Costs and benefits of using rhythmic rate codes

Erik J. Peterson[1,2,*] and Bradley Voytek[3,4,5,6],

1 Department of Psychology,
2 Center for the Neural Basis of Cognition, Carnegie Mellon University, Pittsburgh, PA, USA
3 Department of Cognitive Science,
4 Halıcıoğlu Data Science Institute,
5 Neurosciences Graduate Program, and
6 Kavli Institute for Brain & Mind, University of California, San Diego, CA, USA

* Erik.Exists@gmail.com

Abstract

Neural oscillations are one of the most well known macroscopic phenomena observed in the nervous system, and the benefits of this kind of coding have been the topic of frequent analysis. Many of these studies focused on communication between populations which were already oscillating, and sought to understand how synchrony and communication interact. In this paper, take an alternative approach. We focus on measuring the costs, and benefits, of moving to an from an aperiodic code to a rhythmic one. We utilize a Linear-Nonlinear Poisson model, and assume a rate code. We report that no one factor seems to predict the costs, or benefits, of translating into a rhythmic code. Instead the synaptic connection type, strength, population size, and stimulus and oscillation firing rates interact in a variety of nonlinear ways. We suggest a number of experiments that might be used to confirm these predictions.

Author summary

It's good to oscillate, sometimes.   ₁

Introduction   ₂

In this paper we begin to reconcile the fact that oscillations are often found in the   ₃
nervous system, with the fact that very few aspects of the natural world are oscillatory. ₄
In particular, any theoretical cost(s) of achieving a rhythmic code from an aperiodic   ₅
sensory measurements is not well quantified at present.   ₆

   We presume there *is* a cost because moving from asynchronous sensory signals to   ₇
rhythmic codes is a translation, and translations are rarely free from error, or loss of   ₈
information. For any coding scheme, translation from one coding system to another   ₉
will often introduce artifacts. We are particularly interested in artifacts that come   ₁₀
from physiological properties, like those caused by the mode of synaptic transmission,   ₁₁
population size, firing rate, or the details of synaptic connectivity.   ₁₂

   This gap in our understanding regarding the information cost of oscillations exists   ₁₃
because a typical account assumes populations of interest are already oscillating. Several ₁₄

lines of data however suggest that real oscillations are nonstationary, and bursty [?, ?, 1]. 15
This leads us to hypothesize there is a "translation cost", which needs further study. 16

The benefits of an oscillatory code have been frequently remarked on [2–30]. Here we 17
argue these existing results must be weighed against any information losses, or translation 18
costs, that come from adopting a rhythm code in the first place. 19

In this paper we studied a Linear-Nonlinear-Poisson model (LNP), and measured 20
translation costs using information theory. We assumed 1. a rate code, 2. a single global 21
oscillator, 3. only fast synaptic transmission, and 4. a linear or "arithmetical" model of 22
modulation [31]. 23

We are surprised to report that no one physiological factor predicted the costs, or 24
benefits, we observed by translating to a rhythmic rate code. Instead, the relationships 25
are quite nonlinear, and interdependent. 26

Results 27

We simulated rate coding and oscillatory translation using a LNP system (Linear- 28
Nonlinear Poisson). Despite the oversimplifications present in LNP models, they have 29
proven remarkably useful in developing new theoretical ideas [32, 33]. 30

In our model, a time-varying asynchronous stimulus is modulated by an oscillation 31
and the result is nonlinearly filtered, and Poisson sampled. We studied four kinds of 32
modulation, M, each made using the same global oscillator. 33

1. Model E – additive modulation, mimicking direct excitatory entrainment by 34
   AMPAergic synapses 35
2. Model I – subtractive modulation, mimicking direct inhibitory entrainment by 36
   GABAergic synapses 37
3. Model EI – multiplicative modulation, mimicking the gain effect of balanced 38
   microcircuits [?]. 39
4. Model II – mimicked dual dendritic effects of inhibition, which can be both 40
   subtractive and divisive at the same time [34, 35]. 41

We measured conversion cost/benefit by making two information theory calculations. 42
We first measured the information shared between a time-varying reference stimulus s 43
(whose creation is described in the *Methods*) and a Poissonic sample of that stimulus 44
S. This quantifies the innate cost of translating a stimulus to a rate code by itself. In 45
other words, this is a control condition. This control is the basis for our information 46
measurements in Eq. 1, which we will now explain. 47

Let S be the Poisson rate time course for the control measure. And let it be given by 48
$\hat{S} \sim P(s)$ where s is the original target, and $P(.)$ is a Poisson distribution. We use r_S, 49
etc, to denote the firing rate of samples from P. This sampling to give S then gives us 50
our first mutual information value, $\mathrm{MI}(s, S)$. For the second we generated a "translated" 51
timecourse M, using one of the models above (i.e., Model 1-4). With M we can measure 52
the mutual information between s and M, giving $\mathrm{MI}(s, M)$, and our final equation. 53

$$\Delta\mathrm{MI} = \mathrm{MI}(s, M) - \mathrm{MI}(s, S) \tag{1}$$

Examples of s, S and all models M are shown in Fig. 1. The conversion from 54
timecourses to distribution, and a visual depiction of overlap between S and M is shown 55
in Fig 2 and Fig 3, respectability. 56

