corresponding to 50 neurons at a randomly chosen location of the circuit. This demonstrates that the current firing activity of the neurons in the circuit contains not only information about preceding aperiodic spike input, but also about the current phase of the oscillatory input. Hence the latter may be used to coordinate spatially distributed responses to an aperiodic spike input (that could represent for example proprioceptive or sensory input).

bursting neurons, and the inherent dynamics of depressing synapses in the circuit.

To the best of our knowledge, no attempt was previously made to construct circuits of spiking neurons that combine

input stream. Error bars represent the standard deviation. (b) Same experiment for the same circuit but for the case where all neurons in the circuit receive an additional oscillatory input. (c) Same as in panel b, but for a circuit of twice the size (consisting of 1800 neurons). The relative good classification performance for the pattern in the first time segment of these experiments (compared with that in the second time segment) results from the fact that only the membrane potential but not synaptic states were randomly chosen at the beginning of each simulation.



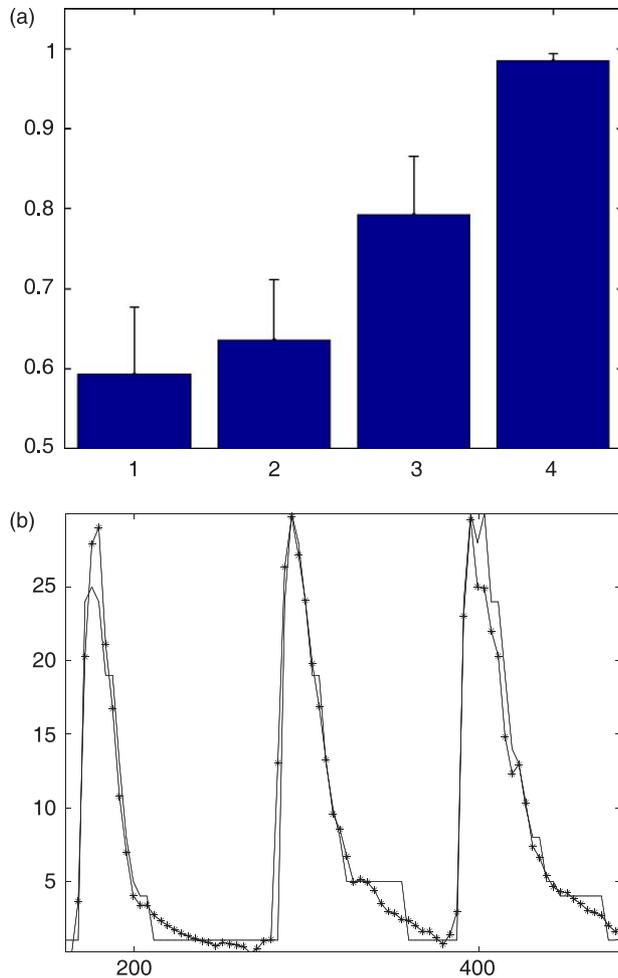


Fig. 6. Information processing capability of a pattern generating circuit. (a) Performance of the same task as considered in Fig. 3 (for an additional input stream consisting of 10 spike trains). (b) Estimate of the current spatial phase (in terms of the order number of the horizontal row of neurons in Fig. 5 that currently has the highest activity) of the traveling wave generated by the circuit. Target values are shown as solid line. The outputs of a linear readout (that received as input at any time t the 450-dimensional liquid state $x(t)$ corresponding to the 450 neurons on layer 1 of the circuit) which was trained to estimate this target value at any time t are marked by stars. The rather good performance of this readout shows that the pattern generated by the circuit is not disturbed too much by the online spike input, and may be viewed as a quasi-independent communication channel within the same neural microcircuit.

