Computation is concentrated in rich clubs of local cortical neurons

Samantha P. Faber[1*], Nicholas M. Timme[2], John M. Beggs[3], Ehren L. Newman[1]

1 Department of Psychological and Brain Sciences, Indiana University Bloomington, 1101 E. 10th St, Bloomington, IN, 47405, U.S.A.

2 Department of Psychology, Indiana University-Purdue University Indianapolis, 402 N Blackford, LD 124, Indianapolis, IN, 46202, U.S.A.

3 Department of Physics, Indiana University Bloomington, Swain Hall West, 727 E. 3rd St, Bloomington, IN, 47405, U.S.A.

* corresponding author – samfaber@indiana.edu

KEYWORDS: networks; effective connectivity; neural computation; rich clubs; transfer entropy

ABSTRACT

To understand how neural circuits process information, it is essential to identify the relationship between computation and circuit topology. Rich-clubs, highly interconnected sets of neurons, are known to propagate a disproportionate amount of information within cortical circuits. Here, we test the hypothesis that rich-clubs also perform a disproportionate amount of computation. To do so, we recorded the spiking activity of on average ~300 well-isolated individual neurons from organotypic cortical cultures. We then constructed weighted, directed networks reflecting the effective connectivity between the neurons. For each neuron, we quantified the amount of computation it performed based on its inputs. We found that rich-club neurons compute ~200% more information than neurons outside of the rich club. Indeed, the amount of computation performed in the rich-club was proportional to the amount information propagation by the same neurons. This suggests that, in these circuits, information propagation drives computation. Comparing the computation-to-propagation ratio inside versus outside of the rich club showed that rich clubs compute at a slightly, though significantly, reduced level (~4% lower). In total, our findings indicate that rich club topology in effective cortical circuits supports not only information propagation but also neural computation.

AUTHOR SUMMARY

Here we answer the question of whether rich club topology in functional cortical circuits supports neural computation as it has been previously shown to do for information propagation. To do so, 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. We found that neurons in rich clubs computed significantly more than neurons outside of rich clubs, suggesting that rich-clubs do support computation in cortical circuits. Indeed, the amount of computation that we found in the rich club was proportional to the amount of information they propagate suggesting that, in these circuits, information propagation drives computation.

INTRODUCTION

The idea that neurons propagate information and that downstream neurons integrate this information via neural computation is foundational to our understanding of how the brain processes, and responds to, the world. Yet, the determinants of such computations remain largely unknown. Advances in data acquisition methods, offering increasingly comprehensive recordings of the activation dynamics that

play out atop of neural circuits, together with advances in data analytics now make it possible to empirically study the determinants of neural computation. Using these tools, we addressed the fundamental question of where, relative to information flow in local cortical networks, the majority of neural computation takes place.

In this context, neural computation is the systematic transformation of information received by a neuron (determined by analyzing its output with respect to its inputs) based on the input of multiple upstream neurons (Timme et al., 2016). This type of computation can be detected empirically when the activity of upstream neurons accounts for the activity of a downstream neuron better when considered jointly than when treated as independent sources of variance. Analytical tools adapted from Shannon's information theory make it possible to track such computation as well as information propagation in networks of spiking neurons (Strong et al., 1998; Borst & Theunissen, 1999; Schreiber, 2000; Williams & Beer, 2010).

Previously, Timme et al. (2016) used these tools to show that computation does not vary systematically with the number of inputs received by a neuron, as might be intuited. Rather, computation correlates with the number of outputs of the upstream neurons. This counter-intuitive finding suggested that the amount of computation a neuron performs is better predicted by its position in the broader topographic structure of the circuit than by the local connectivity. This raises the question of what topological conditions support computation.

Local cortical networks, like many complex networks, contain 'rich clubs.' That is, the most strongly connected neurons interconnect with a higher probability than would be expected by chance. The existence of a rich club in a functional network predicts that a select set of highly-integrated nodes handles a disproportionately large amount of traffic. Indeed, in the local networks of cortical circuits, 20% of the neurons account for 70% of the information propagation (Nigam et al., 2016). As such, rich clubs represent a conspicuous topographic landmark in the flow of information across neural circuits.

Thus, here we addressed the critical question: What is the role of the rich club with respect to neural computation? We tested among three alternatives. First, computation is constant throughout the topology of a network – predicting that rich clubs do not perform more or less computation than would be expected by chance. Second, computation grows with increasing information availability–predicting that rich clubs are rich in computation given their high information density. Third, computation decreases with increasing information availability– predicting that rich clubs are computationally poor. To test among these three hypotheses, we recorded the spiking activity of hundreds of neurons from organotypic cultures of mouse somatosensory cortex and assessed the distribution of computation relative to the topology of information propagation (Figure 1). The results demonstrate that rich clubs account for the majority of network computation (~90%) and that computation occurs at a degree proportional to information propagation. However, we found that computation generally occurs at a stable level relative to information propagation and that this level is slightly (~4%) smaller inside versus outside of the rich club. This result suggests that rich clubs are computationally dense, and yet, because they are even more propagation-dense, the ratio of computation relative to propagation is reduced in rich clubs. Thus, rich clubs support elevated amounts of both computation and propagation relative to the rest of the network; however, they support propagation even more so than computation.