If Eq. 1 gives a negative value, this implies a translation cost. If it gives a positive 57
value, this implies a translation benefit. 58

Fig 1. Examples of rhythmic transformations. The left panels **f**-**j** are power spectra derived from the spiking timecourses on the left. **a.** Example of a typical driving stimulus. **b.** A Poisson sampling of the (unmodulated) control stimulus, which serves as the non-oscillatory control condition for studying encoding between the stimulus and neural response, under a rate-coding regime **c.** The oscillation (grey) and a Poisson sampling of the same (black). This is denoted as model **O**. The remaining panels (*d-g*) show examples of the different kinds of transformation we studied here [31], where the Poisson rate coding system is modulated by different forms of oscillation. Note that all series were normalized between 0-1 in this figure to simplify the comparisons. *Simulation parameters*: the average stimulus firing rate r_s was set to 10 spks/s, the oscillation average rate r_o was 2 spks/s, g was 2, and the population size was 40. The oscillation's oscillation frequency was 6 Hz.

Fig 2. Step 1: estimating mutual information. We used a binning procedure to estimate discrete probability distributions, with m bins. **a**. Example of binning, using reference stimulus from Figure 2. The bins are represented as grey lines. **b**. The probability distribution derived from panel a.

Costs and benefits

We'll now review the translation benefits and costs we observed. The variables of foremost interest were the stimulus firing rate, the oscillation firing rate, the synaptic strength g, and the size of the target population. Other variables which proved less important are discussed later on.

The simple system we studied requires significant caveats to interpret and we'll return to these in the discussion. Despite the limits imposed by such simplicity, the cost/benefit of translation looks to be quite nuanced. "Small" populations of fewer than approximately 100 cells lead to better outcomes, overall. For excitatory modulations (E/EI) translation benefit would be predicted to peak below 100 cells (Figure **??b**). For inhibitory modulations (I/II) the predicted translation cost is minimized below 100 cells (Figure **??b**. There is however a counterexample to this, shown in Figure 4**c**).)

Owing to the importance of population size, we will focus for the remainder of this paper on a "small" population of 40 cells, and "large" population of 10240 cells.

In the small population Model E (additive translation), we report an average benefit of about 0.1 bits over the control (Figure 7**a**). This decreases with stimulus firing rate (Figure 7**b**) and increases with both the oscillation firing rate (Figure 7**b**) and the synaptic strength g (Figure 7**,** all panels).

In the small population with Model EI (multiplicative translation) we report an average benefit of about 0.25 bits (Figure 7**a**). This scaled with both the stimulus and oscillation firing rates (Figure 7**b-d**), until the effect plateaued. As expected for any model of balanced modulation we saw an inverse relationship between synaptic strength and mutual information [**?**]. For weak synapses and strong oscillatory firing we observed a decrease in benefit once the oscillation's firing rate exceeded the stimulus rate. This is 10 spks/s in Figure 7.

With both kinds of inhibitory modulation I and II we report a translation cost, in most conditions. This had an average value of about -0.1 bits (Figure 4)**a**; Figure 7)). The cost was to some degree mitigated by the stimulus strength (Figure 7)**b**, and made worse by the oscillation's firing rate and synaptic strength. In model II, the ratio of divisive to subtractive inhibition did not affect the overall cost. Extensive divisive inhibition did however prevent strong stimulus rates from mitigating the translation cost (Figure 4**b**).

Average, best, and worst case analysis

We have so far characterized rhythmic encoding by assuming the average case is the most relevant. And from the point of view of statistical inference that is often so. But from the point of view of computer science it is often fruitful to consider two other cases:

Fig 3. Step 2: estimating mutual information. Light grey is the stimulus distribution. Dark grey is the Poisson sample distribution. Overlap between the two light and dark provides a visual estimate of their mutual information (MI). That is the more overlap, the higher their MI. **a.** An example distribution of the control model, S. **b-d.** Example distributions for all the transform/modulations. These distributions correspond to the timecourses shown in Figure 1.

Fig 4. The importance of population size. **a)** Total change in mutual information (ΔMI), across all populations. Individual points represent individual trials. Black bar is the average. **e)** The average change in mutual information (ΔMI) plotted as a function of population size. Error bars represent standard deviation. *Simulation parameters*: the average stimulus firing rate r_s was set to 10 spks/s, the average oscillation rate r_o was 2 spks/s, g ranged from 1-8 in one unit increments. The oscillation frequency was 6 Hz.

Fig 5. The effect of firing rate on *small* population sizes. Top panels explore stimulus firing rate. Bottom panels explore the effect of the oscillation's firing rate. **a),c** Total change in mutual information (ΔMI), across all conditions as a function of firing rate. Individual points represent individual trials. Black bar is the average. **b),d** The average change in mutual information with oscillation's firing rate and relative synaptic strength g (blue). Error bars represent standard deviation. *Simulation parameters*: the average stimulus firing rate r_s ranged from 5-30 spks/s, the oscillation's average firing rate r_o ranged from 2-18 spks/s, g ranged from 1-8. The oscillation frequency was 6 Hz.

