

Markus Diesmann points about NeuroMat's Statement of Impact. The notes below have been prepared together with Moritz Helias and Tom Tetzlaff.

Point 1:

"Any abstract theory building requires simplifications and approximations. The question then is whether the results obtained analytically still describe nature or are just an artefact of the sequence of approximations. Today we are in a better position than ever before in validating the theoretical work. We can perform direct simulations of neuronal networks at the resolution of cells and synapses at full scale. Meaning with the number of neurons and synapses found in nature. Also there are single neuron models which are still simple like the variants of the leaky integrate-and-fire model but well validated by electrophysiology.

One example where such a validation is important is if a theory derives some statement in the limit of infinite network size. A direct simulation can assess whether the statement is really compatible with nature given the large but finite number of neurons in a biological neuronal network.

I think that in future reviewers will demand more often that analytical work is backed up by direct simulations of a slightly more realistic model."

This statement does not mean that abstractions and analytical work become less relevant. To the contrary, detailed simulation become easily so complex that they are as hard to understand as nature itself. Simplifications and analytical work is required to gain insight and build up intuition. Furthermore, we do not know the ground truth, also our microscopic models are approximations. Therefore, it is not always the case that a more microscopic model is closer to nature than an abstract one. The comparison of models on different levels of abstraction, however, increases the confidence in the conclusions made.

The simulation engines and data analysis software of computational neuroscience are becoming more mature. This means that the description of new models and the reproduction of their results becomes much easier. The neuroscientist can concentrate on the model instead of basic algorithms of numerical integration and communication. In manuscripts it suffices to reference software versions instead of providing large amounts of problem specific code. In the same way reviewers develop trust in the established systems and concentrate on the modeling results.

Many different models can be simulated with the same simulation engine. This is of advantage because in this way the neuroscientist has access to well validated and optimized code. For some models it is possible to simulate them several simulation engines. This is a good thing because it enables cross-validation on the highest level. An example of this is the implementation of the microcircuit model by Potjans & Diesmann (2014) for the simulation engine Brian by the group of Roque [3].

Technology is now making the next step in the digitization of the scientific workflow. One aspect is version control of model development and the formalization of issue tracking and the review of model descriptions and documentation. Code development platforms like GitHub are a first step in this direction. Often the process leading from the experimental data to the executable model description is more interesting than the model itself. Therefore, another aspect is the

formalization of the workflow with tools like snakemake. These digitized workflows can reproduce the figures of a published paper increasing the trust in the study and its reproducibility.

Students need to be educated in the use of these new tools.

As more executable model descriptions and the respective workflows are published as open source it becomes easier to test new ideas on the basis of published models. For example, complex single neuron models can be replaced by more abstract ones, or subnetworks can be replaced by population models.

Antonio R and his team have already made great contributions in advancing this. There is, for example, some work going on investigating the properties of the microcircuit model by Potjans & Diesmann if subject to other single neuron models.

Point 2:

"The scientific research section [of the Statement of Impact] mentions stochasticity a few times. I think it is important to make clear that NeuroMat has an understanding where the stochasticity in the brain comes from and works on deepening of this understanding. For example, in the cortex a lot of the stochasticity observed in single neurons seem to come from the chaotic or near chaotic dynamics of essentially deterministic neurons. If now a theory describes single neurons by a stochastic process and places them in a recurrent network by itself generating irregular activity, stochasticity may have been accounted for twice."

The spiking activity of a cortical neuron in the intact brain appears highly irregular. However, in vitro experiments show that neurons can respond reliably and with high precision to current transients injected into the soma. A prominent example is the study by Mainen & Sejnowski [1] based on the pioneering work of Bryant & Segundo [2]. These results suggest that the single-neuron behave like a deterministic machine on the voltage scales of interest and that the apparent irregularity is a reflection of the complex network dynamics. The leaky integrate-and-fire (LIF) model of a single neuron is a corresponding formalization.