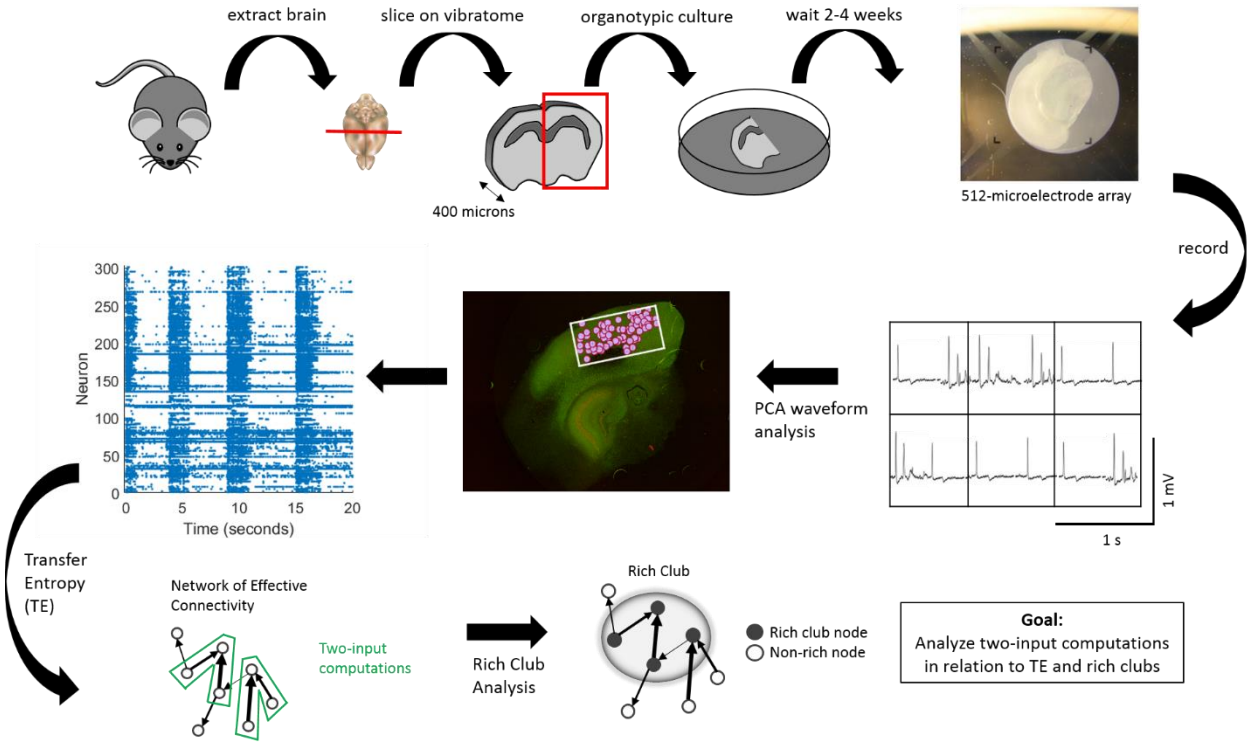

Figure 1. Experimental and data analysis procedure. (Top row, left to right) Brains were extracted from mouse pups and sliced using a vibratome. Slices containing somatosensory cortex were organotypically cultured for up to 2 to 4 weeks. Cultures were then placed on a recording array and recorded for 1 hour. (Middle row, right to left) Recordings yielded neuron spiking dynamics at each electrode – waveforms at six example electrodes shown – which were sorted using PCA in order to isolate individual cells based on their distinct waveforms. Once cells were isolated and localized (pink circles) within the recording area (white rectangle), their corresponding spike trains could be determined. (Bottom Row, left to right) Spike trains were then used to compute Transfer Entropy, at multiple timescales, between each neuron pair in a recording. This resulted in networks of effective connectivity. Computations occurring at neurons receiving two connections were then calculated using Partial Information Decomposition. A rich club analysis was used to detect collections of hub neurons that connect to each other. Finally, we examined the relationship between TE within a triad and two-input computations as well as between two-input computations and rich clubs.

RESULTS

To study the relationship between information processing and the functional network structure of cortical microcircuits, we recorded spiking activity of between 98 and 594 neurons (median = 310) simultaneously from 25 organotypic cultures of mouse somatosensory cortex. We then quantified the information transfer from each neuron to every other neuron at two timescales relevant for synaptic transmission (1.6 – 6.4 ms and 3.5 – 14 ms; Mason et al., 1991; Swadlow, 1994) using transfer entropy (TE; Schreiber, 2000). This resulted in 50 networks (25 recordings at two timescales). Because $TE_{A \rightarrow B}$ quantifies how much information is gained about whether neuron B will spike in the next moment obtained by knowing whether A is now spiking, it provides a directional, weighted effective connection. Significant TE values (determined by comparing to the distribution of TE values obtained when we shuffled the respective spike trains) were kept and used to build a graph of the circuit's functional network. The graph was then used to identify triads–consisting of two neurons with significant connections to a common neuron–for analysis of neural computation. Our analysis of neural computation was predicated on the idea that non-linear integration of multiple inputs implements a form of computation (see also Timme et al., 2016). Thus, we quantified neural computation as the additional information regarding the future state of the receiver that was gained beyond what the sending neurons offered individually (after accounting for the redundancy between the upstream

neurons). We did this using a partial information decomposition approach (Williams & Beer, 2010) to compute the three-way synergy. All TE and synergy values were normalized by the entropy of the receiver neuron in order to cast them in terms of the proportion of the receiver neuron's capacity that is accounted for by the transmitting neuron, or by computation, respectively. Because of this, all TE and synergy values are in terms of bits per bit. All results are reported as medians with 95% bootstrap confidence intervals (computed using 10,000 iterations) presented in brackets.

Computation and information propagation vary widely in cortical microcircuits

When building the networks, we found that 0.56% [0.38% 1.34%] of all possible directed connections between neurons were significant at the $\alpha = 0.001$ level (e.g., 480 of 81510 possible connections, or 0.59%, were significant in a network of 286 neurons). To consider the amount of information that was used in two-input computations, we defined information propagation as the sum of the two inputs (normalized TE values) converging on a neuron. Across neurons, the distribution of propagation values was approximately lognormal (shown in Figure 2B), consistent with previously observed distributions of both structural (Song et al., 2005; Lefort et al., 2009; Ikegaya et al., 2012) and functional connectivity (for a review, see Buzsáki & Mizuseki, 2014). This lognormality indicates that there exists a long tail of large propagation values such that a few neurons propagate particularly large amounts of information. Concretely, we found that the strongest 8.1% [7.2% 9.2%] of neurons propagated as much information as the rest of the neurons combined. Within a network, the difference between the neurons that received the most versus the least information commonly spanned 4.1 [3.8 4.3] orders of magnitude.