Population size: 40

Fig 6. Stimulus firing rate and its effect on oscillatory encoding in Model **II**. **a)** Total change in mutual information (ΔMI). Individual points represent individual trials. Black bar is the average. **e)** Average change in mutual information (ΔMI) plotted as a function of stimulus firing firing rate, and relative synaptic strength g (blue) *Simulation parameters*: the average stimulus firing rate r_s ranged from 5-30 spks/s, the oscillation's average firing rate r_o was 2 spks/s, g ranged from 1-8. The oscillation frequency was 6 Hz.

Population size: 10240

Fig 7. The effect of firing rate on *large* population sizes. Top panels explore stimulus firing rate. Bottom panels explore the effect of the oscillation's firing rate. **a),c** Total change in mutual information (ΔMI), across all conditions. Individual points represent individual trials. **b),d** The average change in mutual information with firing rate and relative synaptic strength g (blue). Error bars represent standard deviation. *Simulation parameters*: the average stimulus firing rate r_s ranged from 5-30 spks/s, the oscillation's average firing rate r_o ranged from 2-18 spks/s, g ranged from 1-8. The oscillation frequency was 6 Hz.

the worst case outcome and the best case outcome. This is because sometimes errors can come with substantial outside consequences. In such cases a worst case disposition is appropriate. While in other situations errors can be inconsequential, and so the focus is on the best case.

When we consider these three different "dispositions"–average, best, and worst–the results we can report change radically. In Figure 8**a-b** we re-plot the average case from Fig. 4**b**. This time averaging over synaptic strengths as well. We then compare this to both the min Figure 8**c-d** and max (Figure 8**e-f**) taken from the same experiment. In the worst case oscillations are a net cost across all conditions. In the best case they are a net gain across all conditions. However these best and worst cases are not uniform across oscillation physiological conditions.

We are aware of no scientific work looking at the error "disposition" of neural communications. Given how severe it changes out results, and given how errors can matter little in play or substantially during precision tasks like surgery, we argue this kind of analysis is relevant, and under-studied.

Control experiments

Oscillation frequency. We have studied a "slow" 6 Hz oscillation so far. When we examined frequency across a typical physiological range, we observed few changes to our results (Figure 9**a**). This comes from our choice of rate models, and linear translation. We would expect a very different result for a biophysical model, or one that considers spike-timing, rather than just rate coding. In both cases a faster oscillation should often synchronize more readily then slow oscillation [**?**].

Background noise (size). We have simulated a background noise level of about 2 spikes/s for all previous simulations, which is consistent with most cortical recordings Varying this more than 10 fold did not affect the benefits or costs we observed (Figure 9**c**).

Bin size. Our calculation of mutual information relies on binned data, and our first choice for the bin number (m) was to have it match the average stimulus firing rate. This was about 10 spikes/s. This makes our entropy's have a consistent 1 spks/s resolution. Doubling or halving our bin number had no qualitative effect on the results we observed (Figure 9**d**). This seems to validate our choice.

Methods

The methods for this work are divided into three parts. The simulations, the information theory calculations.

Simulations

Each model was composed of four parts. The driving stimulus, the oscillation, the target firing population, and the background noise.

To generate the stimulus we simulated a one dimensional diffusion processes, as those have been suggested to act as a reasonable approximation to the "naturalistic" firing patterns observed in early visual areas during passive movie viewing [**?**, 36]. This is given by Equation 2, where $N(0, 1)$ is a *sample* from the normal distribution which gets rescaled by $\sigma_s = 0.01 r_s$, where r_s is the average drive. The stimulus is further constrained such that $r_s^t \geq 2$. Note: that when we refer to value that is a function of time (*i.e.*, $r_s(t)$) we use the shorthand, r_s^t.

In general we use use lower case $\hat{r}$ to refer to individual samples, r for average values, and upper case letters, S, M, E, etc to refer or timecourse of values or to refer to the model classes themselves; their being synonymous for our purposes

Population size: 40

Fig 8. Average, best, and worst case analysis. **a-b)** Total change in mutual information in the *average case. a.*For each model, summed over population size and synaptic strength. *b.* As a function of population size. **c-d)** Total change in mutual information in the *worst case. c.* For each model, summed over population size and synaptic strength. *d.* As a function of population size. **e-f)** Total change in mutual information in the *best case. e.* For each model, summed over population size and synaptic strength. *f.* As a function of population size. *Simulation parameters*: the average stimulus firing rate r_s was set to 10 spks/s, the oscillation's average firing rate r_o was 2 spks/s, g ranged from 1-8. The oscillation frequency was 6 Hz.

Population size: 40

a.

b.

c.

Fig 9. a. Oscillation frequency and its own effect on oscillatory encoding. **b**. Background population size and its effect on oscillatory encoding. **c**. Bin size and its effect on oscillatory encoding.

Oscillatory firing was sinusoidal, and given by Equation 3. Where f is the periodicity in Hz. r_o was fixed at 2 Hz, unless so noted. The third drive is, as noted, the background which is a constant $r_b^t = 2$.

$$r_s^t = r_s^{t-1} + \sigma_s * N(0,1) \tag{2}$$

$$r_o^t = r_o^{t-1} + (\frac{r_o^{t-1}}{2}) \sin 2\pi f t \tag{3}$$

The oscillation rate r_o^t was modulated (or not) and results were independently sampled by N independent Poisson neurons. Prior to Poisson sampling, all firing drives were nonlinearly transformed by a threshold-linear activation function [31, 37]. The combination of linear interactions, nonlinearly filtering, and Poisson sampling is what marks our model as an LNP system [?].