within the same circuit pattern generation with online processing of online input streams, that could represent for example proprioceptive or sensory feedback to the spinal cord. But this question appears to be of interest, since there exists reasonable experimental evidence that pattern

generation and online processing of proprioceptive or sensory feedback cannot be separated into distinct neural circuits in spinal cord (Baev, Esipenko, & Shimansky, 1991; Orlovsky et al., 1999; Perreault, Shefchyk, Jimenez, & McCrea, 1999). We have examined the performance of pattern-generating circuit for the same computational task as considered in Section 2. We have focused on the circuit that generated traveling waves according to the option considered in Fig. 5c (with a spatial gradient in the distribution of the time-invariant tonic input), and injected into such circuit an additional input stream consisting of 10 spike trains over 1000 ms as shown in Fig. 2b. The results depicted in Fig. 6a show that the capability of such circuit for pattern classification and temporal integration is quite impressive, almost the same as for the case with an external oscillatory input considered in Fig. 3b. This is particularly remarkable in view of the fact that no attempt was made to coordinate the timing of the self-generated traveling wave with that of the input stream. This would actually have been very difficult since the periods of the traveling wave were modulated by the online input stream.

As for the case with an external oscillatory input, the self-generated traveling wave opens up a second channel for communication in the same neural circuit. A linear readout can be trained to extract at any time t from the liquid state $x(t)$ of the circuit the current spatial phase of the self-generated traveling wave (see Fig. 6b). Obviously the result of such readout can easily be combined with the results of other readouts that classify spatio-temporal spike patterns in the online input stream. Hence our model provides a simple but biologically not unrealistic computational architecture for the control and modification of periodic movements such as walking in generic neural circuits (e.g. in the spinal cord) that simultaneously receive proprioceptive or sensory feedback.

4. Low dimensional noise in high dimensional state spaces

The capability of a neural circuit to support classification and temporal integration of spike patterns as needed for the computational task considered in this article depends on the amount of information contained in the current firing activity (i.e. the current liquid state $x(t)$) of the circuit at the time t when the readout neuron has to produce an answer. For the case of a linear readout neuron that is considered in this article, this capability depends more precisely on the linear separability of the two classes of

Fig. 5. Generation of periodic patterns by generic neural microcircuits (with excitation reaching farther than inhibition, see text). Light color indicates high firing activity of a neuron at the corresponding position in the circuit (scaled according to current firing rate in Hz). Shown are in each case 8 frames of 450-dimensional liquid states $x(i\Delta)$, $i=0, \dots, 7$, generated by the 30×15 neurons on layer 1 of the circuit. Δ is chosen sufficiently large so that the time interval $[0, 7\Delta]$ covers a little more than one period of the periodic pattern generated by the circuit. All three circuits receive simultaneously spikes of the type shown in Fig. 2b. (a) The same constant input current is injected into all neurons on layer 1. $\Delta=15$ ms. (b) The same constant input current is again injected into all neurons on layer 1. But due to a gradient in the number of synapses from the top to the bottom, a traveling wave emerges that travels from the top to the bottom. $\Delta=25$ ms. (c) There is no gradient in the synaptic connectivity in this circuit, but there exists a gradient in the probability that a neuron receives a constant tonic input current (see Appendix). $\Delta=20$ ms.

liquid states $x(t)$ that emerge at trials when the readout neuron is supposed to output 1 and on trials where the readout neuron is supposed to output 0. Some examples for the diversity of liquid states that may fall into the same equivalence class were given in (Maass et al., 2004).

When periodic patterns are superimposed (with a random phase) on the spatio-temporal spike pattern that are to be classified, the diversity of liquid states belonging to one equivalence class appears to grow dramatically. This can be seen when one tries to visualize the difference in liquid states resulting from the firing activity, shown in panel a of Fig. 1 with that of panel c in Fig. 1. But visual appearance of diversity of liquid states is of course not really relevant for the performance of linear readouts from such circuit with several hundred neurons. From their perspective the added diversity of liquid states resulting from periodic background activity looks relatively harmless as long as this periodic background activity is inherently low dimensional, i.e. if it can be characterized after a suitable coordinate transformation (for example principal component analysis) by a low dimensional dynamical system. In this case a linear classifier has a chance to find a separating hyperplane that is essentially parallel to the low dimensional subspace spanned by the background activity, so that the resulting equivalence classes become largely invariant to such superimposed background activity. This is clearly demonstrated by the results in Fig. 3 where the separating hyperplanes were computed by linear regression, and are therefore optimal (with regard to mean squared errors). The last bar in panel a of Fig. 3 indicates that even without background activity the correspondent equivalence classes in the high dimensional state space are not linearly separable. But the last bar in panel b shows that the best separating hyperplane has managed to become invariant to the additional noise imposed by the periodic background input. This analysis is oversimplified since it ignores the facts that the periodic background noise interacts in the circuit nonlinearly with the signal, and that in particular the impact of earlier segments of the periodic background input can no longer be confined to such low dimensional subspaces. On the other hand panel c of Fig. 3 indicates that not the absolute dimensionality of the noise dynamics is essential for the performance of linear readouts, rather the relationship between the noise-dimension and the total dimension of the circuit dynamics: If the latter is increased, then the impact of the same low dimensional background noise on the performance of optimal linear readouts is reduced.³