Computation, like propagation, varied in a lognormal fashion (Figure 2B). The top 8.2% [7.3% 9.1%] of neurons computed the same amount of information as the rest of the neurons combined. The normalized synergy typically spanned 4.1 [3.7 4.4] orders of magnitude over neurons in individual networks. Notably, computation was reliably smaller than propagation (0.14 vs 0.026; $Z_{s.r.} = 6.15$, n = 50 networks, $p = 7.6 \times 10^{-10}$) indicating that, despite finding substantial computation, most information was accounted for by neuron-to-neuron propagation. The variability in computation motivated us to test the hypothesis that the dense inter-connectivity of rich clubs serves as a hub, not only for propagation, but also for computation.

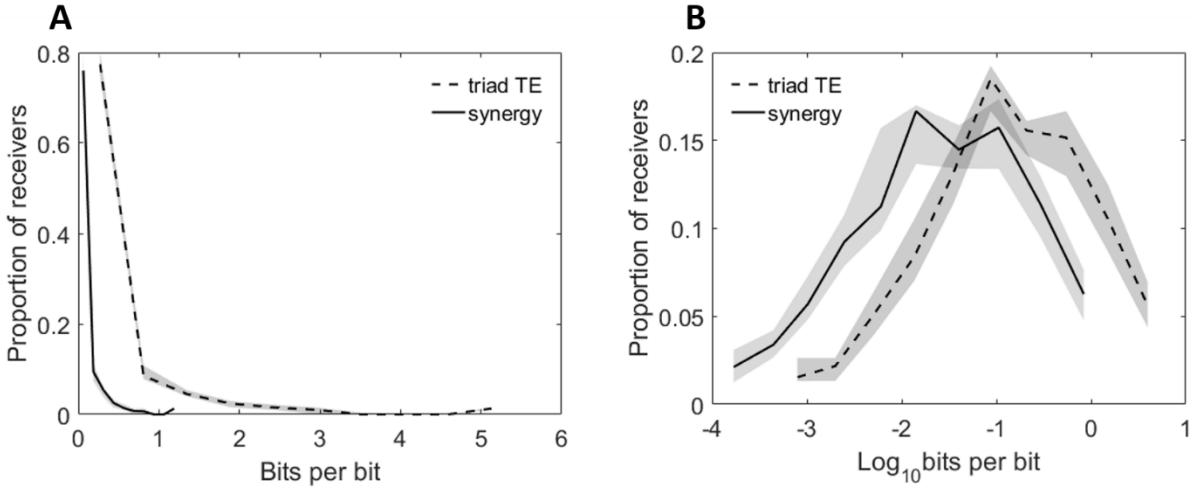

Figure 2. Distributions of neuron computation (normalized synergy) and propagation (normalized triad TE) are highly varied. Histograms of synergy and triad TE values for all receivers in all networks at both timescales. **A.** Distributions of measured values. **B.** Distributions of log-scaled values to emphasize variability. Solid lines depict the median across networks and shaded regions depict 95% bootstrap confidence intervals around the median.

Rich clubs perform a majority of the network-wide computation

In every network, we found significant rich clubs at multiple thresholds consisting of 10% - 40% of the neurons in each network (Figure 3). We investigated the relationship between these rich clubs and computation using multiple approaches. In the first approach, we asked if the expected (i.e., mean) normalized synergy-per-triad was significantly greater for triads with receivers inside, versus outside, of a randomly selected rich club from each network. Indeed, we found that expected synergy was greater inside of the rich club (0.03 vs. 0.01, $Z_{s.r.} = 5.8$, n = 50 networks, $p = 5.7 \times 10^{-9}$; Figure 4A). Rich clubs consist of a minority of the neurons in a network, so we next asked what percentage of the network-wide computation was performed inside of the rich club. We found that, in randomly selected rich clubs, 89.8% [80.8% 95.6%] of all computation in a network was performed by rich club neurons (Figure 4B).

Our second approach examined how computation varied over all possible richness thresholds. The results of these analyses recapitulate the findings reported above. That is, across thresholds, richer neurons consistently had greater expected synergy-per-triad (Figure 4C). Likewise, at most thresholds, rich neurons accounted for a majority of the network-wide computation (Figure 4D). Because the richest neurons are likely to participate in a larger number of triads, rich clubs would be expected to compute a large percentage of the network-wide computation even if the expected