Firing in the target population was based on a Poisson process. This is given by Equation 4, where r is the firing rate, or drive, and k is the number of spikes per time step (1 ms for all models here). Firing was limited to be no faster than once every 2 milliseconds, the absolute factory period [?].

For each trial in every experiment a new random seed was chosen. This generated a new naturalistic stimulus, as well as a new population of Poisson neurons. Every experiment consisted of 100 trials. Before analysis all spike trains for individual Poisson neurons were transformed to population average firing rate, denoted by a hat symbol. For example, $\hat{r}_s^t$ is the instantaneous Poisson firing rate from the stimulus drive r_s^t. The temporal resolution used was 1 ms.

$$P(r) = \frac{r^k}{k!}e^{-r} \tag{4}$$

We assume that "synapses" for the stimulus into the target population were fixed, and unitary. This means the equations for modulation are as follows. The oscillation drive r_o^t modulates the stimulus r_s^t with a variable "synaptic strength" g. Modulation comes in one of four kinds – additive (Eq 5, subtractive (Eq 6, multiplicative (Eq 7, and subtractive/divisive (Eq 8. As a shorthand we refer to these as **E**, **I**, **EI**, and **II** respectively. Examples of each are shown in Fig. 1.

$$r_{\mathrm{E}}^t = r_s^t + g r_o^t \tag{5}$$

$$r_{\mathrm{I}}^t = r_s^t - g r_o^t \tag{6}$$

$$r_{\mathrm{EI}}^t = r_s^t g r_o^t \tag{7}$$

$$r_{\mathrm{II}}^t = \frac{r_s^t - (1-q)\ g r_o^t}{q g r_o^t} \tag{8}$$

Our first model of inhibition **I** is perhaps too simple, given what is presently understood about the complexity of cortical inhibition [34, 38–40]. This is because inhibition in the cortex, our prime area of interest, is targeted to both the dendrites and the soma. With recent work suggesting the soma experiences divisive inhibition at the soma, while dendrite targeting interneurons act subtractively. Papasavvas *et al* [35] developed a rate model of this two part inhibition scheme, and we adapted it to fit our LNP system. Here we implement a linear trade-off between the amount of divisive and subtractive inhibition, controlled by parameter q (Eq. 8).

Fig 10. Neural arithmetic, as neural modulation.

Examples of all arithmetic codes are shown in Figure 10 where they are presented in terms of the classic FI-curve.

Parameters and definitions for all equations, both the model and the metrics below, are summarized in Table **??**

Mutual information

To calculate entropy and mutual information we took a discretization approach, in particular a binning approach. This worked as follows. First all signals were normalized between $[0, 1]$. The $[0, 1]$ range was then divided into m bins. We set the bin number m to be equal to the average stimulus drive, $\bar{r}_s$. That is 10 spks/s, which is the same as 1 Hz. See Figures 2-3 for an example. Our rationale was that 1 Hz is a reasonable first pass estimate for how precise a downstream decoding population can be. Doubling or halving m did not qualitatively change our findings (Figure 9).

After discretization we constructed histograms separately from the reference signal and target firing. Normalizing these histograms by the total count generated discrete probability distributions from which we can calculate entropy H and mutual information MI (Equations 9-10). Examples of these distributions can be found in Fig. 3.

$$\mathrm{H}_{\hat{X}} = \sum_{i=1}^{m} = p_i \, \log_{10} \, p_i \tag{9}$$

$$\Delta\mathrm{MI} = \mathrm{MI}(s, M) - \mathrm{MI}(s, S) \tag{10}$$

To measure translation cost/benefit we made two of these information theory calculations. We first measure the information shared between a time-varying stimulus s (whose creation is described in the *Methods*) and a Poissonic firing-rate sample of that stimulus $\hat{r}_S \sim P(r_s s)$, where r_s is some average firing rate. That is, $\mathrm{MI}(s, \hat{r}_S)$. Second we measured the information shared between the stimulus s and a sample of the *modulated* stimulus. That is, $\hat{r}_M \sim P(s \circledast g o r_o)$ where o is the time-varying oscillator (see *Methods*), g represents the fixed "synaptic" strength of the modulation, r_o is the average firing rate, and $\circledast$ represents one of the four modulations outlined above. This gives, $\mathrm{MI}(s, \hat{r}_M)$. For notational purposes we replace the M the $\hat{r}_M$ with the notation for modulation itself. For example, in model E we would denote the M to read as r_E. The final cost/benefit measure is given by Eq. 1 (*Results*).

$$\Delta\mathrm{MI} = \mathrm{MI}_{\mathrm{S}, \hat{s}} - \mathrm{MI}_{\mathrm{S}, \hat{X}} \tag{11}$$