³This explanation suggests that our results do not depend on the particular neuron model or synapse model that are used. Partial verifications of this conjecture can be found in more recent studies of wave generation (Kaske & Bertschinger, 2005) and pattern classification (Häusler & Maass, submitted) in circuits consisting of Hodgkin-Huxley neurons.

5. Discussion

We have demonstrated in Section 2 that previously proposed paradigms for temporal integration of information and real-time classification of spatio-temporal spike patterns are compatible with biologically realistic assumptions about an oscillatory background activation of neural circuits and systems. Note that the firing activity in such circuits does not cease during periods where no phasic input arrives. Our computer simulations suggest that such oscillatory background drive reduces somewhat the capability of neural microcircuits to hold and process information contained in input segments that occurred several hundred ms ago. But we have also shown that this poses no serious limitation for the computational power of neural microcircuits in the presence of oscillatory background activation, since a circuit with twice the number of neurons achieves an equally good performance in spite of additional oscillatory input compared with a circuit that receives no oscillatory input. This holds even if the phase of the oscillatory background drive has a random relation to the timing of the information-carrying input stream. On the other hand, additional oscillatory input makes new information available in all parts of the circuit, such as the current phase of the oscillatory background drive, that could potentially be used to coordinate information processing or muscle activation in spatially segregated parts of a larger neural system.

In Section 3 we have demonstrated that temporal integration of information and classification of spatio-temporal spike patterns is also compatible with the generation of periodic patterns within the same circuit. Again, such multiplexing reduces somewhat the computing capabilities of a circuit, but it facilitates computational tasks where a modulation of a generated pattern in response to online streams of proprioceptive or sensory feedback is desirable. Hence our model predicts that such multiplexing of heterogeneous tasks could for example take place in neural circuits of the spinal cord. The analysis given in Section 4 suggests that the feasibility of such multiplexing is no mystery, since it depends only on the low dimensionality of the superimposed periodic activation patterns. The examples presented in this article demonstrate that superimposed periodic activation patterns may drastically change the appearance of transient circuit states, but tend to reduce the performance of suitably trained readouts only moderately.

6. Methods: specifications of models and simulations

Generic neural microcircuit models were constructed as described in (Maass et al., 2002). Neurons were modeled as leaky integrate-and-fire neurons with a membrane time constant of 30 ms and randomly chosen initial states. Transmission delays between neurons were chosen to be

quite small (1.5 ms for connections between excitatory neurons, 0.8 ms for the other connections) in accordance with biological data.

We modeled the synaptic dynamics according to the model proposed in (Markram et al., 1998), with synaptic parameters U, D, F . The model predicts the amplitude A_k of the EPSC for the k th spike in a spike train with interspike intervals $\Delta_1, \Delta_2, \dots, \Delta_{k-1}$ through the equations

$$\begin{aligned} A_k &= w u_k R_k & u_k &= U + u_{k-1}(1 - U)\exp(-\Delta_{k-1}/F) \\ R_k &= 1 + (R_{k-1} - u_{k-1}R_{k-1} - 1)\exp(-\Delta_{k-1}/D) \end{aligned} \quad (1)$$

with hidden dynamic variables $u \in [0, 1]$ and $R \in [0, 1]$ whose initial values for the first spike are $u_1 = U$ and $R_1 = 1$ (see (Maass & Markram, 2002) for a justification of this version of the equation, which corrects a small error in (Markram et al., 1998)).