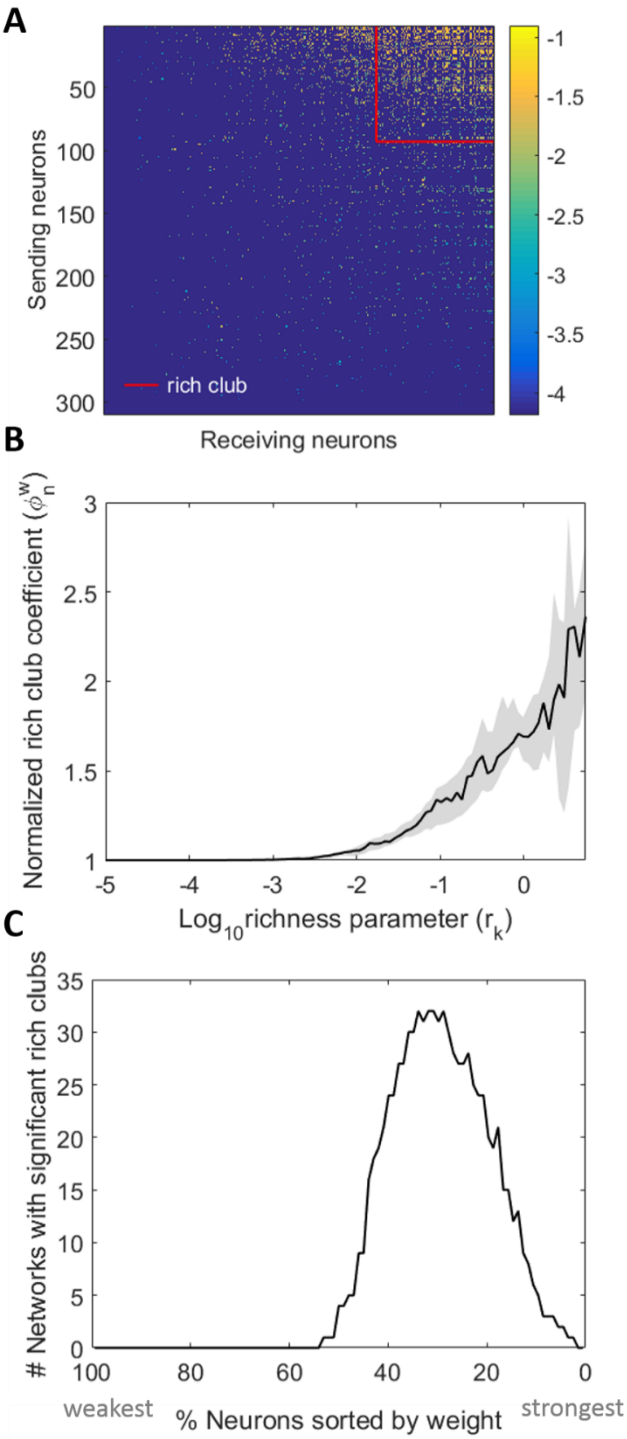

Figure 3. Networks reliably show rich clubs. **A.** Adjacency matrix of a representative 310-neuron network with rich clubs. Rich club of top 30% of neurons depicted. Neurons sorted in order of increasing richness from left to right and bottom to top. TE-values are log-scaled. **B.** Normalized, weighted rich club coefficients for all networks. X-axis is richness parameter level, where the richness parameter is the sum of the weighted connections for each neuron. Solid line represents median across all networks; shaded region is 95% bootstrap confidence interval around the median. In order for a rich club to be recruited into the synergy analysis, coefficients were required to be significant ($p < 0.01$) when compared to those from randomized networks. **C.** Distribution of the number of networks with significant rich clubs at each threshold. The majority of networks had significant rich clubs comprised of the top 40% to 10% of the network.