Term	Description	Value or formula
B	the background drive	-
r_b	the average firing rate of B	2
S	the stimulus drive	-
r_s	the average firing rate of s	$2 - 30$ (10)
σ_r	variance of s	$0.01 r_s$
$\hat{S}$	Poisson firing; a sample of S	Eq. 4
O	the oscillating drive	-
r_o	the average firing rate of O	$2 - 8$ (2)
f	oscillation frequency	$4 - 30$ (6)
$\hat{O}$	Poisson firing; a sample of O	Eq. 4
g	modulation strength	$1 - 8$
E	additive modulation (a drive)	Eq. 5
I	subtractive modulation (a drive)	Eq. 6
EI	multiplicative modulation (a drive)	Eq. 7
II	subtractive/division modulation (a drive)	Eq. 8
$\hat{E}$	Poisson firing; a sample of E	Eq. 4
$\hat{I}$	Poisson firing; a sample of I	Eq. 4
$\hat{E}I$	Poisson firing; a sample of EI	Eq. 4
$\hat{I}I$	Poisson firing; a sample of II	Eq. 4
m	Bin size	r_s
H	Shannon entropy	Eq. 9
MI	mutual information	Eq. 10
PAC	phase amplitude coupling	Eq. **??**
$N(0,1)$	a normal sample	-
τ_α	the alpha kernel	TODO

Table 1. Model parameters, and general mathematical terms. Parameter ranges are listed first with a typical value, when appropriate, shown in parenthesis.

The central idea behind our use of ΔMI is that if oscillatory modulation is truly beneficial it should do better, that is have higher MI, than just sampling r_s alone.

Discussion

We have made a beginning of a theory to understand benefits and costs of oscillatory translation, or coding, in a rate setting. Our LNP model is simple, and costs not a common consideration. As a result our discussion takes the form of some general questions, and answers. We conclude by making several experimental predictions.

So physiological details matter, a lot?

The results we find depend on factors hidden in a typical electrophysiology experiment. The population size, for example.

A debate, which has been happening since oscillations were first noted can be colloquially summarized by the question, "do oscillations do anything, or are they a byproduct of some other process?". In some ways this question is settled. We have clear examples of oscillations seeming to cause neural pathology [41, 42]. But for much of the literature the link to a direct functional role rests on correlative findings. So in some ways the basic question posed above is still an important one.

This basic question is why we are so interested in costs here. Oscillations arise naturally in most any system of equations which has additive, subtractive, or balanced interactions [9, 34, 38, 39]. It is natural to wonder if evolution has taken these general mathematical phenomena to harness. If so, then costs are a central issue because to take advantage of rhythms, their use would develop in accordance with their limits.

What are the central limitations of this model?

In this paper, we do not capture any one neural circuit. This needs to be addressed in future work. We instead tried to capture the broader phenomenon of oscillatory translation, in the abstract. We have some faith in our predictions because other LNP models have born out in the past [?].

Why would we use a rate model?

Why would we use a rate model when biophysical models are available? Rate models are used to fit a very large range of electrophysiology data [?] and in some cases to pre-test and validate clinical results [?]. Epilepsy is a well known example [?], but there are others [?]. Despite their simplicity then, or perhaps because of it, rate models used in the clinic and for other interventions require that they capture something real in the data.

We have not so far discussed more complex attempts to answer the same basic question we posed. This is because in the few examples there are, biophysical models have required *very specific* conditions to see benefits [19]. But we argue one the most interesting things about demonstrating oscillations is that they are ubiquitous across circuits, and even animal phyla. We argue that this ubiquity argues for simplicity in mechanism, not high precision.

Even if we are wrong about oscillation's mechanism as rooting in simplicity, then at the very least the simple view was worth considering first, before developing overly-complex or complicated theories about what oscillations might do for neural coding.

Why not Wilson-Cowan?

Many of the more popular rate models, like the Wilson-Cowan equations, are dynamical systems models. They show far more complex patterns than our system. But all of this complexity boils down to addition, subtraction, and a nonlinear transform at their synapses. Our model can be thought of as a zoomed in view, a localized approximation, of just the interactions between individual populations.

Are costs always costs? Are benefits always benefits?

The answer to these questions is a definitive no. The cost/benefit labels we have been using follow from our assumption that oscillations aid communications. We could have adopted the opposite assumption, which is just as supported by theory and experiments. That is, oscillations can gate out information, or otherwise suppress communications [43–45]. So from this view what we have been calling a cost (-ΔMI) is not, but can be seen as an optimal result. And what we have been calling a benefit must be seen as a kind of cost, or error.

Inhibitory oscillations are "bad"?

Over the widest range of parameters, excitatory oscillations are of consistent benefit. So why would the nervous system use anything else? We think the answer to this question might have three parts.

First, a loss of information might not mean a loss of function. For example, we can think of inhibitory oscillations as a *sampling* [?], or a *compression algorithm*? In this view the "cost", or loss of information, simplifies the message.

Secondly we know mathematically that inhibition interactions form dynamical systems are very robust to perturbation [24, 26, 46]. Perhaps then they are nature's way to create time-stable activity, that can spread over spatial distances. This, in turn, may lay the groundwork for forming stable "cognitive maps", that are robust to perturbations, just like the kind hypothesized to happen in the hippocampus and its inhibitory theta rhythm.

Thirdly, simulations of hippocampus have also shown that an inhibitory code is also stable in terms of intrinsic homeostasis [?]. This work predicted that one cannot use excitatory instead, their tonic excitement causes a homeostatic response that ends up suppressing all possible input. This, in effect, can silence the cell completely.