The parameters U, D , and F were randomly chosen from Gaussian distributions that were based on empirically found data for such connections (Gupta et al., 2000). Depending on whether the input was excitatory (E) or inhibitory (I), the mean values of these three parameters (with D, F expressed in seconds) were chosen to be 0.5, 1.1, 0.05 (E), 0.25, 0.7, 0.02 (I). The SD of each parameter was chosen to be 10% of its mean (with negative values replaced by values chosen from an uniform distribution between 0 and two times the mean).

The neurons were located on the grid points of a three-dimensional $2 \times 30 \times 15$ grid (exception: the larger circuits used for generating the results shown in Fig. 3c employed a $2 \times 30 \times 30$ grid), yielding a basically two-dimensional circuit consisting of 2 layers. 20% of the neurons were randomly chosen to be inhibitory for the circuits considered in Section 2. The parameter λ in the formula $C \exp(-D^2(a, b)/\lambda^2)$ for the connection probability for neurons a and b (see Section 1) was set to be 2 for the circuits considered in Section 2.

Spike trains from input streams as shown in Fig. 2b generated EPSPs in a randomly chosen subset of neurons on layer 1 (each of the 10 spike trains caused EPSPs in a neuron of layer 1 chosen with probability 0.6), with randomly chosen amplitudes. The distribution of amplitudes was not optimized for any of the tasks or circuits discussed in this paper (although a proper balance with oscillatory input is likely to improve performance in each case). New noisy variations of each spike pattern were generated for each run of the simulation by moving each spike in a spike pattern by a different amount drawn from a Gaussian distribution with mean 0 and SD 4 ms. The pattern classification results reported in Fig. 3 and Fig. 6 were achieved by 4 perceptrons (i.e. linear weighted sums with a decision threshold) that received as input the liquid state $x(1000)$ of the neurons on layer 1 of the 2-layer circuit (see Section 2). A separate perceptron was used for the retroactive classification of each of the 4 segments of the preceding input stream. Reconstruction of the phase of oscillations, as reported in

Figs. 4 and 6, was carried out by an additional linear readout that received as input at any time t the liquid state $x(t)$ of the neurons of layer 1.

All readouts were trained by linear regression for 2000 training runs, and the performance of the trained readouts was recorded for 1200 test runs. In each run the initial state of neurons in the circuit, the templates and spikes jitter in the input stream, and (for the simulations discussed in Section 2) the phase of the oscillatory input relative to the beginning of the 1000 ms spike input were randomly chosen. This protocol was repeated for 10 randomly generated circuits, and Figs. 3 and 6 give the mean and SD resulting from these 10 repetitions of the whole experiment.

The pattern generating circuits considered in Section 3 were constructed slightly differently. Pattern generation emerges if the range of excitatory connections is on average wider than that of inhibitory connections (Golomb & Amitai, 1997; Golomb & Ermentrout, 1999). In order to achieve this, for section 3 all neurons on layer 1 of the $2 \times 30 \times 15$ grid were chosen to be excitatory, and all neurons on layer 2 were chosen to be inhibitory. The parameter λ had a value of 3 in the formula $C \exp(-D^2(a, b)/\lambda^2)$ for the probability of a synaptic connection from an excitatory neuron a to an excitatory or inhibitory neuron b . For connections from inhibitory neurons a to excitatory neurons b this parameter λ had a value of 1.5 (synaptic connections among inhibitory neurons were not modeled in Section 3). Deviating from (Maass et al., 2002) and Section 2 the scaling parameter A for the amplitude ('weight') of a synapse was chosen in section 3 from a Gaussian distribution with a mean of 3 and SD 0.75 for synapses between excitatory neurons, and a mean of 1 and SD 0.5 for all other synapses. In order to achieve a more regular circuit structure (in particular in order to avoid border effects) synaptic connections were randomly drawn until each neuron had a prespecified fan-in of 32 (of which 11 were reserved for inhibitory synapses in the case of an excitatory neuron). Only the circuit considered for Fig. 5b had a gradient in the synaptic density, achieved by linearly decreasing the fan-in from 32 to 11 along the longest axis, i.e., along the 30 2×15 slices of the $2 \times 30 \times 15$ circuit.