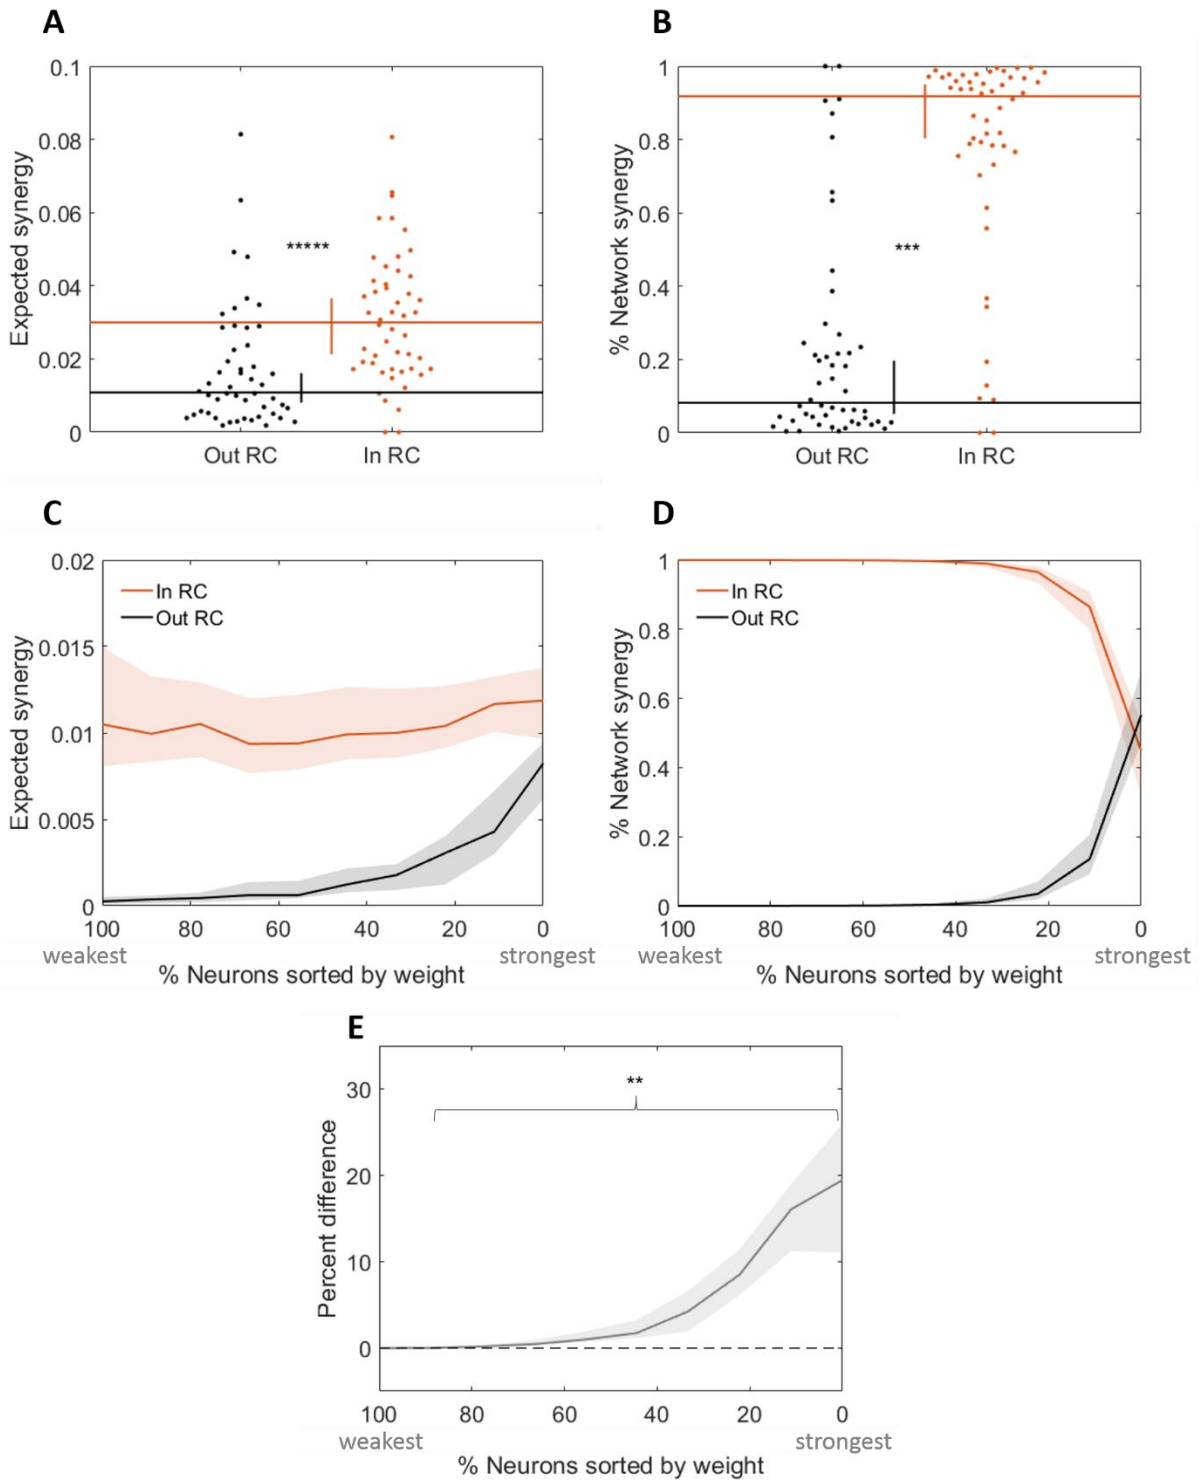

Figure 4. Rich club neurons compute more than non-rich club neurons. **A.** Triads with receivers in the rich club (RC) have greater median expected normalized synergy (compute more) than those with receivers outside the rich club. **B.** Triads with receivers in the RC perform a significantly larger percentage of the total network computation than triads with receivers outside the RC. Distributions shown here are complementary; values sum to 1. **C.** At all possible rich club levels, triads with receivers in the rich club have greater median expected synergy than those with receivers outside the rich club. **D.** At all significant rich club levels, a greater percentage of network synergy occurs in the rich clubs than outside the rich clubs. Distributions shown here

are complementary; values sum to 1. **E.** At all significant rich club levels, a greater percentage of synergy is accounted for by a smaller percentage of triads in the rich clubs. Significance indicators: *****$P < 1 \times 10^{-8}$; ***$P < 1 \times 10^{-5}$; **$P < 1 \times 10^{-4}$

synergy-per-triad was not significantly greater than that found elsewhere in the network. Thus, we also asked how the percentage of network-wide computation varied as a function of the percentage of triads that are included in the rich club. As shown in Figure 4E, the share of computation performed by the richest neurons was consistently greater than the percentage of above threshold triads. Notably, the percent difference drops off as the threshold decreases, reflecting that as the rich clubs become less rich, they perform less relative amounts of computation.

Our third approach compared the relationship between the strength of the rich club for any given threshold, as indicated by the normalized rich club coefficient, to the amount of computation performed by that rich club. Specifically, we correlated the normalized rich club coefficient with the expected normalized synergy of triads inside of the rich club across thresholds for each network separately and asked if there was a consistent trend across networks (Figure 5). In most networks, computation and normalized rich club coefficient were positively correlated (43 of 50 networks) such that the median correlation coefficient of 0.79 [0.62 0.89] was significantly greater than zero ($Z_{s.r.} = 5.36$, n = 50 networks, $p = 8.2 \times 10^{-8}$; Figure 5B). These results indicate that rich club strength was strongly predictive of expected computation.

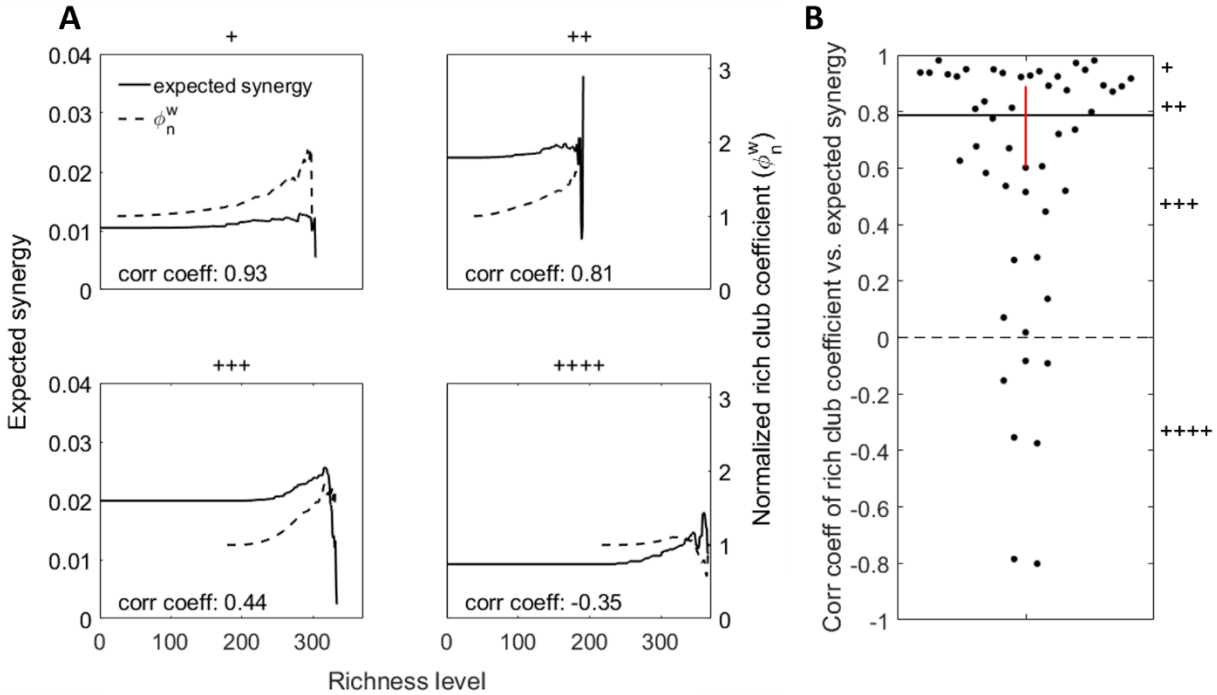

Figure 5. Normalized rich club coefficient correlates with synergy. **A.** Normalized rich club coefficients and expected normalized synergy at increasing richness levels for four representative networks. Negative correlations are observed in networks that have poor rich clubs, or in which the expected synergy decreases as we consider fewer, richer neurons. The second case is observed in networks whose top neurons participate in many triads with synergy values that are highly variable. **B.** Distribution of correlation coefficients for correlations between rich club coefficients and expected triad synergy at all richness levels. Most network rich club coefficients are positively correlated with expected triad synergy. This shows that rich clubs are predictive of increased synergy levels.

