Dear Marco Leite and reviewers,

Thank you for the fair and considerate review of our manuscript. We agree with the overall comments of both reviewers, and have made changes accordingly. Below, we outline in detail the changes made, with reference to each reviewer's specific comments (in gray), and also include a diff pdf file which presents all added and removed text since the previous version. Furthermore, in consideration of the overall comments of both reviewers, we have made additional modifications to the abstract and main text where we felt that more clarity was warranted. The manuscript has benefited greatly from the reviewer comments, both in coherence and comprehension. We thank you again for the time put into this review, and we await your decision on this new version.

Kind regards,
Michele Nardin, James Phillips, William Podlaski & Sander Keemink

Detailed response:

Review by anonymous reviewer, 2021-05-04 16:52

Overall, I like the paper very much. I think it is a very nice extension of the SCN framework. The maths is thorough and the paper is written clearly (although a few clarifications are necessary).

Most of my criticism revolves around the lack of a clear biological mechanism behind multiplicative synapses. Even though it is not strictly necessary that the authors deliver a precise biological interpretation, the paper could benefit from a more detailed examination of possibilities. It would moreover strengthen the validity of multiplicative synapses.

Major

- I am still struggling to think of an implementation of multiplicative synapses that is biologically feasible. For example, in 1B last row, how would the blue neuron know about the coincident firing of red and orange? In this simple illustration, one might think that a simple spiking threshold could do the trick. Basically, only when two neurons are activated at the same time, blue reaches the threshold. However, I am not sure if this interpretation can be generalized easily. If so, it would make strong predictions about spiking thresholds etc. Anyway, it would be nice if the authors could go more into the details how neurons can implement the multiplicative synapses (The authors state in the Methods „[…] each neuron in Gd needs to keep track of coincident firing of any neuron in Ge with any other in Gf.“ And I think it would be nice to have a clearer picture of how this can be achieved in biological systems)

  We agree with the reviewer that not enough detail was given on the biological implementation of multiplicative synapses and related predictions. To address this, we rewrote the explanation in Section 3 of the results, which we repeat here for convenience:

With this factorization we demonstrate how any given multiplicative interaction of state-variables can be accurately implemented through multiplicative synapses between neurons (Fig.1B+C~bottom, Methods~7.4). The matrix Ω_{md} then represents d-th degree multiplicative interactions between cells. In particular, Ω_{m2} represents the connectivity required for each cell to multiply each pair of their inputs (Fig. 1B+C bottom row), with the synapse essentially acting as a coincidence detector. Higher order synapses would behave similarly but with three or more coincident spikes (Supp. Fig. S1). While these higher order interactions are unlikely to be biologically feasible, lower order multiplicative interactions may indeed be possible in biology, and have been hypothesized before (Koch and Poggio,1992), as we will discuss further in the discussion. In principle, increasingly complex nonlinear dynamics may be implemented through the inclusion of higher-order terms in eq. (5) (Ω_{md} for d>2), though this flexibility comes with increased cost on the number of synapses and neural interactions. In a later section we will show how to avoid interactions beyond pair-wise synapses, and we will derive the expected number of connections for each type of synapse.

Compared to linear dynamics, multiplicative synapses enable nonlinear computations such as AND gates (Fig. 1C bottom). Overall, the above derivation demonstrates that the presence of multiplicative synapses arises naturally in mSCNs from extending the spike coding framework to polynomial dynamical systems.

- 4.2: To prove that higher-order multiplicative synapses are not strictly necessary, you stack two networks together that can compute third-order multiplications with pairwise interactions.
    - I think this would make strong predictions about the brain's connectivity structure that might be visible in one way or another. Is there evidence for such a hypothesis?
    - In the example you show, there is no difference between triplet and no triplet suggesting that both have the very same performance. However, I guess there are examples where the triplet (or even higher-order multiplicative synapse) is necessary to achieve a reasonable performance. It would be nice if this could be discussed more so that the reader gets a better understanding of the limitations.

We are intrigued by the reviewer's suggestion that stacking networks together for higher-order multiplications could lead to specific predictions. While we feel a detailed exploration of these points is out of the scope of the paper, we attempt to address both of them at a high level with a new paragraph in the discussion, again repeated here for convenience:

Finally, while higher order multiplicative interactions are increasingly difficult to implement biologically, we demonstrated that by stacking networks with lower-order interactions one can achieve the same computations. This makes the concrete prediction that the connectivity between areas should be of the same dimensionality *as the signal being transferred. This does indeed seem to be the case to some degree, with the communication between some areas being low-dimensional (Semedo et al., 2019). There is also a likely limit to how many networks can be effectively stacked in this way to perform higher-order interactions, as each stacked network introduces a delay.*

- Section 5 needs clarification: This section is very important but is not explained well (I feel it should be extended significantly).
  - Also, Fig. 4 is not completely clear. Dashed lines are derived from theory!? Refer to the eq. In the methods section. Solid lines are mean derived from the simulations?
  - It would be nice to show performances (maybe for the example shown before) in the case when no one-to-one connectivity is used. This would make it easier for the reader to evaluate how well the approach works even in more biologically plausible regimes.
  - The authors use phrases like „signal dimensions“ or „dynamical system density“ for the very first time in section 5. Even though it becomes clear over time, it would be nice, if these terms could be introduced properly.