Transformation, not transmission

We assumed oscillations are a communications channel. This view of oscillations is common today, so we are justified in the work of others [33, 47]. However from another view the "goal" of neurons is not broadly speaking to transmit information, but to compute with it [?, 48–52]. If computation is a rhythms' primary role, then the analysis in this paper does not seem to apply.

What are the new predictions?

We assume in all cases the ability to make dual recordings of spiking at two locations in the central nervous system, and that the connection type(s) between these locations are known. For example, AMPAergic (E) or GABAergic (I) connections. With the ability to make dual recordings, any experimentalist can construct rate-based information theory measures, like those we studied in this paper.

Our predictions are qualitative, and seek to establish a basic concordance between our model and a real nervous system, to be measured. We assume also that passive

recording of cortical areas can generate sufficient variability to test these predictions. 287
We predict, 288

1. A rise and then decline in $|\Delta MI|$ as the size of the entrained population increases 289
   from a 100 cells, to 10000 or more. This is predicted to occur regardless of 290
   connection type. 291

2. Direct (monosynaptic) inhibitory I oscillatory coding should uniformly show nega- 292
   tive ΔMI, with the exception of large (ξ10,000) populations subjected to a strong 293
   driving stimulus. 294

3. Direct (monosynaptic) excitatory E oscillatory coding should uniformly show 295
   positive ΔMI. 296

4. Oscillations that affect both E and I connections, that are balanced [?], lead to 297
   qualitative distinct cost/benefits when compared to E or I direct connections. 298
   Namely, 299

5. Increases to the stimulus drive increase ΔMI for EI connections, but 300

6. Increases to the stimulus drive decrease ΔMI for E connections. 301

7. For large populations, increases to the amplitude of the oscillation will shift ΔMI 302
   from a positive value for in EI connections, to a negative value. 303

8. But this will not happen for E connections. 304

9. When a population of cells experiences both dendritic and somatic inhibition, 305
   relative increases to divisive/somatic component will further decrease ΔMI, and 306

10. this last effect will be amplified as the strength of the driving stimulus rises. 307

11. Increases or decreases to the background activity will not affect ΔMI. 308

Summary 309

We have made the beginnings of a cost-benefit theory for oscillatory codes, and shown 310
how the effect of oscillations can be quite intricate, when measured in information 311
theoretic terms. A fully realized version of this theory would let us predict when and 312
where oscillation would aid or harm communication, and by exactly how much. 313

References

1. Cole S, Voytek B. Cycle-by-Cycle Analysis of Neural Oscillations. Journal of Neurophysiology. 2019;122(2):849–861.

2. Siegel M, Donner TH, Engel AK. Spectral Fingerprints of Large-Scale Neuronal Interactions. Nature Reviews Neuroscience. 2012;13(2):121–134. doi:10.1038/nrn3137.

3. Buehlmann A, Deco G. Optimal Information Transfer in the Cortex through Synchronization. PLoS Comput Biol. 2010;6(9):e1000934. doi:10.1371/journal.pcbi.1000934.

4. von Nicolai C, Engler G, Sharott A, Engel AK, Moll CK, Siegel M. Corticostriatal Coordination through Coherent Phase-Amplitude Coupling. Journal of Neuroscience. 2014;34(17):5938–5948. doi:10.1523/JNEUROSCI.5007-13.2014.

5. Voytek B, Knight RT. Dynamic Network Communication as a Unifying Neural Basis for Cognition, Development, Aging, and Disease. Biological Psychiatry. 2015;77(12):1089–1097. doi:10.1016/j.biopsych.2015.04.016.

6. Womelsdorf T, Valiante TA, Sahin NT, Miller KJ, Tiesinga P. Dynamic Circuit Motifs Underlying Rhythmic Gain Control, Gating and Integration. Nature Neuroscience. 2014;17(8):1031–1039. doi:10.1038/nn.3764.

7. Mejias JF, Murray JD, Kennedy H, Wang XJ. Feedforward and Feedback Frequency-Dependent Interactions in a Large-Scale Laminar Network of the Primate Cortex. Science Advances. 2016;2(11):e1601335.

8. Sohal VS. How Close Are We to Understanding What (If Anything) Oscillations Do in Cortical Circuits? Journal of Neuroscience. 2016;36(41):10489–10495. doi:10.1523/JNEUROSCI.0990-16.2016.

9. Peterson EJ, Voytek B. Alpha Oscillations Control Cortical Gain by Modulating Excitatory-Inhibitory Background Activity. bioRxiv. 2017;185074. doi:10.1101/185074.

10. Fries P. Rhythms for Cognition: Communication through Coherence. Neuron. 2015;88(1):220–235. doi:10.1016/j.neuron.2015.09.034.

11. Brunet N, Vinck M, Bosman CA, Singer W, Fries P. Gamma or No Gamma, That Is the Question. Trends in Cognitive Sciences. 2014;18(10):507–509. doi:10.1016/j.tics.2014.08.006.

12. Saalmann YB, Pinsk MA, Wang L, Li X, Kastner S. The Pulvinar Regulates Information Transmission Between Cortical Areas Based on Attention Demands. Science. 2012;337(6095):753–756. doi:10.1126/science.1223082.