Finally, we asked how expected synergy depends upon which triad members were in the rich club. The analyses described above categorize a triad as inside of the rich club if the receiver is a member of the rich club. Here, we computed the expected normalized synergy separately for each of the six possible

alignments of the triad members to the rich club. The results, shown in Figure 6, demonstrate that the single configuration with the largest expected computation occurred in triads wherein all three members were in the rich club. The two configurations with the next greatest expected computation both had two members in the rich club. Configurations with one member in the rich club had the next greatest expected computation. Triads that had no neurons in the rich club had the lowest expected computation. Consistent with the above results, this pattern indicates that the greater the involvement in the rich club, the greater the expected computation.

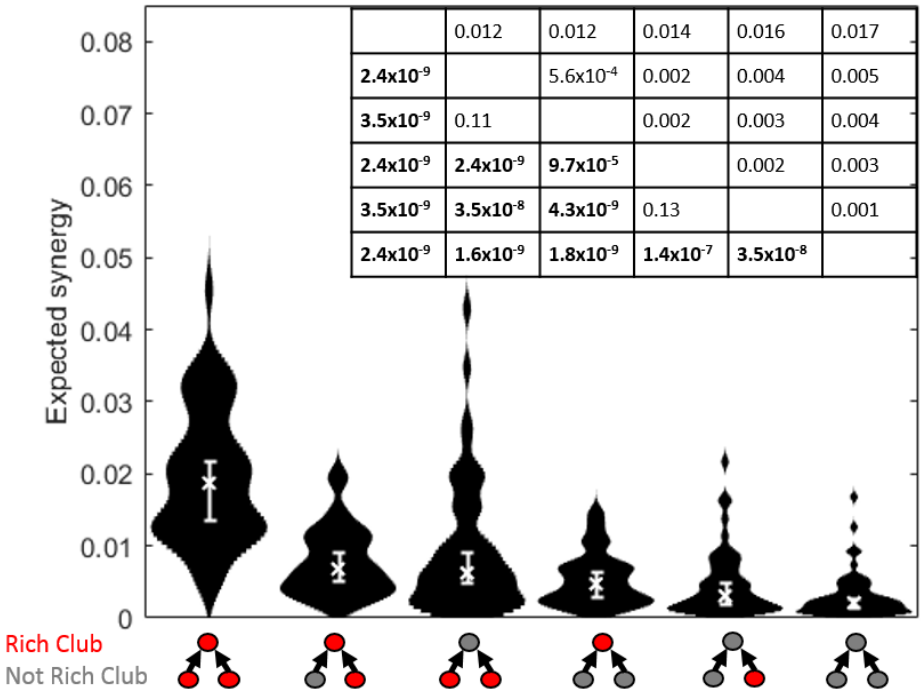

	0.012	0.012	0.014	0.016	0.017
2.4×10^{-9}		5.6×10^{-4}	0.002	0.004	0.005
3.5×10^{-9}	0.11		0.002	0.003	0.004
2.4×10^{-9}	**2.4×10^{-9}**	**9.7×10^{-5}**		0.002	0.003
3.5×10^{-9}	**3.5×10^{-8}**	**4.3×10^{-9}**	0.13		0.001
2.4×10^{-9}	**1.6×10^{-9}**	**1.8×10^{-9}**	**1.4×10^{-7}**	**3.5×10^{-8}**	

Figure 6. Greater computation is performed by triads with greater numbers of neurons in the rich club. Distributions of expected synergy for each of all possible triad interactions with the rich club. Triads that have all members in the rich club compute the most. Triads with a single transmitter and the receiver in the rich club, and both transmitters in the rich club compute equally well. Triads with only the receiver in the rich club compute better than triads with a single transmitter in the rich club and. All triads with any member in the rich club compute better than triads with no members in the rich club. Medians, denoted by 'x', and 95% bootstrap confidence intervals are shown. Table shows p-values (lower diagonal) and differences of medians (upper diagonal) of pairwise comparisons between the conditions, which are sorted by median expected synergy. Significant p-values are bolded.

Stable computation-to-propagation ratio accounts for the high density of computation in rich clubs

Our results show that rich club neurons both propagate a majority of the information (Nigam et al., 2016) and perform a majority of the computation in the network. A possible explanation for the co-localization of information propagation and computation is that propagation drives computation. To investigate this, we asked how correlated information propagation and computation were across triads for each network. In every network, propagation was strongly correlated with computation ($\rho = 0.75$ [0.74 0.77], minimum $\rho = 0.57$, $Z_{s.r.} = 6.15$, n = 50 networks, $p = 7.6\times10^{-10}$; Figures 7A and 7B).

Seeing that computation and propagation were highly correlated across triads, we asked what range of computation ratios (i.e., computation / propagation) occurred across networks. As shown in Figures 7C and 7D, the computation ratio was highly stereotyped over networks with a median of 0.2358 [0.2353

0.2364]. In contrast to the 4.1 [3.8 4.3] and 4.1 [3.7 4.4] orders of magnitude over which normalized triad TE and synergy varied, respectively, the computation ratio varied by 1.5 [1.4 1.6] orders of magnitude. By way of comparison, randomizing the alignment of synergy to triad TE across triads results in computation ratios that span 6.1 [5.5 6.5] orders of magnitude. The significant reduction in variance over what would be expected by chance ($Z_{s.r.}$ = -6.15, n = 50 networks, p = 7.6×10^{-10}) suggests that the computation ratio is a relatively stable property of neural information processing in such networks. The implication of a stable computation ratio across triads is that information propagation will reliably be accompanied by computation, thereby accounting for the high density of computation in propagation-dense rich clubs.