  We agree that this section of the paper suffered from a lack of clarity, and have now comprehensively reworded and expanded it. In particular, we have moved some detailed parts from the methods section (as also suggested by Reviewer 2), which we think has improved the overall intelligibility. We have also added a new analysis showing how connectivity density relates to network performance (Supp.Fig.S2C).

Minor

- In abstract: „Synapse type“. For readers that are not familiar with SCN, this might be misleading. They could think of excitatory vs. inhibitory, and not fast, slow etc. as you have in mind
  We have reworded much of the abstract, among other parts of the text, to clarify this point.
- Page 3, (b): „iff“
  This is an abbreviation for if and only if -- we have kept it in the text, but can change it if the reviewer thinks it is too unclear.
- The last row in Fig. 1B is very nice, but to make sure that the reader does understand the general concept, an illustration beyond m=2 could be helpful.
  This is a great suggestion, and we have added a new supplementary figure (S1) illustrating higher-order interactions, and refer to it from the legend of Fig 1 and the main text (repeated here):
  Higher order synapses would behave similarly but with three or more coincident spikes (Supp. Fig. S1).
- Below Eq. 7 —> could you please briefly state how you ended up with equation (log2(g)-1) support networks
  A footnote has been added here to explain this.
- In Eq. (8): Is the Heaviside step function hidden in the sum (the exponentials should not have an impact before the spike was emitted) … sorry it is late and I am tired. ;)

We believe that it is not necessary to include a Heaviside function considering how we defined the sum (by writing t_ki<t).

- You introduce the Kronecker product on page 16 but it has been used throughout the Methods section. It would make sense to restructure the Methods part and introduce this important concept before.
  We have now given the Kronecker product its own methods subsection, and refer to it in the text.
- Method section 7.7 needs more clarification (and please also use equation numbers here). Especially, the equations E(F), E(S), and E(Q) need more explanations.
  We have clarified the equations and text, in addition to moving some of this content to the main text in section 5.
- Page 5: you reference to Fig. 1C, D —> there is no D
- Page 6: your reference to Fig. 2Aiii (see second paragraph, end) but mean 2Aiv
- Page 6: your reference to Fig. 2Biii (see third paragraph, middle) but mean ii
- Page 10 (before Discussion): You refer to FIg. 3 but mean Fig. 4
- Before Eq. (14) in the text: should be - D_i^T y = …)
- Before 7.3: You refer to Supp. Fig. 1 but you mean 2
- First line page 17: the
- Before Eq. 23: double the
  We thank the reviewer for catching these mistakes! They have all been fixed.

Review by anonymous reviewer, 2021-06-02 13:10

In this manuscript, the authors provide an important theoretical contribution to the understanding of spiking neural networks, by describing mathematically how to design neural networks to implement any type of polynomial dynamics. The study is based on previous work on the spike coding framework (Boerlin et al. 2013), initially based on generating linear dynamical systems, that has been then extended in several publications to account for different biological constraints and non-linear dynamics. In this study, the authors explore an alternative top-down extension of the framework to generate non-linear dynamics based on multiplicative synapses. After defining the mathematical framework, the study illustrates the application of the method to a standard benchmark: the chaotic dynamics of a Lorenz attractor. Next, they show an alternative implementation, based on a hierarchical multi-network architecture, that can implement polynomial dynamics using only pairwise multiplicative synapses. Finally, they consider how the top-down assumptions of the model relate to biological constraints in brain wiring, and compare their advantages and disadvantages with respect to other existing methods. This effort in the discussion is necessary, since how is the main premise of this theoretical work -multiplicative synapses- remains an open question.

This work is quite technical, and builds on previous theoretical work. Nevertheless, the authors make an effort to explain previous findings and make this study accessible to readers familiar with standard methods in computational neuroscience methods, but not necessarily with the spike coding framework. The findings of the study constitute a solid advancement in the understanding of computations in spiking networks, and open new paths for potential impactful applications in the field.

Comments

- Introduction
  In the second sentence of the first paragraph, it remains unclear to me why the implementation of non-linear dynamics through the recurrent connectivity, leads to ("accordingly") low-dimensional internal representations. I believe this statement makes sense based on the assumption that the implemented non-linear dynamics are low-dimensional. If this is the case, I would suggest to specify it. Otherwise, these two ideas (non-linear computation implemented by recurrent wiring and low-dimensional representations) should be presented separately.

  This is a very good point, which we agree was not clear in the original version. We have now rewritten the sentence as follows:

  One promising hypothesis is that networks represent relatively low-dimensional signals (compared to network size) (Cunningham and Yu, 2014; Keemink and Machens, 2019), and in this lower-dimensional space implement nonlinear dynamical systems through recurrent connectivity (Eliasmith,2005; Mante et al., 2013; Sussillo, 2014; Abbott et al., 2016; Barak, 2017; Mastrogiuseppeand Ostojic, 2018).