Nevertheless, the experiment is limited in that the somatic injection of the current already effectively reduces the neuron to a point and ignores the dynamics of the synapses. In the cortex, synapses exhibit a rich set of dynamics on different time scales. On the time scales of minutes and above, synapses change the amplitude of the post-synaptic potential depending on the pre- and the post-synaptic activity and third factors like the neuromodulator dopamine. This process is called synaptic plasticity and fundamental for system-level learning. Also depending on the correlation of pre- and postsynaptic activity and other triggers, synapses can be removed and creating in a process called structural plasticity. On shorter time scales and a spike-by-spike basis two other phenomena are of interest. Short-term (SP) plasticity describes the effect that the amplitude of the post-synaptic potential (PSP) depends on the history of spikes arriving at this synapse. Facilitating synapses become stronger with subsequent spikes and depressing synapses get weaker, in the absence of spikes the amplitude tends towards a resting value. The PSP amplitude in addition has a biophysically stochastic component; the merging of a finite number of vesicles with the pre-synaptic membrane. This can cause a variability of the synaptic amplitude of the same magnitude as its mean value. In addition fluctuations of the membrane potential in the sub-millivolt range originate from the stochastic opening and closing of ion channels. The mechanism of action potential (AP) generation also has properties that are

suppressed by strong transients of an injected current. The spike threshold changes depending on the spike history by processes called adaptation and intrinsic plasticity.

Hartmann et al. [4] discuss various sources of stochasticity in the network context.

The detailed biophysics underlying these sources of variability on the single-neuron level are partially unknown or require complex systems of differential equations with many parameters. As in other areas of physics the neuroscientist may therefore decide to replace the complexity by a stochastic model. The randomness does, therefore, not necessarily describe a random process in nature but is an expression of our ignorance of details and the hope that the conclusions drawn from the stochastic model are sound. The gain is mathematical compactness, accessibility to analytical tools, and the chance to get insight.

Linear approximation

The first question arising is if and how the dynamics emerging in a network of stochastic model neurons is identical to its deterministic counterpart. Grytsyky et al. [6] investigate this question for the case of recurrent networks of leaky integrate-and-fire (LIF) model neurons. The authors provide a unified view on correlated activity in networks of LIF, model neurons, stochastic Hawkes processes, and stochastic rate equations.

There is a quantitative mapping between these three models capturing the global network states as characterized by Brunel [12]. The mapping is more accurate if the deterministic LIF network receives more external noise.

Generally replacing the deterministic units by stochastic ones should increase the variability and decrease correlations because the stochastic dynamics is added to the fluctuations generated by the recurrent network dynamics.

However Grytsyky et al. show that if the stochastic units are adjusted such that they exhibit the same linear response as the deterministic units also the cross-correlations are preserved to linear order.

Schematically:

$\text{autocorr(LIF)} = \text{autocorr(stochastic unit)}$ and $\text{linear approx(LIF)} = \text{linear approx(stochastic unit)}$

=> same correlation structure to linear order.

In Helias et al. [9] it is made explicit that in this approximation the equations for the cross-covariance have the same form.

Several laboratories have investigated this approximation. See for example also the works of the group of Stefan Rotter at Freiburg University [10] and [11]. The work of Senk et al. [7] contains a recent summary of the mapping from LIF to stochastic neuron models.

The clearest view on how effectively stochastic equations of motion arise from deterministic networks is, to our view, is obtained by applying methods from disordered systems. In Schuecker et al. [8] such methods are applied to neuronal networks. One finds that even if the noise ξ_i in the individual neurons is put to zero, one can derive an effective equation of motion that is stochastic (eq. 5, a noise η appears). The way how this stochasticity arises is quite interesting: First one finds that networks with different realizations of the random connectivity collectively have the same dynamics (with only minor realization-dependent variations); this is what is called "self-averaging" in spin glasses. Now, if one is only interested in the "average" network (see last paragraph on page 2), one throws away the knowledge about the individual

realization. Technically, this is called a quenched average over the connections, so connections are held fixed over time. The subsequent derivation in [8] shows that this ignorance of the detailed connectivity structure translates into stochasticity in time, namely the appearance of the noise η in eq. (5).

This view is related to Section 3.5 in Brunel & Hakim 1999: The analysis of finite size fluctuations performs an "annealed" approximation of the connectivity (redrawing it anew whenever a spike is generated). They achieve a quantitative agreement between power spectra in full simulations and in this analytical approximation. This implies that, if one would reinterpret the annealed approach as a stochastic neuron model, one would get the same power spectrum in such a stochastic network model as in the original, deterministic network.