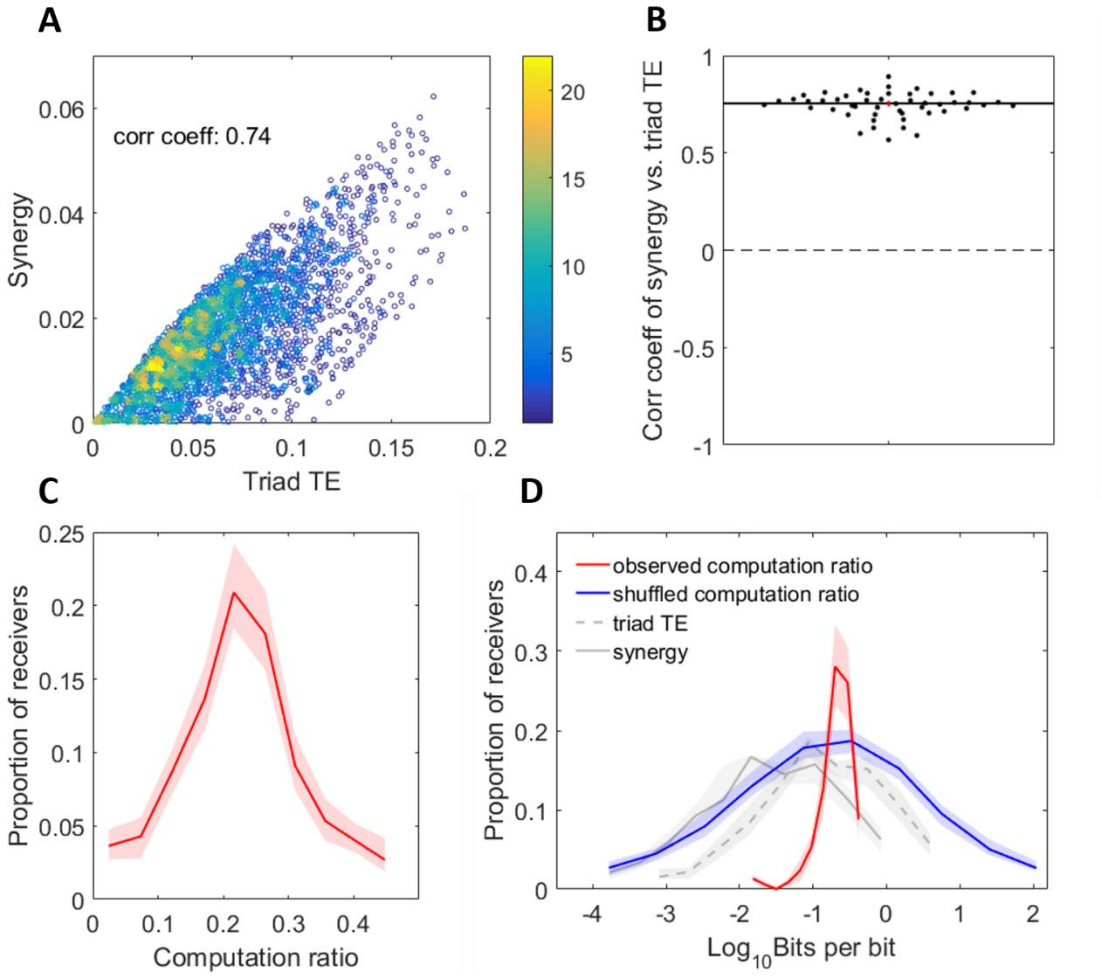

Figure 7. Propagation is highly predictive of computation. **A.** Scatterplot of normalized synergy (computation) versus normalized triad TE (propagation) in a representative network with 3448 triads. Colorbar depicts point density. Also shown is the correlation coefficient. **B.** Distribution of network correlations between synergy and triad TE. This shows that computation was strongly, positively correlated with propagation across all networks. **C.** Histogram of computation ratio values for all receivers in all networks. **D.** Histogram of log-scaled computation ratio values for all receivers in all networks. Gray lines are replotted here from Figure 2 for ease of comparison. The blue line represents the distribution of computation ratios that results from shuffling the alignment of triad synergy to TE. Thus, the span of observed computation ratios is significantly smaller than what we might have observed by chance. For C and D, Solid and dashed lines depict the medians across networks and shaded regions depict 95% bootstrap confidence intervals around the medians.

Because the computation ratio allows computation to be predicted from propagation, it is informative to ask if this ratio varies as a function of rich club membership. We found that the ratio was very slightly, though significantly, lower for triads inside, versus outside, of rich clubs (0.22 vs. 0.24; $Z_{s.r.}$ = -2.5, n = 50

networks, p = 0.014; Figures 8A). This difference was reliable across the same range of richness thresholds that were associated with significant rich clubs (Figure 8B). Indeed, the strength of the rich club was significantly negatively correlated with the computation ratio (ρ = -0.50 [-0.64 -0.16], $Z_{s.r.}$ = -3.1, n = 50 networks, p = 0.002; Figure 8C). These results suggest that there is a trade-off between computation and propagation, with reduced levels of computation at high levels of propagation. Importantly, however, computation ratios were only slightly reduced (~4%) relative to the substantially greater information propagation values found in the rich club (~200%), thereby leaving rich clubs overall dense in computation (see Figure 10).

