

Synergistic neural computation is greater downstream of recurrent connectivity in organotypic cortical cultures

AUTHORS

Samantha P. Sherrill^{1,4,*}, Nicholas M. Timme², John M. Beggs³, and Ehren L. Newman¹

AFFILIATIONS

¹ Department of Psychological and Brain Sciences & Program in Neuroscience, Indiana University Bloomington, Bloomington, IN 47405, USA

² Department of Psychology, Indiana University-Purdue University Indianapolis, Indianapolis, IN 46202, USA

³ Department of Physics & Program in Neuroscience, Indiana University Bloomington, Bloomington, IN 47405, USA

⁴ Lead Contact

*Correspondence: samfaber@indiana.edu (S.S.)

ABSTRACT

Cortical information processing requires synergistic integration of input. Understanding the determinants of synergistic integration—a form of computation—in cortical circuits is therefore a critical step in understanding the functional principles underlying cortical information processing. We established previously that synergistic integration varies directly with the strength of feedforward connectivity. What relationship recurrent and feedback connectivity have with synergistic integration remains unknown. To address this, we analyzed the spiking activity of hundreds of well-isolated neurons in organotypic cultures of mouse somatosensory cortex, recorded using a high-density 512-channel microelectrode array. We asked how empirically observed synergistic integration, quantified through partial information decomposition, varied with local functional network structure. Toward that end, local functional network structure was categorized into motifs with varying recurrent and feedback connectivity. We found that synergistic integration was elevated in motifs with greater recurrent connectivity and was decreased in motifs with greater feedback connectivity. These results indicate that the directionality of local connectivity, beyond feedforward connections, has distinct influences on neural computation. Specifically, more upstream recurrence predicts greater downstream computation, but more feedback predicts lesser computation.

KEYWORDS

Neural information processing; recurrent connectivity; feedback connectivity; in vitro networks; effective connectivity; cortex; information theory; synergy; neural computation

INTRODUCTION

Feedforward, recurrent and feedback connections are important for information processing in both artificial and biological neural networks (Basheer & Hajmeer, 2000; Kriegeskorte, 2015). Whether these connections represent the strength of a synapse, or the amount of information transmission between two nodes, the directionality of these connections—*feedforward*, recurrent (*lateral*) or *feedback*—influences how the network processes information. A component of information processing that is central to both biological and artificial neural networks is their ability to perform synergistic integration, a form of computation. Feedforward connectivity has been previously shown to be a strong predictor of synergistic integration (Faber et al., 2019). However, the influence of recurrent and feedback connectivity on synergistic integration is unclear. Understanding how each of these connectivity types influences the computational properties of neural networks is a critical step in understanding how neural networks compute. Here, we examine this in the context of cortical networks, using a motif-style, information theoretic analysis of high-density *in vitro* recordings of spiking neurons.

Synergistic integration refers to the synergistic combination of existing information to derive new information which is “greater than the sum of the parts.” Thus, it is a proxy for a form of non-trivial computation. Synergistic integration can be measured as the synergy (Williams & Beer, 2011) that emerges when a given neuron integrates input from two other neurons (Timme et al., 2016; Wibral et al., 2017). This approach has been used effectively before (Faber et al., 2019; Sherrill et al., 2019; Timme et al., 2016; Wibral et al., 2017). Here, we leveraged this approach to determine how the amount of both recurrent and feedback connections relates to synergistic integration.

Recurrent connections are believed to implement memory processes (e.g. recollection, recognition) due to their generation of attractor-like, pattern completion activity (Douglas et al., 1995; Douglas & Martin, 2007; Hopfield, 1982; Leutgeb et al., 2005; Neunuebel & Knierim, 2014; Rolls, 2007; Rolls, 2013; Tang et al., 2018; Treves et al., 1997). This type of activity involves the combination of diverse features to form representations, also contributing to the interpretation and categorization of representations (Brincat & Connor, 2006; Carlson et al., 2013; Cichy et al., 2014; Clarke et al., 2015; Freiwald & Tsao, 2010; Sugase et al., 1999; Tang et al., 2014). These studies also show that greater interpretability of images and object categories occurs at latencies beyond those of known feedforward connections. Relatedly, in artificial neural networks, recurrent connections serve to expand computational power by extending operations in time, requiring a smaller network to carry out the same operations as a larger, purely feedforward network (Kriegeskorte, 2015; Spoerer et al., 2017). Thus, controlling for the size of a network, the use of recurrent connections over some feedforward connections can improve network operation.

Feedback connections are believed to implement top-down, goal-driven attention and perception, which involves the preferential activation of lower level neurons by higher level neurons (e.g., Boly et al., 2011; Kwon et al., 2016; Manita et al., 2015; for reviews, see Gilbert &

Li., 2013 and Sikkens et al., 2019). Due to its top-down nature, feedback connectivity also plays a role in the gating and rerouting of information flow, as well as error prediction (for related review see Clark, 2013; Bastos et al., 2012; Gilbert & Sigman, 2007; Grace, 2000; Lillicrap et al., 2016). Relatedly, feedback is associated with increased surround suppression, reducing the range of stimuli to which lower-level neurons respond (Nassi et al., 2013; Nurminen et al., 2018). From this perspective, feedback reduces the variance with which lower-level neurons can account for variance in higher-level neurons.

Here, we tested how recurrent and feedback connections relate to synergistic integration in cortical networks (Figure 1). To do this, we analyzed the spiking activity of hundreds of simultaneously recorded neurons from each of 25 organotypic cultures of mouse somatosensory cortex. We found that motifs with more recurrent connections had greater synergy compared to motifs without recurrent connections, and compared to what might be expected by chance. We also found that motifs with more feedback connections had less synergy compared to motifs without feedback connections, and compared to what might be expected by chance but that this negative relationship was accounted for by concurrent shifts in feedforward connectivity.

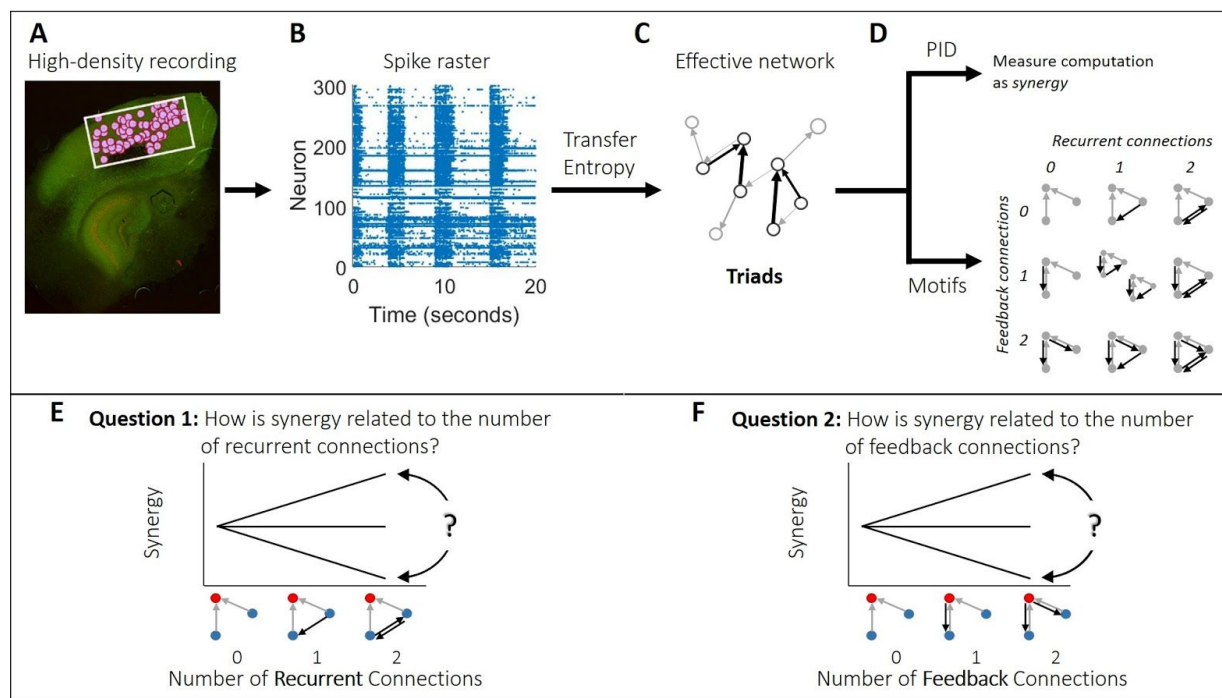


Figure 1. Methodological approach taken to ask how synergy is related to the number of recurrent and feedback connections in organotypic cultures of mouse cortex. **(A)** Hour-long recordings of spiking activity were collected *in vitro* from organotypic cultures of mouse somatosensory cortex using a high density 512-channel multielectrode array. **(B)** Spike sorting yielded spike trains of hundreds of well-isolated individual neurons per recording. **(C)** Effective connectivity between neurons was determined by quantifying transfer entropy between spike trains. The resulting effective networks were analyzed to identify all triads consisting of two