13. Helfrich RF, Fiebelkorn IC, Szczepanski SM, Lin JJ, Parvizi J, Knight RT, et al. Neural Mechanisms of Sustained Attention Are Rhythmic. Neuron. 2018;99(4):854–865.e5. doi:10.1016/j.neuron.2018.07.032.

14. Fukunaga I, Herb JT, Kollo M, Boyden ES, Schaefer AT. Independent Control of Gamma and Theta Activity by Distinct Interneuron Networks in the Olfactory Bulb. Nature Neuroscience. 2014;17(9):1208–1216. doi:10.1038/nn.3760.

15. Roach JP, Pidde A, Katz E, Wu J, Ognjanovski N, Aton SJ, et al. Resonance with Subthreshold Oscillatory Drive Organizes Activity and Optimizes Learning in Neural Networks. Proc Natl Acad Sci USA. 2018;115(13):E3017–E3025. doi:10.1073/pnas.1716933115.

16. Tort ABL, Komorowski R, Eichenbaum H, Kopell N. Measuring Phase-Amplitude Coupling Between Neuronal Oscillations of Different Frequencies. Journal of Neurophysiology. 2010;104(2):1195–1210. doi:10.1152/jn.00106.2010.

17. Onslow ACE, Jones MW, Bogacz R. A Canonical Circuit for Generating Phase-Amplitude Coupling. PLoS ONE. 2014;9(8):e102591. doi:10.1371/journal.pone.0102591.

18. Jensen O. Oscillations in the Alpha Band (9-12 Hz) Increase with Memory Load during Retention in a Short-Term Memory Task. Cerebral Cortex. 2002;12(8):877–882. doi:10.1093/cercor/12.8.877.

19. Tiesinga PHE, Fellous JM, José JV, Sejnowski TJ. Optimal Information Transfer in Synchronized Neocortical Neurons. Neurocomputing. 2001;38-40:397–402. doi:10.1016/S0925-2312(01)00464-7.

20. Nyhus E, Curran T. Functional Role of Gamma and Theta Oscillations in Episodic Memory. Neuroscience & Biobehavioral Reviews. 2010;34(7):1023–1035. doi:10.1016/j.neubiorev.2009.12.014.

21. Lisman JE, Jensen O. The Theta-Gamma Neural Code. Neuron. 2013;77(6):1002–1016. doi:10.1016/j.neuron.2013.03.007.

22. Jensen O, Lisman JE. An Oscillatory Short-Term Memory Buffer Model Can Account for Data on the Sternberg Task. The Journal of Neuroscience. 1998;18(24):10688–10699. doi:10.1523/JNEUROSCI.18-24-10688.1998.

23. Mizuseki K, Buzsaki G. Theta Oscillations Decrease Spike Synchrony in the Hippocampus and Entorhinal Cortex. Philosophical Transactions of the Royal Society B: Biological Sciences. 2013;369(1635):20120530–20120530. doi:10.1098/rstb.2012.0530.

24. Wang XJ, Buzsáki G. Gamma Oscillation by Synaptic Inhibition in a Hippocampal Interneuronal Network Model. The Journal of Neuroscience. 1996;16(20):6402–6413. doi:10.1523/JNEUROSCI.16-20-06402.1996.

25. Wang HP, Spencer D, Fellous JM, Sejnowski TJ. Synchrony of Thalamocortical Inputs Maximizes Cortical Reliability. Science. 2010;328(5974):106–109. doi:10.1126/science.1183108.

26. Wang XJ. Neurophysiological and Computational Principles of Cortical Rhythms in Cognition. Physiological Reviews. 2010;90(3):1195–1268. doi:10.1152/physrev.00035.2008.

27. Lundqvist M, Rose J, Herman P, Brincat SL, Buschman TJ, Miller EK. Gamma and Beta Bursts Underlie Working Memory. Neuron. 2016;90(1):152–164. doi:10.1016/j.neuron.2016.02.028.

28. Miller EK, Buschman TJ. Cortical Circuits for the Control of Attention. Current Opinion in Neurobiology. 2013;23(2):216–222. doi:10.1016/j.conb.2012.11.011.

29. Sherman MA, Lee S, Law R, Haegens S, Thorn CA, Hämäläinen MS, et al. Neural Mechanisms of Transient Neocortical Beta Rhythms: Converging Evidence from Humans, Computational Modeling, Monkeys, and Mice. Proceedings of the National Academy of Sciences. 2016;113(33):E4885–E4894. doi:10.1073/pnas.1604135113.

30. Haegens S, Nacher V, Luna R, Romo R, Jensen O. -Oscillations in the Monkey Sensorimotor Network Influence Discrimination Performance by Rhythmical Inhibition of Neuronal Spiking. Proceedings of the National Academy of Sciences. 2011;108(48):19377–19382. doi:10.1073/pnas.1117190108.

31. Silver RA. Neuronal Arithmetic. Nat Rev Neurosci. 2010;11(7):474–489. doi:10.1038/nrn2864.

32. Gerstner W, Kistler W, Naud R, Paninski L. Neuronal Dynamics: From Single Neurons to Networks and Models of Cognition. vol. 1. Cambridge University Press; 2014.