The spatial gradient in the probability that a neuron received tonic inputs for the circuits discussed in Figs. 5c and 6 was implemented as follows. A source of tonic input currents was placed in the middle of the 2×15 slice that appears as the top row in the 30×15 pixel matrices shown in Fig. 5. This input source a was connected to neurons b in the circuit with probability $C \exp(-D^2(a, b)/\lambda^2)$ with $\lambda = 5$. Thus only neurons that correspond to the upper rows of the pixel matrices shown in Fig. 5 had a high probability of receiving this constant input current.

References

- Baev, K. V., Esipenko, V. B., & Shimansky, Y. P. (1991). Afferent control of central pattern generators: experimental analysis of locomotion in the decerebrate cat. *Neuroscience*, *43*(1), 237–247.
- Buchanan, J. T. (1982). Identification of interneurons with contralateral, caudal axons in the lamprey spinal cord: synaptic interactions and morphology. *Journal of Neurophysiology*, *47*(5), 961–975.
- Golomb, D., & Amitai, Y. (1997). Propagating neuronal discharges in neocortical slices. *Journal of Neurophysiology*, *78*, 1199–1211.
- Golomb, D., & Ermentrout, G. B. Golomb and Ermentrout. (1999). Continuous and lurching traveling pulses in neuronal networks with delay and spatially decaying connectivity. *Proceedings of the National Academy of Sciences USA*, *96*(23), 13480–13485.
- Gupta, A., Wang, Y., & Markram, H. (2000). Organizing principles for a diversity of GABAergic interneurons and synapses in the neocortex. *Science*, *287*, 273–278.
- Häusler, S., Maass, W. (submitted). *A statistical analysis of information processing properties of lamina-specific cortical microcircuit modes.*
- Kaske, A., & Bertschinger, N. (2005). Travelling wave patterns in a model of the spinal pattern generator using spiking neurons. *Biological Cybernetics*, *92*(3), 206–218.
- Kaske, A., Weinberg, G., & Coester, J. (2003). Emergence of coherent traveling waves controlling quadrupled gaits in a two-dimensional spinal cord model. *Biological Cybernetics*, *88*, 20–32.
- Legenstein, R. A., Markram, H., & Maass, W. (2003). Input prediction and autonomous movement analysis in recurrent circuits of spiking neurons. *Reviews in the Neurosciences (Special Issue, on Neuroinformatics of Neural and Artificial Computation)*, *14*(1–2), 5–19.
- Maass, W., & Markram, H. (2002). Synapses as dynamic memory buffers. *Neural Networks*, *15*, 155–161.
- 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.
- Maass, W., Natschläger, T., & Markram, H. (2004). Computational models for generic cortical microcircuits. In J. Feng (Ed.), *Computational neuroscience: a comprehensive approach* (pp. 575–605). Boca Raton: Chapman and Hall/CRC. Chapter 18.
- Markram, H., Wang, Y., & Tsodyks, M. (1998). Differential signaling via the same axon of neocortical pyramidal neurons. *PNAS*, *95*, 5323–5328.
- Orlovsky, G. N., Deliagina, T. G., & Grillner, S. (1999). *Neuronal control of locomotion*. Oxford: Oxford University Press.
- Osan, R., & Ermentrout, B. (2001). Two dimensional synaptically generated traveling waves in a theta-neuron neural network. *Neuro-computing*, *38*(40), 789–795.
- Perreault, M. C., Shefchyk, S. J., Jimenez, I., & McCrea, D. A. (1999). Depression of muscle and cutaneous afferent-evoked monosynaptic field potentials during fictive locomotion in the cat. *Journal of Physiology*, *521*(Pt3), 691–703.
- Roberts, A., Soffe, S. R., Wolf, E. S., Yoshida, M., & Zhao, F. Y. (1998). Central circuits controlling locomotion in young tadpoles. *Annals of the New York Academy of Sciences*, *860*, 19–34.