effective connections to a common receiver. **(D)** For each triad, we quantified the amount of synergy via partial information decomposition (PID). We also identified all possible triad-based motifs, and arranged them according to the number of recurrent and feedback connections they contained. **(E-F)** We sought to answer two questions. **(E)** Is synergy positively related, negatively related, or unrelated to the number of recurrent connections? **(F)** Is synergy positively related, negatively related, or unrelated to the number of feedback connections? Triads consist of two transmitter neurons (blue), each sending input (gray arrows) to a receiver neuron (red). Black arrows depict recurrent and feedback connections on the left and right, respectively. In this study, we ask how the number of recurrent and feedback connections relates to synergy.

RESULTS

We asked how the number of recurrent and feedback connections in motifs is related to computation by those motifs in cortical microcircuits by analyzing hour long recordings of spiking activity from organotypic cultures of mouse somatosensory cortex ($n = 25$), as summarized in Figure 1. Recordings contained between 98 and 594 well-isolated neurons (median = 310). We identified effective connections between neurons in each recording as those that had significant transfer entropy. We then identified all computational 3-node motifs. Computational motifs were those which included two transmitter nodes sending inputs to the same receiver-node. Motifs without this structure were excluded because we were only concerned with the motifs' ability to compute. The set of motifs included in our analyses are shown in Figure 2. We quantified the amount of computation performed by the receiver based on its inputs using 'synergy,' a term derived from partial information decomposition. Synergy was normalized to reflect the proportion of the receiving neuron entropy for which it accounted (pH^{rec}) and to control for variability across networks. Across triads, we asked whether synergy was positively or negatively related to the number of recurrent and feedback connections. This analysis was repeated at three timescales relevant to synaptic transmission, as determined by the granularity of the data binning and the delay between bins. All summary statistics are reported as medians or means followed by 95% bootstrap confidence intervals in brackets.




Figure 2. Set of computational 3-node motifs. Computational motifs were those in which both transmitters (blue dots) sent input (gray arrows) to the same receiver (red dots). Motifs were arranged in order (1-10) of the number of feedback and recurrent connections they contain (black arrows); either 0, 1 or 2.

Recurrence predicts increased synergy, feedback predicts decreased synergy

To examine the pattern of synergy across all 10 motifs, in all networks, we quantified the mean synergy for each motif type within each network (Figure 3). We then compared the mean synergy in recurrent motifs (those with more recurrent than feedback connections) to the mean synergy in feedback motifs (those with more feedback than recurrent connections). We observed significantly greater synergy in recurrent motifs (Fig 3, orange) compared to feedback motifs (Fig 3, green) (mean = 0.011 vs. 0.007, $z_{s,r.} = 6.31$, $n=75$, $p < 1 \times 10^{-9}$). To determine how recurrent and feedback connections affect synergy relative to baseline levels, we compared the observed synergy in each motif to the synergy in the default motif (with 0 recurrent and 0 feedback connections; Fig 3). We found that triads with recurrent motifs had significantly greater synergy than those with the default motif ($z_{s,r.} = 5.80$, $n=75$, $p < 1 \times 10^{-8}$). Conversely, triads with feedback motifs had significantly less synergy than those with the default motif ($z_{s,r.} = -2.61$, $n=75$, $p=0.009$). In addition to the comparison to baseline synergy, we compared the observed synergy values to those observed when motif labels were randomly permuted across triads. We observed the same qualitative pattern of results here as in the comparison to baseline synergy levels. We found that recurrent motifs had significantly greater synergy than expected by chance ($z_{s,r.} = 5.93$, $n=75$, $p < 1 \times 10^{-8}$). Conversely, feedback motifs had significantly less synergy than expected by chance ($z_{s,r.} = -3.96$, $n=75$, $p < 1 \times 10^{-4}$).