33. Dayan P, Abbott LF. Theoretical Neuroscience. MIT press; 2005.

34. Jadi M, Polsky A, Schiller J, Mel BW. Location-Dependent Effects of Inhibition on Local Spiking in Pyramidal Neuron Dendrites. PLoS Computational Biology. 2012;8(6):e1002550. doi:10.1371/journal.pcbi.1002550.

35. Papasavvas CA, Wang Y, Trevelyan AJ, Kaiser M. Gain Control through Divisive Inhibition Prevents Abrupt Transition to Chaos in a Neural Mass Model. Physical Review E. 2015;92(3). doi:10.1103/PhysRevE.92.032723.

36. Barbieri F, Mazzoni A, Logothetis NK, Panzeri S, Brunel N. Stimulus Dependence of Local Field Potential Spectra: Experiment versus Theory. Journal of Neuroscience. 2014;34(44):14589–14605. doi:10.1523/JNEUROSCI.5365-13.2014.

37. Curto C, Degeratu A, Itskov V. Encoding Binary Neural Codes in Networks of Threshold-Linear Neurons. Neural Computation. 2013;25(11):2858–2903. doi:10.1162/NECO$_a$0504.

38. Jadi MP, Sejnowski TJ. Cortical Oscillations Arise from Contextual Interactions That Regulate Sparse Coding. Proceedings of the National Academy of Sciences. 2014;111(18):6780–6785. doi:10.1073/pnas.1405300111.

39. Jadi MP, Sejnowski TJ. Regulating Cortical Oscillations in an Inhibition-Stabilized Network. Proceedings of the IEEE. 2014;102(5):830–842. doi:10.1109/JPROC.2014.2313113.

40. Mel BW, Schiller J. On the Fight Between Excitation and Inhibition: Location Is Everything. Science Signaling. 2004;2004(250):pe44–pe44. doi:10.1126/stke.2502004pe44.

41. Kühn AA, Doyle L, Pogosyan A, Yarrow K, Kupsch A, Schneider GH, et al. Modulation of Beta Oscillations in the Subthalamic Area during Motor Imagery in Parkinson's Disease. Brain. 2006;129(3):695–706. doi:10.1093/brain/awh715.

42. Tinkhauser G, Pogosyan A, Tan H, Herz DM, Kühn AA, Brown P. Beta Burst Dynamics in Parkinson's Disease OFF and ON Dopaminergic Medication. Brain. 2017;140(11):2968–2981. doi:10.1093/brain/awx252.

43. Bonnefond M, Jensen O. Alpha Oscillations Serve to Protect Working Memory Maintenance against Anticipated Distracters. Current Biology. 2012;22(20):1969–1974. doi:10.1016/j.cub.2012.08.029.

44. Vijayan S, Kopell NJ. Thalamic Model of Awake Alpha Oscillations and Implications for Stimulus Processing. Proceedings of the National Academy of Sciences. 2012;109(45):18553–18558. doi:10.1073/pnas.1215385109.

45. Davis ZW, Muller L, Martinez-Trujillo J, Sejnowski T, Reynolds JH. Spontaneous Travelling Cortical Waves Gate Perception in Behaving Primates. Nature. 2020;doi:10.1038/s41586-020-2802-y.

46. Curto C, Itskov V. Cell Groups Reveal Structure of Stimulus Space. PLoS Computational Biology. 2008;4(10):e1000205. doi:10.1371/journal.pcbi.1000205.

47. Serrano E, Nowotny T, Levi R, Smith BH, Huerta R. Gain Control Network Conditions in Early Sensory Coding. PLoS Computational Biology. 2013;9(7):e1003133. doi:10.1371/journal.pcbi.1003133.

48. Sohn H, Narain D, Meirhaeghe N, Jazayeri M. Bayesian Computation through Cortical Latent Dynamics. Neuron. 2019;103(5):934–947.e5. doi:10.1016/j.neuron.2019.06.012.

49. Enel P, Procyk E, Quilodran R, Dominey PF. Reservoir Computing Properties of Neural Dynamics in Prefrontal Cortex. PLOS Computational Biology. 2016;12(6):e1004967. doi:10.1371/journal.pcbi.1004967.

50. Lagorce X, Benosman R. STICK: Spike Time Interval Computational Kernel, a Framework for General Purpose Computation Using Neurons, Precise Timing, Delays, and Synchrony. Neural Computation. 2015;27(11):2261–2317. doi:10.1162/NECO$_{a_0}$0783.

51. Marcus G, Marblestone A, Dean T. The Atoms of Neural Computation. Science. 2014;346(6209):551–552. doi:10.1126/science.1261661.

52. Tripp B, Eliasmith C. Function Approximation in Inhibitory Networks. Neural Networks. 2016;77:95–106. doi:10.1016/j.neunet.2016.01.010.

Addendum

This work is supported by a Sloan Research Fellowship (FG-2015-66057), the Whitehall Foundation (2017-12-73), the National Science Foundation under grant BCS-1736028, and a Halıcıoğlu Data Science Institute Fellowship (to B.V.). The authors declare that they have no competing financial interests. Correspondence should be addressed to E.J.P. (email: erik.exists@gmail.com).