The reason, we believe, is the self-averaging property of these networks, as explained above.

This is not to say that we believe that deterministic networks are equivalent to deterministic ones in every respect; the crucial point is that this identification works only for quantities that are self-averaging, too - for example the correlation functions averaged over many pairs of cells etc.

Beyond linear approximation

The linear response approximation described above may not be sufficient to describe higher-order correlations like the complexity distributions. This is the distribution of spike counts in a certain narrow bin across a population of neurons. It may also not be sufficient to describe the dynamics in networks deviating from the self-averaging property like feed-forward structures.

In particular Goedeke & Diesmann [13] show that for transient input the instantaneous spike rate cannot be described as function of the membrane potential excursion $U(t)$ alone but by the product of the derivative of U and an expansive function of U

$$\rho = \dot{U} * F(U)$$

It is unclear whether the existing stochastic neuron models are rich enough to capture this phenomenon.

The model of Galves and Loecherbach is similar to a Hawkes process. Okker et al [5] derive a related model using methods from field theory.

References

- [1] Mainen, Zachary F. and Sejnowski, Terrence J. (1995) Reliability of Spike Timing in Neocortical Neurons. *Science* 268:1503-1506
- [2] Bryant, Hugh L. and Segundo, Jose. P. (1976) Spike Initiation by transmembrane current: a white-noise analysis. *J. Physiol. Lond.* 260: 279-314
- [3] Shimoura, Renan O., Kamiji, Nilton L., Pena, Rodrigo F.O., Cordeiro, Vinicius L., Ceballos, Cesar C., Romaro Cecilia, & Roque, Antonio C. (2018). [Re] The cell-type specific cortical microcircuit: relating structure and activity in a full-scale spiking network model (Version 1.0). *Rescience*, 4(1). <http://doi.org/10.5281/zenodo.1244116>

- [4] Hartmann C, Lazar A, Nessler B, Triesch J (2015) Where's the Noise? Key Features of Spontaneous Activity and Neural Variability Arise through Learning in a Deterministic Network. PLoS Comput Biol 11(12): e1004640. <https://doi.org/10.1371/journal.pcbi.1004640>
- [5] Ocker GK, Josić K, Shea-Brown E, Buice MA (2017) Linking structure and activity in nonlinear spiking networks. PLoS Comput Biol 13(6): e1005583. <https://doi.org/10.1371/journal.pcbi.1005583>
- [6] Grytskyy D., Tetzlaff T., Diesmann M., Helias M. (2013). A unified view on weakly correlated recurrent networks. Frontiers in Computational Neuroscience 7:131.
- [7] Johanna Senk, Karolína Korvasová, Jannis Schuecker, Espen Hagen, Tom Tetzlaff, Markus Diesmann, Moritz Helias (2018) Conditions for traveling waves in spiking neural networks. arXiv:1801.06046 [q-bio.NC]
- [8] Jannis Schuecker, Sven Goedeke, and Moritz Helias (2018) Optimal sequence memory in driven random networks. PRX (accepted) <https://arxiv.org/abs/1603.01880>
- [9] Helias M, Tetzlaff T, Diesmann M (2013) Echoes in correlated neural systems. New J. Phys. 15:023002. DOI:10.1088/1367-2630/15/2/023002
- [10] Pernice V, Staude B, Cardanobile S, Rotter S (2011) How Structure Determines Correlations in Neuronal Networks. PLoS Comput Biol 7(5): e1002059. <https://doi.org/10.1371/journal.pcbi.1002059>
- [11] Volker Pernice, Benjamin Staude, Stefano Cardanobile, and Stefan Rotter (2012) Recurrent interactions in spiking networks with arbitrary topology. Phys. Rev. E 85:031916
- [12] Brunel N (2000) Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. J. Comput. Neurosci. 8 183-208
- [13] Nicolas Brunel and Vincent Hakim (1999) Fast Global Oscillations in Networks of Integrate-and-Fire Neurons with Low Firing Rates. Neural Computation 11(7):1621-1671
- [14] Goedeke S, Diesmann M (2008) The mechanism of synchronization in feed-forward neuronal networks. New Journal of Physics 10:015007. DOI:10.1088/1367-2630/10/1/015007