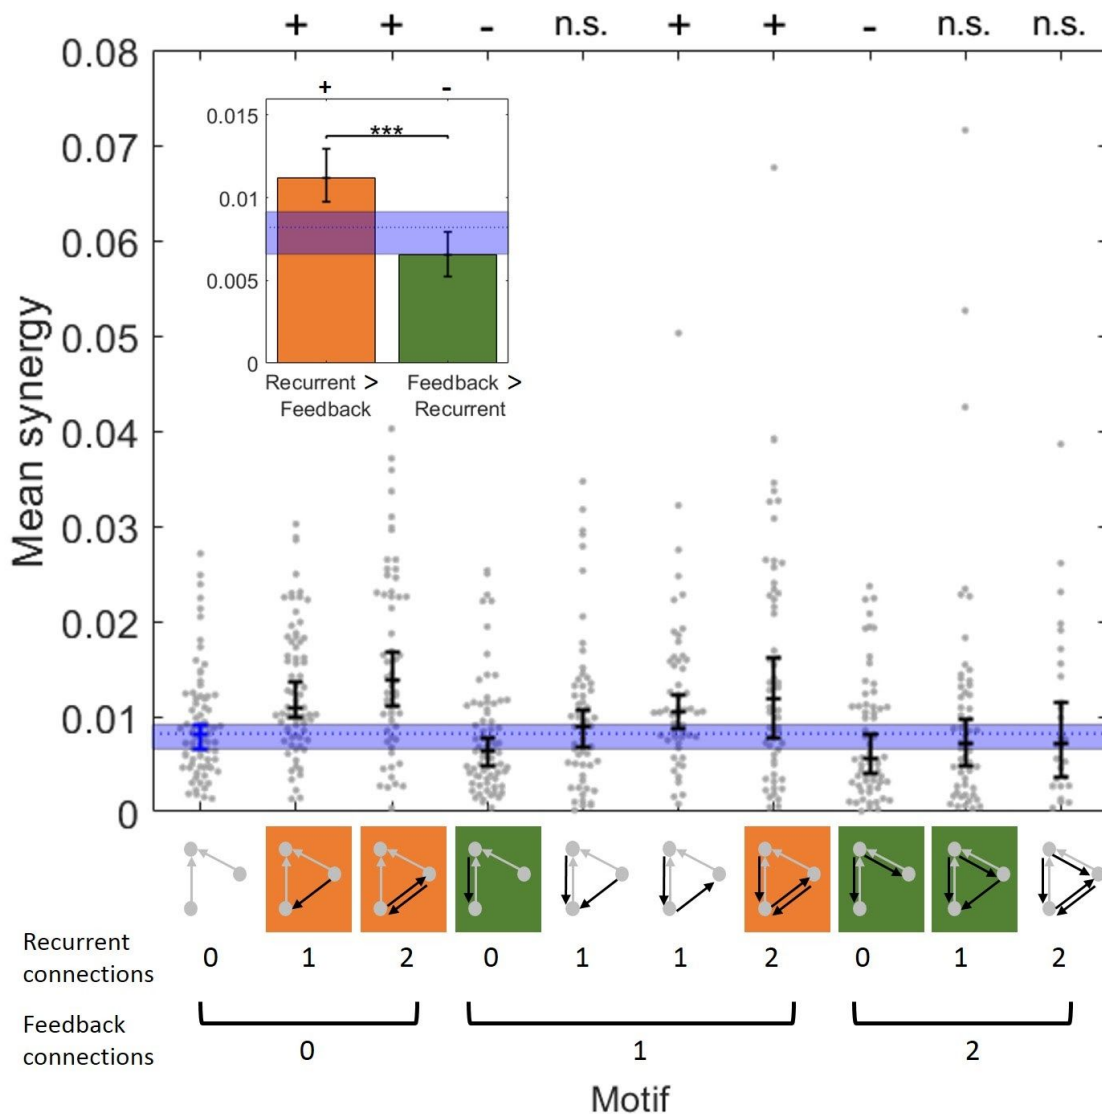


Figure 3. Synergistic integration is greater in recurrent motifs than in feedback motifs. Point clouds show the mean synergy value for each of the 75 networks analyzed for each type of motif. For distributions in which all networks did not exhibit the motif, $n < 75$. Central tendency and error bars depict the median and the 95% bootstrap confidence interval around the median. The motifs are graphically depicted below the x-axis and are organized by the number of feedback and recurrent connections. Motifs with more recurrent than feedback connections are indicated in orange. Motifs with more feedback than recurrent connections are indicated in green. Inset: Synergy values from motifs with greater recurrence or greater feedback connections were aggregated to directly compare the mean synergy. In both panels, the median (dotted line) and 95% bootstrap confidence interval (blue region) for baseline synergy values in default motifs (with 0 recurrent and 0 feedback connections) is shown. Significance indicators: '+' and '-' indicates $p < 0.01$ by a two-tailed test wherein '+' indicates significantly more than baseline and '-' indicates significantly less than baseline; *** $p < 1 \times 10^{-9}$.

To assess how synergy varies across the multiple levels of recurrence and feedback, we grouped the 10 motifs into 9 categories based on the number of feedback and recurrent connections they contain (Fig 4A). A two-factor ANOVA (recurrent vs. feedback), with three levels of each factor (0, 1 or 2 connections), was conducted to examine the main effects of recurrent and feedback connections on synergy and to test for an interaction effect. The main effect of recurrent connections was significant ($F(2,75)=17.53$, $p<0.0001$). The mean synergy increased as the number of recurrent connections increased (0.008 [0.007 0.009] vs. 0.012 [0.010 0.014] vs. 0.014 [0.012 0.017]; Fig 4B), reflected by a significant positive correlation between synergy and number of recurrent connections (Spearman $r = 0.25$, $n=675$, $p<1\times 10^{-8}$). The main effect of feedback connections was also significant ($F(2,75)=5.77$, $p=0.003$). The mean synergy decreased as the number of feedback connections increased (0.012 [0.011 0.014] vs. 0.011 [0.009 0.013] vs. 0.008 [0.007 0.010]; Fig 4C), reflected by a significant negative correlation between synergy and number of feedback connections (Spearman $r = -0.22$, $n=675$, $p<1\times 10^{-6}$). There was no significant interaction between the effects of recurrent and feedback connections on synergy ($F(2,75)=1.19$, $p=0.31$). The non-significant interaction effect between recurrence and feedback indicates that recurrence and feedback have predominantly independent effects on synergy. Taken together, these results show that, across these networks, motifs with greater upstream recurrence have greater synergy, and motifs with greater feedback have lesser synergy.

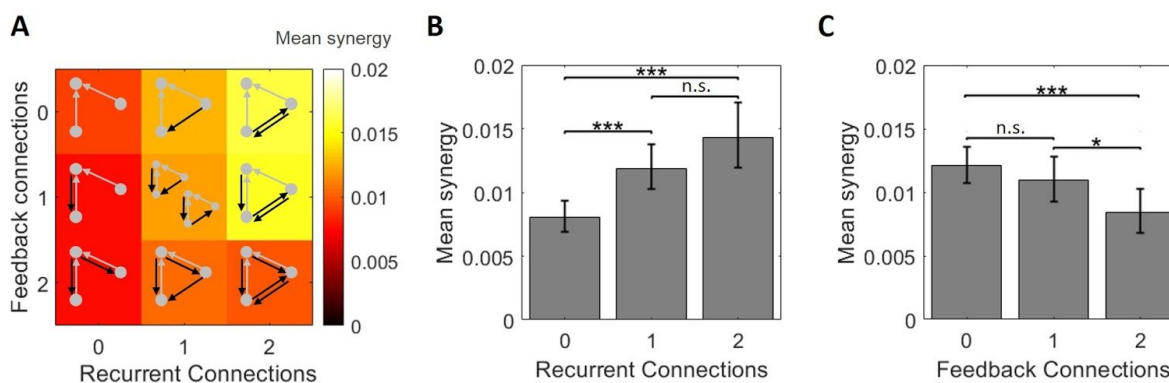


Figure 4. Synergy increases with greater recurrence and decreases with greater feedback. **(A)** Motifs are ordered based on the number of recurrent (columns) and feedback (rows) connections. The background heatmap, wherein brighter yellow reflects larger values, replots the medians shown in Figure 3. **(B)** Means of columns shown in A, plotted with errorbars computed across networks, show that synergy increases as the number of recurrent connections increases. **(C)** Means of rows shown in A, plotted with errorbars computed across networks, show that synergy decreases as the number of feedback connections increases. Significance indicators: * $p<0.05$, *** $p<1\times 10^{-3}$. These p-values were Bonferroni-Holm corrected for multiple comparisons. Errorbars are 95% bootstrap confidence intervals around the mean.

The strength of feedforward connections is a strong predictor of synergy (Faber et al. 2019). To address whether the relationship between synergy and either recurrent or feedback connectivity reported here is accounted for by the influence of the strength of feedforward connections, we

performed a control analysis wherein we regressed out the variance in synergy that was associated with feedforward connectivity strength. We then ran an ANOVA in which the residuals were the output variable and the number of recurrent and feedback connections were the predictor variables. This enabled us to determine if recurrence and/or feedback could account for variance in synergy after regressing out the effect of feedforward connection strengths on synergy. The ANOVA again revealed a significant main effect of the number of recurrent connections ($F(2,75)=12.09$, $p<0.0001$). This result was again supported by a significant positive correlation between residual synergy and number of recurrent connections (Spearman $r = 0.32$, $n=675$, $p<1\times 10^{-13}$). Thus, the number of recurrent connections between senders remained a significant predictor of synergy after controlling for the strength of feedforward connectivity. However, there was no significant main effect of the number of feedback connections after regressing out the strength of the feedforward connections ($F(2,75)=0.24$, $p=0.79$). This indicates that variance associated with the number of feedback connections was entangled with the variance in feedforward strength. Indeed, upon testing, we found that feedforward connectivity was negatively correlated with the number of feedback connections (Spearman $r = -0.22$, $n=675$, $p<1\times 10^{-6}$). The correlational nature of this analysis prohibits us from concluding whether the feedback connections reduced feedforward connectivity or vice versa. What can be said is that synergy is greater when the ratio of feedforward to feedback connectivity is larger. Thus, modulation of either can likely impact net synergy. Finally, as before, there was no significant interaction effect between recurrence and feedback in predicting the synergy residuals ($F(2,75)=1.19$, $p=0.31$). Taken together, these results show that the number of recurrent connections among upstream neurons contributes a novel, independent source of synergy beyond the variance accounted for by the strength of the feedforward connections.

Recurrent and feedback motifs are rare but overrepresented

To gain perspective as to how our findings regarding the influence of recurrent and feedback connectivity on synergy relate to network-wide processing, we asked how prevalent each type of connectivity was in our networks. To do this, we calculated the percentage of network-wide triads accounted for by each motif (Fig 5).