- 2 Spike coding networks
  It seems that there are missing references to some of the panels in Fig 1 in the main text (check as well for other figures). Fig 1A for example is not referred to anywhere. It could be included in Section 2.1.

  We now cite Fig 1A at the beginning of Section 2.1. Furthermore we double-checked other figure references to ensure that all were cited accordingly in the main text.

- In Section 2.2, I suggest to give a more precise intuition about why and what is "fast" in fast connections and what is "slow" in slow connections. Although this terminology is commonly used in the spike coding framework, it would help understanding the difference between Fig 1B top and middle, which is currently not commented in the main text.
  We thank the reviewer for this suggestion. This has now been clarified as follows:

  where so-called "fast" connections $\Omega_f = -D^T D$ keep the error constrained on a short time-scale (Fig. 1B+C top row), and "slow" connections $\Omega_s = D^T(A+\lambda I)D$ implement the dynamical computation using the filtered spikes r (Fig. 1B+C middle row). Here "fast" and "slow" refer to the rise-time of the synaptic PSPs (Fig. 1B). While technically an approximation, this implementation works well in practice and can closely reproduce a given linear dynamical system (Fig. 1B+C middle row)

- If I understood correctly, the middle row in Fig 1B+C has both fast and slow synapses.
  This is correct. We clarify this at the end of the figure legend as:
  For both linear and nonlinear dynamics the fast synapses are also required.

- Also, there is no reference to Fig1C middle in section 2.2.
  This has been fixed: we now reference figure 1C properly in Sections 2.2 and 3.

- 3 Nonlinear dynamics
  In Section 2.3, I suggest highlighting the fact that the Lorenz attractor is based on pairwise multiplications of the state-variables, since this concept is useful for the rest of the paper.
  Good point! We now clarify that the multiplications in the Lorenz attractor are pairwise:
  Notably, this system contains pairwise multiplicative terms of the state-variables, thereby making it a polynomial (and nonlinear) dynamical system.

- 4 Higher-order polynomials
  Reference to the Supp. Fig. 2 needs more context. Since Supp. Fig 2 does not show any combination of networks, it is worth explaining in a sentence why it is relevant here.
  To address this we adjusted the end of Section 4.1:
  We illustrate the resulting network and its inputs and outputs in Supp.~Fig.~S3.

  As well as the beginning of Section 4.2:
  Now, the combination of a standard mSCN (eq. 5) with a network which calculates the square of its inputs (Supp. Fig. S3), results in a system with the ability of computing third-order multiplications with only pairwise (second-order) synapses (Fig. 3A). Notably, one network computes the pairwise multiplications, and the other computes the desired third-order dynamic equation using another pairwise multiplication of the squared network output ($x \otimes x$) and x.

- In Section 4.3 and corresponding Fig. 3, it remains unclear whether the double pendulum implementation corresponds to the approximated dynamics (where sin theta = theta) for all lines, or whether this approximation is used only for the spike coding networks.
  All implementations of the double pendulum use the approximated dynamics. We have now clarified this in the text as follows:
  We implemented the first-order approximated double pendulum system in three distinct ways.

- 5 On the number of required connections
  It is currently hard to understand this section and the content of Fig. 6, just by reading the caption and the main text, without looking at the methods. Some sentences would require more explanations (“The fast connections only depend on the density of the decoder, and any two neurons are connected whenever they share a decoding-dimension”). I suggest moving part of the Methods 7.7 into the main text.
  We agree with the reviewer that this section specifically lacked sufficient explanation, as the other reviewer also pointed out. We have significantly updated and expanded this section, and hope that it is now clear enough.

- Discussion, Section 6.2

Reference to Fig. 3 should be Fig. 4, I believe.
This has been fixed.

- Methods
I suggest including Section 7.3 before Section 7.2. The Kronecker product is first explained in Section 7.3, although already used in Section 7.2. Explain what Section 7.2 wants to calculate before starting calculating (introduce matrix W, for example). included before S the Methods. In
We have now given the Kronecker product its own section, which we hope addresses this point.

Misc errors

- I believe matrix D should be in bold in Eq. 3
- The colors of the two neurons in Figure 1 look to me red and orange, instead of red and green as indicated in the caption.
- Section 3.1: Reference to Fig. 2Aiii at the end of the second paragraph should actually be Fig. 2A iv ?
- Section 3.1: the network simulation still display the… -> the network simulation still displays the
- Section 3.1: (Fig. 2top) -> (Fig. 2 top)
- Section 3.2: high pair-wise connectivity -> dense pair-wise connectivity? high-density pair-wise connectivity?
- Methods 7.3: “among the other properties, the one that will be used is…”-> (suggestion) We use the mixed-product property
- Methods 7.5: nonliarity -> non-linearity or nonlinearity
- Methods 7.7.2: the second part-> the second term.
- So the interesting part is… -> We focus on the first term …
We thank the reviewer for the thorough reading of the text. All of these mistakes have been corrected according to the reviewer's suggestions.