Consistent with the sparsity of these networks (average connection density: 1.14% [0.83% 1.54%]), the rate of incidence of each motif decreased rapidly as a function of the number of edges contained in the motif. The first motif, containing only 2 edges was most prevalent, accounting for 70.12% [67.15% 73.07%] of the computational 3-node triads. Motifs with 3 edges, whether recurrent or feedback, accounted for 23.94% [21.76% 26.14%] of the computational 3-node triads. This is significantly greater than the 2.12% [1.60% 2.83%] that would be expected by chance given random networks with the same sparsity ($t = 21.24$, $n = 75$, $p < 1\times 10^{-32}$). Motifs with 4, 5, and 6 edges were similarly over-represented from what would have been expected in random networks, but progressively decreased in prevalence (4 edge motifs: 4.93% [4.03% 5.92%] vs. 0.13% [0.06% 0.26%], $t = 10.27$, $n = 75$, $p < 1\times 10^{-15}$; 5 edge motifs: 0.66% [0.47% 0.88%] vs. 0.0039% [0.0008% 0.01%], $t = 6.39$, $n = 75$, $p < 1\times 10^{-7}$; 6 edge

motifs: 0.36% [0.12% 0.87%] vs. 0.0002% [0.0000% 0.0004%], $t = 1.99$, $n = 75$, $p = 0.051$. These results agree with findings in similar networks generated from the same data (Shimono & Beggs, 2015). These results are shown in Figure 5. Importantly, all motifs with recurrent and feedback edges, with the exception of the 6-edge motif, occurred more frequently than expected given network connection densities. Thus, the sparsity of our networks did not preclude our ability to detect recurrent and feedback motifs.

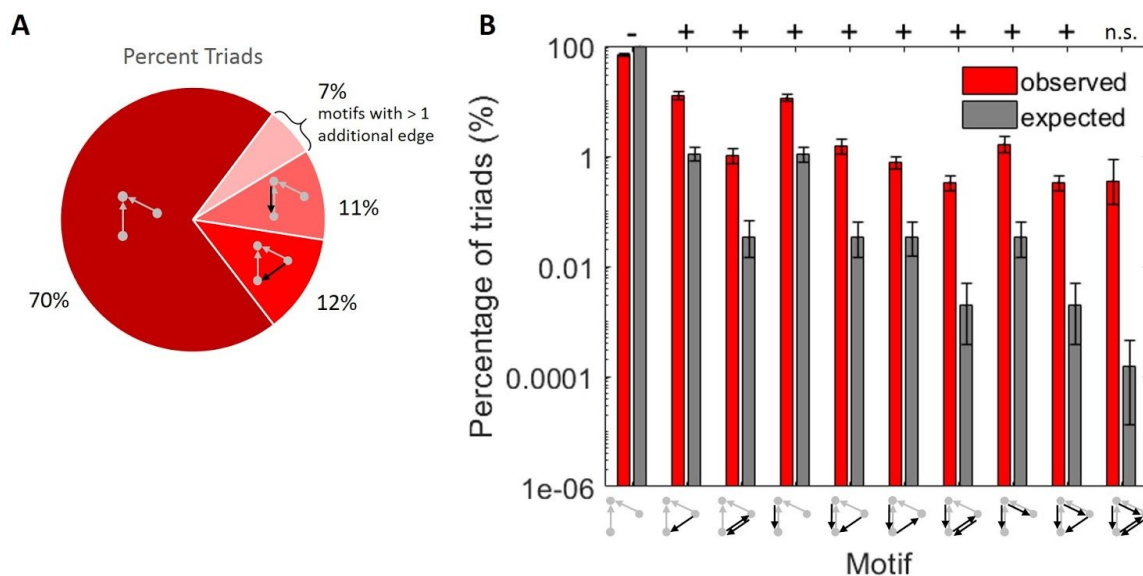


Figure 5. Recurrent and feedback motifs are rare, but occur more than expected given connection density. **(A)** Percent of network triads accounted for by each motif type. Motifs with greater connectivity are more rare. Values indicate means across all networks. **(B)** Log₁₀ scaled observed percentages of triads compared to expected percentages of triads per motif. Expected percentages obtained by raising the probability of observing a connection to the power of the number of connections in the motif, for each network. Significance indicators: '+' indicates significantly more than expected and '-' indicates significantly less than expected. For all significant values, $p < 1 \times 10^{-6}$.

To test for evidence of selection bias toward or away from triads with recurrent or feedback connectivity, we tested whether one type of connectivity was more or less prevalent among the triads containing a given number of edges. The null distribution would be an equal number of each. Among 3-edge computational motifs, the extra edge was recurrent in 50.7% [44.2% 57.1%] of the triads. This was not significantly different from 50% ($t = 0.22$, $n = 75$, $p = 0.83$). Likewise, across triads with 4 and 5 edges, containing differing numbers of recurrent and feedback connections, we found no evidence of bias toward one type of connectivity versus the other (4-edge motifs: 43.12% [35.12% 51.13%], $t = -1.67$, $n = 75$, $p = 0.10$; 5-edge motifs: 50% [50% 50%], $t = 0$, $n = 75$, $p = 1$; 6-edge motifs were not included in this as they contain the same number of feedback and recurrent connections by definition).

Finally, given the similar incidence of motifs containing recurrent and feedback edges, but significant differences in the synergy observed for each motif type, computational triads containing recurrent edges can be expected to account for a larger percentage of the network-wide synergy (Fig 6). Indeed, recurrent motifs comprised 13.79% [11.52% 16.28%] of triads and accounted for 20.43% [17.26% 23.85%] of network-wide synergy. Feedback motifs comprised 13.47% [11.41% 15.70%] of triads and only 10.12% [8.37% 12.11%] of network-wide synergy (Fig 6A, inset). Thus, although recurrent and feedback motifs accounted for similar percentages of network triads ($z_{s,r} = 0.09$, $n=75$, $p=0.92$), recurrent motifs accounted for a significantly higher percentage of network synergy than feedback motifs ($z_{s,r} = 4.18$, $n=75$, $p < 1 \times 10^{-4}$).

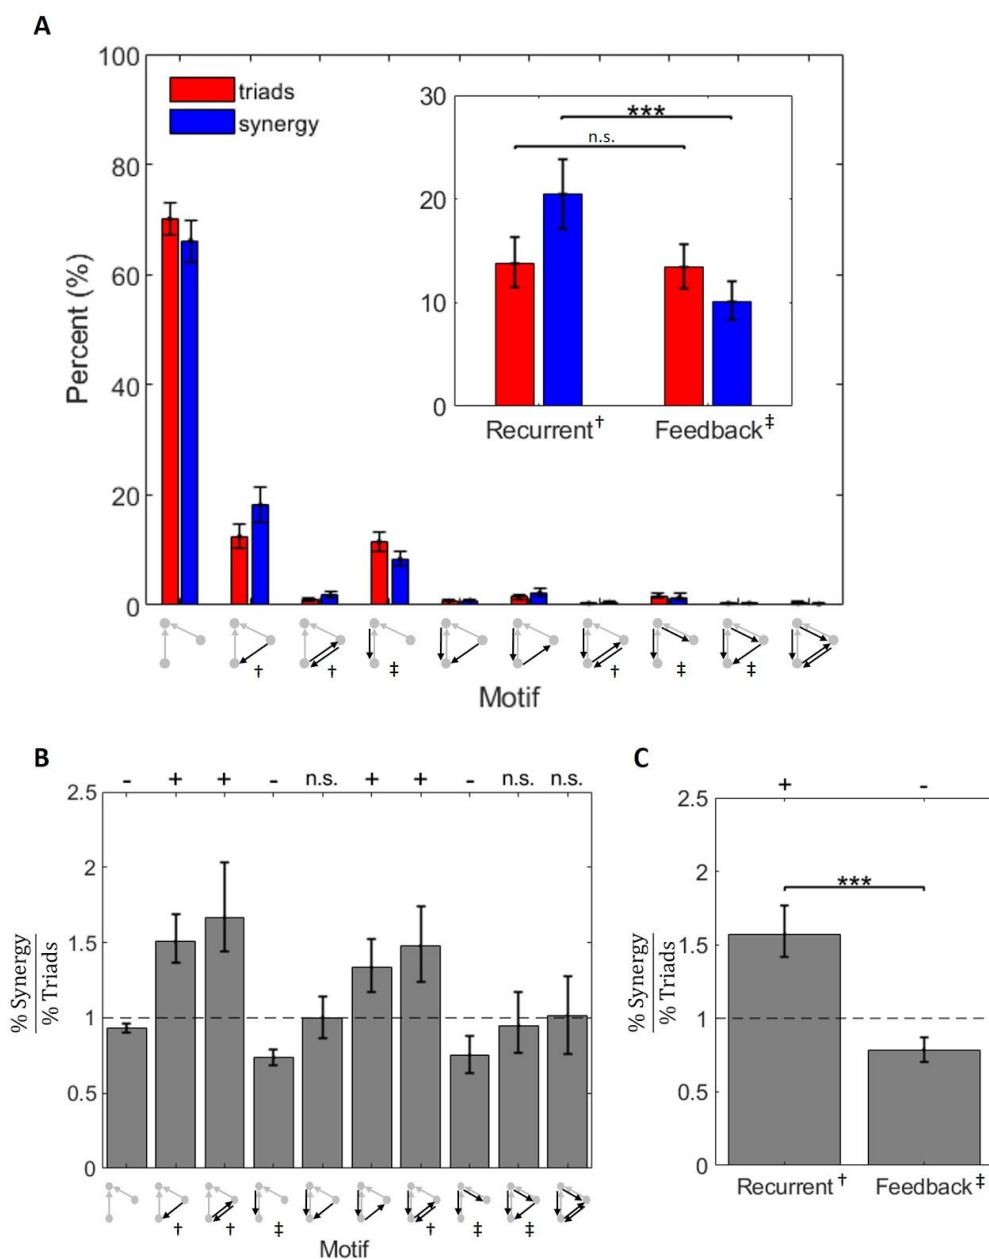


Figure 6. Recurrent and feedback motifs account for more and less network-wide synergy than expected, respectively. **(A)** Both recurrent motifs and feedback motifs are relatively rare, and they account for a relatively small proportion of overall synergy. Red bars from Fig 5B are replotted on a linear scale here for comparison. Inset: Recurrent motifs are as common as feedback motifs, but they account for significantly more synergy than feedback motifs. **(B)** The ratio of percent synergy to percent triads is shown per motif. Values above one indicate that the motif accounts for greater network-wide synergy than it does triads. Values less than one indicate that the motif accounts for less network-wide synergy than it does triads. **(C)** Recurrent motifs account for more synergy than expected given their frequency. Conversely, feedback motifs account for less synergy than expected given their frequency. Significance was determined by asking whether the distribution of ratios for each motif came from a distribution whose mean is equal to 1 (t-test). Significance indicators: '+' indicates significantly more than expected and '-' indicates significantly less than expected. For all significant values, $p < 0.001$. Central tendency shown in each figure is mean and error bars are 95% bootstrap confidence intervals around the mean. Mean was selected over median to ensure that percentages sum to 100. Significance indicator: *** $p < 0.001$. Motif indicator: † Recurrent, ‡ Feedback.

To determine whether motifs accounted for more synergy than expected given their frequency, we calculated the ratio of percent synergy to percent triads for each motif (Fig 6B-C). Values greater than one indicate that the motif accounts for more synergy than expected given its frequency. Values less than one indicate that the motif accounts for less synergy than expected given its frequency. We observed that recurrent motifs accounted for significantly greater network-wide synergy than expected given their frequency ($z_{s,r} = 6.28$, $n=75$, $p < 1 \times 10^{-9}$), and feedback motifs accounted for significantly less network-wide synergy than expected given their frequency ($z_{s,r} = -4.35$, $n=75$, $p < 1 \times 10^{-4}$; Fig 6C).

DISCUSSION

Understanding the relationship between specific connectivity types (feedforward, feedback, and upstream recurrence) for synergistic processing in cortical networks is essential for understanding how neural networks compute. We previously showed that synergistic processing varies directly with feedforward connectivity (Faber et al., 2019). Here, we examined the influence of recurrent and feedback connectivity on synergistic information processing in organotypic cortical cultures. Using information theoretic and network analyses of the spiking activity of hundreds of simultaneously recorded neurons from organotypic cultures of mouse somatosensory cortex, we showed for the first time that the number of recurrent and feedback connections in functional local microcircuits predicts the amount of synergy performed by those microcircuits. Specifically, we found that greater recurrence predicted greater synergy, but greater feedback predicted lesser synergy (Figure 7). Interestingly, the strength of feedforward connections, a covariate of synergy, explained the feedback-synergy relationship, but not the recurrence-synergy relationship. Thus, recurrence predicts synergistic processing above and beyond that predicted by the strength of inputs. Additionally, we found that, although recurrent motifs (those with more recurrent than feedback connections) were somewhat rare in our

networks--comprising 14% of all motifs--they account for 20% of the total network-wide synergy. Feedback motifs (those with more feedback than recurrent connections) were matched for prevalence with recurrent motifs--comprising 13% of all motifs--but only accounted for 10% of the total network-wide synergy. Thus, with similar prevalence, recurrent motifs accounted for twice as much synergy as feedback motifs.

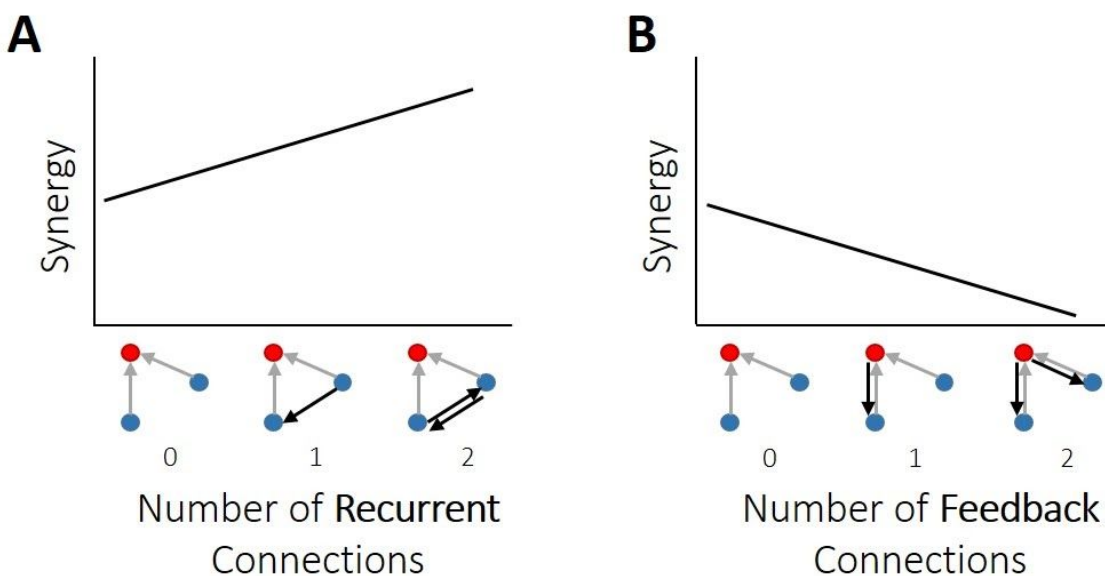


Figure 7. Summary of our findings regarding the relationships between recurrence and synergy and feedback and synergy. Synergy had a positive relationship with the number of recurrent connections and a negative relationship with the number of feedback connections. That is, synergy was elevated where there was greater upstream recurrence. Synergy was diminished where there was greater feedback.

Our finding that synergy increased with greater recurrence is consistent with previous work showing that recurrent connections are necessary for pattern completion tasks, both in biological (Douglas et al., 1995; Douglas & Martin, 2007; Leutgeb et al., 2005; Neunuebel & Knierim, 2014; Rolls, 2007; Tang et al., 2018; Treves et al., 1997) and artificial networks (Hopfield, 1982; Tang et al., 2018). Such tasks involve the integration of multiple, distinct features to generate a coherent representation, a process that involves some form of synergistic processing. Our finding that synergy decreased with greater feedback agrees with theoretical frameworks (Bastos et al., 2012; Clark, 2013; Gilbert & Sigman, 2007; Sikkens et al., 2019) and experimental studies (Bastos et al., 2015; Boly et al., 2011; Grace, 2000; Kwon et al., 2016; Manita et al., 2015) suggesting that feedback connections serve to reduce the variance with which lower-level neurons can account for variance in higher-level neurons, thereby reducing the strength of feedforward connectivity, and resulting in reduced synergy.

Our finding that increased recurrent connectivity corresponded to greater synergistic processing is also consistent with previous analyses of the topological determinants of synergistic processing in cortical cultures. For example, one such analysis found that synergistic

processing was directly related to the 'out-degree' of the upstream neurons (Timme et al., 2016). That is, the more neurons that a given upstream neuron made effective connections with, the greater the resulting synergy was in the recipient neurons. Similarly, we have previously shown that neurons in the rich clubs of cortical micro-circuits (i.e., highly-interconnected neurons) do about twice as much synergistic processing as neurons outside of the rich clubs (Faber et al., 2019). We have also shown that greater similarity (i.e. synchrony) of transmitters, such as might be generated by strong inter-connectivity, predicts greater synergy at synaptic timescales (Sherrill et al., 2019).

The strength of feedforward connectivity was an important consideration when analyzing the relationship between the number of recurrent/feedback connections and synergy. We have shown previously that the strength of feedforward connections is a strong, positive predictor of the amount of synergy (Faber et al., 2019). Here, we performed a control analysis where this relationship was first regressed out of the synergy values before asking whether recurrence and/or feedback connectivity were predictive of synergy. In our control analysis, we found that the number of feedback connections no longer accounted for a significant portion of the variance in synergy after accounting for the variance related to feedforward connectivity. This suggests that feedforward and feedback connectivity account for common variance in the resulting synergy. The positive relationship between recurrence and synergy, however, persisted after regressing out the influence of feedforward connectivity, suggesting that recurrence reflects a novel source of explanatory power over the generation of synergy. We hypothesize that this additional synergy emerges because recurrence increases the capacity of the transmitter neurons to jointly predict the behavior of the receiver, resulting in more synergy than if it just increased the amount of bivariate transfer entropy.

Network connection density was another important consideration in studying the influence of the amount of recurrent or feedback connectivity on synergy. Our networks were sparse, consistent with those observed in previous studies of biological neural networks (Hubel and Wiesel, 1959; Olshausen and Field, 2004; Mason et al., 1991; Markram et al., 1997; Thom & Palm, 2013). Thus, our results might have been skewed by the lack of connectivity, which would translate to a lack of observations for motifs with greater connectivity (i.e. recurrent and feedback motifs). We investigated the influence of sparsity on our results by asking how the expected frequency of motifs, given the probability of a single connection, compared to the frequency of motifs that we observed in our networks. We found that our networks had significantly more instances of both recurrent and feedback motifs than expected by chance. Thus, we concluded that the sparsity of our networks did not curtail our ability to observe these motifs. Moreover, the fact that recurrent and feedback motifs occurred more than expected by chance may indicate that such motifs, which evolve from network dynamics, are important for network processing.

The use of organotypic cultures in the present work facilitated the recording of hundreds of neurons simultaneously. While organotypic cultures naturally differ from intact *in vivo* tissue, organotypic cultures nonetheless exhibit synaptic structure and electrophysiological activity very similar to that found *in vivo* (Beggs & Plenz, 2004; Bolz et al., 1990; Caeser et al., 1989; Götz &

Bolz, 1992; Ikegaya et al., 2004; Klostermann & Wahle, 1999; Plenz & Aertsen, 1996). For example, the distribution of firing rates observed in cultures is lognormal, as seen in vivo (Nigam et al., 2016), and the strengths of functional connections are lognormally distributed, similar to the distribution of synaptic strengths observed in patch clamp recordings (reviewed in Buzsáki & Mizuseki, 2014; Song et al., 2005). These features indicate that organotypic cortical cultures serve as a reasonable model system for exploring local cortical networks, while offering unique accessibility to large neuron count and high temporal resolution recordings. However, additional work will need to be done to understand how the relationships between synergy and recurrence and synergy and feedback observed in vitro differ from what may exist in vivo, particularly in the context of behavior.

While stimulus-driven activity has been favored in research for its ability to provide insight into neural coding mechanisms, such studies assume that the brain is primarily reflexive and that internal dynamics are not informative with regard to information processing. However, internally-driven spontaneous activity of neurons, or activity that does not track external variables in observable ways, has been repeatedly shown to be no less cognitively interesting than stimulus-linked activity (Johnson et al., 2009; Raichle, 2010; for a review see Tozzi et al., 2016; Tsodyks et al., 1999). Not only is spontaneous activity predominant throughout the brain, but it also drives critical processes such as neuronal development (Cang et al., 2005; Chiappalone et al., 2006; Wibrál et al., 2017).

This work could inform future research on the importance of the topology of biological networks. The functional topology of biological neural networks has already been shown to influence neural information processing (Nigam et al., 2016; Timme et al., 2016; Faber et al., 2019). The present results add to our growing understanding of how the structure of neuronal interactions shapes neuronal behavior. These findings could also inform further research on artificial intelligence. Specifically, our results could be used in applied efforts to design engineered systems for the optimization of computational power and efficiency.

In summary, the present study demonstrates that, in *in vitro* local cortical networks, the number of upstream recurrent connections is positively related to the amount of downstream computation. And, the number of feedback connections from a downstream receiver to its upstream transmitters is negatively related to the amount of downstream computation. We also show that, although motifs with recurrent or feedback connections do not dominate the network, they account for more and less synergy than expected, respectively. These results agree with a number of previous studies arguing that network topology predicts neural information processing. Taken together, these findings provide increasing evidence of the influence of recurrence and feedback on neural information processing.

MATERIALS & METHODS

To answer the question of how computation is related to feedback and recurrence in cortical circuits, we combined network analysis with information theoretic tools to analyze the spiking

activity of hundreds of neurons recorded from organotypic cultures of mouse somatosensory cortex. Due to space limitations, here we provide an overview of our methods and focus on those steps that are most relevant for interpreting our results. A comprehensive description of all our methods can be found in the Supplemental Materials.

All procedures were performed in strict accordance with guidelines from the National Institutes of Health, and approved by the Animal Care and Use Committees of Indiana University and the University of California, Santa Cruz.

Electrophysiological recordings

All results reported here were derived from the analysis of electrophysiological recordings of 25 organotypic cultures prepared from slices of mouse somatosensory cortex. One hour long recordings were performed at 20 kHz sampling using a 512-channel array of 5 μm diameter electrodes arranged in a triangular lattice with an inter-electrode distance of 60 μm (spanning approximately 0.9 mm by 1.9 mm). Once the data were collected, spikes were sorted using a PCA approach (Ito et al., 2014; Litke et al., 2004; Timme et al., 2014) to form spike trains of between 98 and 594 (median = 310) well isolated individual neurons depending on the recording.

Network construction

Networks of effective connectivity, representing global activity in recordings, were constructed following the methods described by Timme et al. (2014, 2016). Briefly, weighted effective connections between neurons were established using transfer entropy (TE; Schreiber, 2000). To consider synaptic interactions, we computed TE at three timescales spanning 0.05 – 14 ms, discretized into overlapping bins of 0.05-3 ms, 1.6-6.4 ms, and 3.5-14 ms, resulting in 75 different networks. Only significant TE, determined through comparison to the TE values obtained with jittered spike trains ($\alpha = 0.001$; 5000 jitters), were used in the construction of the networks. TE values were normalized by the total entropy of the receiving neuron so as to reflect the proportion of the receiver neuron's capacity that can be accounted for by the transmitting neuron. Note, due to the sparse firing of our recordings, transfer entropy is biased towards detecting excitatory, rather than inhibitory, interactions. This is because transfer entropy grows with the probability of observing spike events. And in sparse spike time series it is statistically easier to detect an increase in the number of spikes (an excitatory effect) than it is to detect a decrease in the number of spikes (an inhibitory effect). Thus, here we assume connections are excitatory.

Identifying motifs

Computational motifs were identified using code inspired by the Matlab Brain Connectivity toolbox (Rubinov & Sporns, 2010). The code was written to categorize all computational triads—those in which two transmitters send edges to the same receiver node—according to the set of ten possible computational motifs, containing up to four additional edges. Because we were only interested in computational motifs, we did not consider the entire set of 3-node motifs. In addition, although motifs 5 and 6 (in this paper) would normally be considered conformationally equivalent, here they are distinct due to the consideration of transmitter and receiver node roles.

Quantifying computation

Computation was operationalized as synergy. Synergy measures the additional information regarding the future state of the receiver, gained by considering the prior state of the senders jointly, beyond what they offered individually, after accounting for the redundancy between the sending neurons and the past state of the receiver itself. Synergy was calculated according to the partial information decomposition (PID) approach described by Williams and Beer (2011), including use of the I_{min} term to calculate redundancy (see Supplemental Material). PID compares the measured bivariate TE between neurons $TE(J \rightarrow I)$ and $TE(K \rightarrow I)$ with the measured multivariate TE (the triad-level information transmission) among neurons $TE(\{J, K\} \rightarrow I)$ to estimate terms that reflect the unique information carried by each neuron, the redundancy between neurons, and the synergy (i.e., gain over the sum of the parts) between neurons. Redundancy was computed as per Supplemental equations 8-10. Synergy was then computed via:

$$\begin{aligned} Synergy(\{J, K\} \rightarrow I) = & TE(\{J, K\} \rightarrow I) - TE(J \rightarrow I) - TE(K \rightarrow I) \\ & + Redundancy(\{J, K\} \rightarrow I) \end{aligned} \quad (2)$$

Although there are other methods for calculating partial information terms (Bertschinger et al., 2014; Lizier et al., 2018; Pica et al., 2017; Wibral et al., 2017), we chose this measure because it is capable of detecting linear and nonlinear interactions and it has been shown to be effective for our datatype (Timme et al., 2016; Faber et al., 2019). In addition, unlike other methods (Lizier et al., 2011; Stramaglia et al., 2012), PID of mvTE can decompose the interaction into non-negative and non-overlapping terms. However, to address previously raised concerns that PID overestimates the redundancy term (Bertschinger et al., 2014; Pica et al., 2017), and consequently synergy, we also used an alternate implementation of PID that estimates synergy based on the lower bound of redundancy. In this implementation, the effective threshold for triads to generate synergy is higher. This approach yielded the same qualitative pattern of results.

Note, we did not examine interactions larger than triads due to the multi-fold increase in the computational burden that arises in considering higher order synergy terms. In addition to the combinatorial explosion of increased numbers of inputs, the number of PID terms increases rapidly as the number of variables increases. However, based on bounds calculated for the highest order synergy term by Timme et al. (2016), it was determined that the information gained by including an additional input beyond two either remained constant or decreased. Thus, it was inferred that lower order (two-input) computations dominated. In addition, although we did not consider more than two inputs at a time, because we considered all possible triads in each network, we effectively sub-sampled the entire space of inputs for each neuron.

Statistics

All results are reported as medians or means followed by the 95% bootstrap confidence limits (computed using 10,000 iterations) reported inside of square brackets. Accordingly, figures depict the medians or means with errorbars reflecting the 95% bootstrap confidence limits. Comparisons between conditions or against null models were performed using the

nonparametric Wilcoxon signed-rank test, unless specified otherwise. The threshold for significance was set at 0.05, unless indicated otherwise in the text. Bonferroni-Holm corrections were used in cases of multiple comparisons.

ACKNOWLEDGEMENTS

We thank Blanca Gutierrez Guzman for helpful comments and discussion.

FUNDING INFORMATION

Ehren L. Newman, Whitehall Foundation (<http://dx.doi.org/10.13039/100001391>), Award ID: 17-12-114. John M. Beggs, National Science Foundation (<http://dx.doi.org/10.13039/100000001>), Award ID: 1429500. John M. Beggs, National Science Foundation (<http://dx.doi.org/10.13039/100000001>), Award ID: 1513779. Samantha P. Faber, National Science Foundation (<http://dx.doi.org/10.13039/100000001>), Award ID: 1735095; Samantha P. Faber, Indiana Space Grant Consortium.

REFERENCES

- Basheer, I. A., & Hajmeer, M. (2000). Artificial neural networks: fundamentals, computing, design, and application. *Journal of microbiological methods*, 43(1), 3-31.
- Bastos, A. M., Usrey, W. M., Adams, R. A., Mangun, G. R., Fries, P., & Friston, K. J. (2012). Canonical microcircuits for predictive coding. *Neuron*, 76(4), 695-711.
- Bastos, A. M., Vezoli, J., Bosman, C. A., Schoffelen, J. M., Oostenveld, R., Dowdall, J. R., De Weerd, P., Kennedy, H., & Fries, P. (2015). Visual areas exert feedforward and feedback influences through distinct frequency channels. *Neuron*, 85(2), 390-401.
- Beggs, J. M., & Plenz, D. (2004). Neuronal avalanches are diverse and precise activity patterns that are stable for many hours in cortical slice cultures. *Journal of neuroscience*, 24(22), 5216-5229.
- Bertschinger, N., Rauh, J., Olbrich, E., Jost, J., & Ay, N. (2014). Quantifying unique information. *Entropy*, 16(4), 2161-2183.
- Boly, M., Garrido, M. I., Gosseries, O., Bruno, M. A., Boveroux, P., Schnakers, C., ... & Friston, K. (2011). Preserved feedforward but impaired top-down processes in the vegetative state. *Science*, 332(6031), 858-862.
- Bolz, J., Novak, N., Götz, M., & Bonhoeffer, T. (1990). Formation of target-specific neuronal projections in organotypic slice cultures from rat visual cortex. *Nature*, 346(6282), 359.

- Brincat, S. L., & Connor, C. E. (2006). Dynamic shape synthesis in posterior inferotemporal cortex. *Neuron*, 49(1), 17-24.
- Buzsáki, G., Kaila, K., & Raichle, M. (2007). Inhibition and brain work. *Neuron*, 56(5), 771-783.
- Buzsáki, G., & Mizuseki, K. (2014). The log-dynamic brain: how skewed distributions affect network operations. *Nature Reviews Neuroscience*, 15(4), 264.
- Caeser, M., Bonhoeffer, T., & Bolz, J. (1989). Cellular organization and development of slice cultures from rat visual cortex. *Experimental brain research*, 77(2), 234-244.
- Carlson, T., Tovar, D. A., Alink, A., & Kriegeskorte, N. (2013). Representational dynamics of object vision: the first 1000 ms. *Journal of vision*, 13(10), 1-1.
- Cichy, R. M., Pantazis, D., & Oliva, A. (2014). Resolving human object recognition in space and time. *Nature neuroscience*, 17(3), 455.
- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and brain sciences*, 36(3), 181-204.
- Clarke, A., Devereux, B. J., Randall, B., & Tyler, L. K. (2015). Predicting the time course of individual objects with MEG. *Cerebral Cortex*, 25(10), 3602-3612.
- Douglas, R. J., Koch, C., Mahowald, M., Martin, K. A., & Suarez, H. H. (1995). Recurrent excitation in neocortical circuits. *Science*, 269(5226), 981-985.
- Douglas, R. J., & Martin, K. A. (2007). Recurrent neuronal circuits in the neocortex. *Current biology*, 17(13), R496-R500.
- Faber, S. P., Timme, N. M., Beggs, J. M., & Newman, E. L. (2019). Computation is concentrated in rich clubs of local cortical networks. *Network Neuroscience*, 3(2), 384-404.
- Freiwald, W. A., & Tsao, D. Y. (2010). Functional compartmentalization and viewpoint generalization within the macaque face-processing system. *Science*, 330(6005), 845-851.
- Gilbert, C. D., & Li, W. (2013). Top-down influences on visual processing. *Nature Reviews Neuroscience*, 14(5), 350-363.
- Gilbert, C. D., & Sigman, M. (2007). Brain states: top-down influences in sensory processing. *Neuron*, 54(5), 677-696.
- Götz, M., & Bolz, J. (1992). Formation and preservation of cortical layers in slice cultures. *Journal of neurobiology*, 23(7), 783-802.

Grace, A. A. (2000). Gating of information flow within the limbic system and the pathophysiology of schizophrenia. *Brain Research Reviews*, 31(2-3), 330-341.

Hubel, D. H., & Wiesel, T. N. (1959). Receptive fields of single neurones in the cat's striate cortex. *The Journal of physiology*, 148(3), 574-591.

Ikegaya, Y., Aaron, G., Cossart, R., Aronov, D., Lampl, I., Ferster, D., & Yuste, R. (2004). Synfire chains and cortical songs: temporal modules of cortical activity. *Science*, 304(5670), 559-564.

Ito, S., Yeh, F. C., Hiolski, E., Rydygier, P., Gunning, D. E., Hottowy, P., Timme N., Litke A.M. & Beggs, J. M. (2014). Large-scale, high-resolution multielectrode-array recording depicts functional network differences of cortical and hippocampal cultures. *PloS one*, 9(8), e105324.

Johnson, A., Fenton, A. A., Kentros, C., & Redish, A. D. (2009). Looking for cognition in the structure within the noise. *Trends in cognitive sciences*, 13(2), 55-64.

Klostermann, O., & Wahle, P. (1999). Patterns of spontaneous activity and morphology of interneuron types in organotypic cortex and thalamus–cortex cultures. *Neuroscience*, 92(4), 1243-1259.

Kriegeskorte, N. (2015). Deep neural networks: a new framework for modeling biological vision and brain information processing. *Annual review of vision science*, 1, 417-446.

Kwon, S. E., Yang, H., Minamisawa, G., & O'Connor, D. H. (2016). Sensory and decision-related activity propagate in a cortical feedback loop during touch perception. *Nature neuroscience*, 19(9), 1243.

Leutgeb, S., Leutgeb, J. K., Moser, M. B., & Moser, E. I. (2005). Place cells, spatial maps and the population code for memory. *Current opinion in neurobiology*, 15(6), 738-746.

Lillicrap, T. P., Cownden, D., Tweed, D. B., & Akerman, C. J. (2016). Random synaptic feedback weights support error backpropagation for deep learning. *Nature communications*, 7(1), 1-10.

Litke, A. M., Bezayiff, N., Chichilnisky, E. J., Cunningham, W., Dabrowski, W., Grillo, A. A., ... & Kalmar, R. S. (2004). What does the eye tell the brain?: Development of a system for the large-scale recording of retinal output activity. *IEEE Transactions on Nuclear Science*, 51(4), 1434-1440.

Lizier, J. T., Heinzle, J., Horstmann, A., Haynes, J., & Prokopenko, M. (2011). Multivariate information-theoretic measures reveal directed information structure and task relevant changes in fMRI connectivity. *Journal of Computational Neuroscience*, 30: 85–107.

Lizier, J. T., Prokopenko, M., & Zomaya, A. Y. (2014). A framework for the local information dynamics of distributed computation in complex systems. In *Guided self-organization: inception* (pp. 115-158). Springer, Berlin, Heidelberg.

Lizier, J.T., Bertschinger, N., Jost, J., & Wibral, M. (2018). Information Decomposition of Target Effects from Multi-Source Interactions: Perspectives on Previous, Current and Future Work. *Entropy*, 20(4), 307.

Manita, S., Suzuki, T., Homma, C., Matsumoto, T., Odagawa, M., Yamada, K., ... & Ohkura, M. (2015). A top-down cortical circuit for accurate sensory perception. *Neuron*, 86(5), 1304-1316.

Markram, H., Lübke, J., Frotscher, M., Roth, A., & Sakmann, B. (1997). Physiology and anatomy of synaptic connections between thick tufted pyramidal neurones in the developing rat neocortex. *The Journal of physiology*, 500(2), 409-440.

Mason, A., Nicoll, A., & Stratford, K. (1991). Synaptic transmission between individual pyramidal neurons of the rat visual cortex in vitro. *Journal of Neuroscience*, 11(1), 72-84.

Nassi, J. J., Lomber, S. G., & Born, R. T. (2013). Corticocortical feedback contributes to surround suppression in V1 of the alert primate. *Journal of Neuroscience*, 33(19), 8504-8517.

Neunuebel, J. P., & Knierim, J. J. (2014). CA3 retrieves coherent representations from degraded input: direct evidence for CA3 pattern completion and dentate gyrus pattern separation. *Neuron*, 81(2), 416-427.

Nurminen, L., Merlin, S., Bijanzadeh, M., Federer, F., & Angelucci, A. (2018). Top-down feedback controls spatial summation and response amplitude in primate visual cortex. *Nature communications*, 9(1), 1-13.

Olshausen, B. A., & Field, D. J. (2004). Sparse coding of sensory inputs. *Current opinion in neurobiology*, 14(4), 481-487.

Pica, G., Piasini, E., Chicharro, D., & Panzeri, S. (2017). Invariant components of synergy, redundancy, and unique information among three variables. *Entropy*, 19(9), 451.

Plenz, D., & Aertsen, A. (1996). Neural dynamics in cortex-striatum co-cultures—II. Spatiotemporal characteristics of neuronal activity. *Neuroscience*, 70(4), 893-924.

Rolls, E. T. (2007). An attractor network in the hippocampus: theory and neurophysiology. *Learning & Memory*, 14(11), 714-731.

Rolls, E. (2013). The mechanisms for pattern completion and pattern separation in the hippocampus. *Frontiers in systems neuroscience*, 7, 74.

Schreiber, T. (2000). Measuring information transfer. *Phys Rev Lett* 85:461– 464

Sherrill, S. P., Timme, N. M., Beggs, J. M., & Newman, E. L. (2019). Correlated activity favors synergistic processing in local cortical networks at synaptically-relevant timescales. *bioRxiv*, 809681.

Shimono, M., & Beggs, J. M. (2015). Functional clusters, hubs, and communities in the cortical microconnectome. *Cerebral Cortex*, 25(10), 3743-3757.

Sikkens, T., Bosman, C. A., & Olcese, U. (2019). The role of top-down modulation in shaping sensory processing across brain states: implications for consciousness. *Frontiers in systems neuroscience*, 13, 31.

Song, S., Sjöström, P. J., Reigl, M., Nelson, S., & Chklovskii, D. B. (2005). Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS biology*, 3(3), e68.

Spoerer, C. J., McClure, P., & Kriegeskorte, N. (2017). Recurrent convolutional neural networks: a better model of biological object recognition. *Frontiers in psychology*, 8, 1551.

Sugase, Y., Yamane, S., Ueno, S., & Kawano, K. (1999). Global and fine information coded by single neurons in the temporal visual cortex. *Nature*, 400(6747), 869-873.

Tang, H., Buia, C., Madhavan, R., Crone, N. E., Madsen, J. R., Anderson, W. S., & Kreiman, G. (2014). Spatiotemporal dynamics underlying object completion in human ventral visual cortex. *Neuron*, 83(3), 736-748.

Tang, H., Schrimpf, M., Lotter, W., Moerman, C., Paredes, A., Caro, J. O., ... & Kreiman, G. (2018). Recurrent computations for visual pattern completion. *Proceedings of the National Academy of Sciences*, 115(35), 8835-8840.

Thom, M., & Palm, G. (2013). Sparse activity and sparse connectivity in supervised learning. *Journal of Machine Learning Research*, 14(Apr), 1091-1143.

Timme, N., Ito, S., Myroshnychenko, M., Yeh, F. C., Hiolski, E., Hottowy, P., & Beggs, J. M. (2014). Multiplex networks of cortical and hippocampal neurons revealed at different timescales. *PLoS ONE*, 9(12), e115764.

Timme, N. M., Ito, S., Myroshnychenko, M., Nigam, S., Shimono, M., Yeh, F. C., Hottowy, P., Litke, A.M., & Beggs, J. M. (2016). High-degree neurons feed cortical computations. *PLoS Computational Biology*, 12(5), e1004858.

Treves, A., Rolls, E. T., & Simmen, M. (1997). Time for retrieval in recurrent associative memories. *Physica D: Nonlinear Phenomena*, 107(2-4), 392-400.

Wibral, M., Priesemann, V., Kay, J. W., Lizier, J. T., & Phillips, W. A. (2017). Partial information decomposition as a unified approach to the specification of neural goal functions. *Brain and cognition*, 112, 25-38.

Williams, P. L., & Beer, R. D. (2011). Generalized measures of information transfer. *arXiv preprint arXiv:1102.1507*.