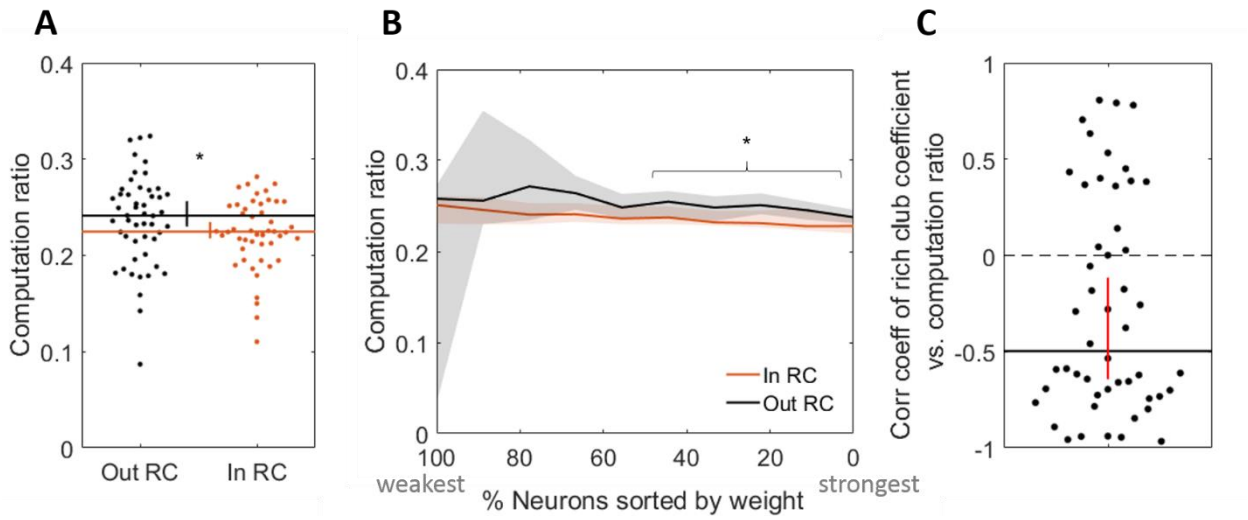

Figure 8. Rich clubs are predictive of a reduced ratio of computation relative to propagation (computation ratio). **A.** Triads with receivers in the rich club (RC) have a lower median expected computation ratio (compute less relative to how much they propagate) than those with receivers outside the rich club. **B.** At all significant rich club levels, triads with receivers in the rich club have a lower median expected computation ratio than those with receivers outside the rich club. **C.** Distribution of correlation coefficients for correlations between expected computation ratio and normalized rich club coefficient at all richness levels. Significance indicators: *P < 0.05.

Operationalization of computation was not critical for present results

To investigate whether our findings are robust to the method of quantifying computation, we implemented two alternate methods of identifying computation to test if the same results were obtained. In the first, we used an alternate implementation of partial information decomposition (PID). Unlike the standard PID approach, which effectively computes the upper bound on synergy (by assuming maximum redundancy between senders), our alternate implementation effectively computes the lower bound of synergy (by assuming no redundancy between senders). When using this approach, we find the same pattern of results. That is, expected synergy-per-triad is significantly greater inside versus outside of the rich clubs (0.01 vs. 0.005; $Z_{s.r.}$ = 4.2, n = 50 networks, p = 2.5×10^{-5}; data not shown); the computation is significantly positively correlated with information propagation (ρ = 0.53 [0.44 0.59]; $Z_{s.r.}$ = 6.1, n = 50 networks, p = 1.23×10^{-9}); and the computation ratio is slightly, but significantly, lower inside versus outside of the rich club (0.20 vs. 0.21; $Z_{s.r.}$ = -1.66, n = 50 networks, p = 0.01; data not shown).

Our second alternate method of identifying computation estimated the input-output transfer functions of individual neurons as described by Chichilnisky (2001). In the prior analyses, using PID, we took

synergy as evidence of computation because it quantifies how much more information is carried by multiple neurons when considered together than the sum of information carried by the same neurons when considered individually. Likewise, in this analysis, we took nonlinear input-output functions as evidence of computation because they indicate that a neuron does not simply echo its inputs but, rather, responds based on patterns of upstream inputs. Accordingly, we performed a median split across neurons based on the linearity of their estimated input-output functions. Those with the most nonlinear functions were identified as neurons that likely perform more computation. Comparing this classification to the normalized synergy values obtained via PID, we found that neurons with nonlinear transfer functions were also found to have significantly greater amounts of synergy than those with linear transfer functions (0.116 vs. 0.056; $Z_{s.r.} = 4.73$, n = 50 networks, $p = 2.19 \times 10^{-6}$). As such, use of this classification regime provides an independent, yet related, approach to identifying computation. Consistent with our main results, the concentration of neurons exhibiting nonlinear transfer functions was significantly greater inside versus outside of the rich clubs (68.3% vs. 41.6%; $Z_{s.r.} = 6.15$, n = 50 networks, $p = 7.56 \times 10^{-10}$; Figure 9).

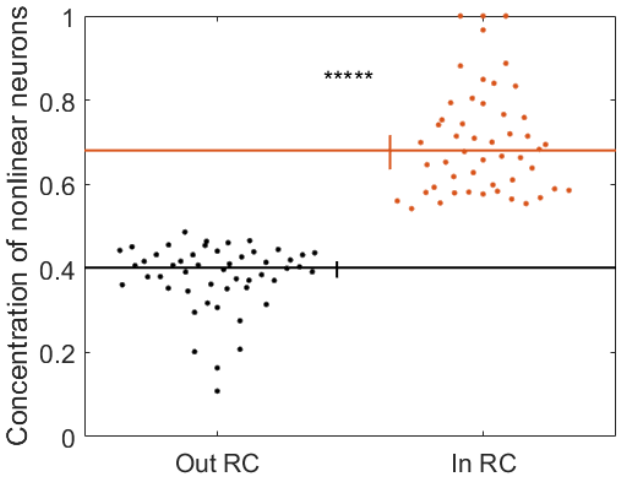

Figure 9. Alternative measure of neural computation reveals results that correspond to those obtained using PID. Neurons with nonlinear transfer functions are represented more inside rich clubs than they are outside rich clubs. Significance indicators: *****$P < 1 \times 10^{-8}$

DISCUSSION

Our goal in this work was to test the hypothesis that neural computation in cortical circuits varies in a systematic fashion with respect to the functional topology of those circuits. Specifically, we asked if neurons in rich clubs perform more computation than those outside of the rich clubs. To answer this question, we recorded the spiking activity of hundreds of neurons in organotypic cultures simultaneously and compared the information processing qualities of individual neurons to their relative position in the broader functional network of the circuit. We found that neurons in rich clubs computed ~200% more than neurons outside of rich clubs. The amount of computation that we found in the rich club was proportional to the amount of information they propagate suggesting that, in these circuits, information propagation drives computation. These results are summarized in Figure 10.

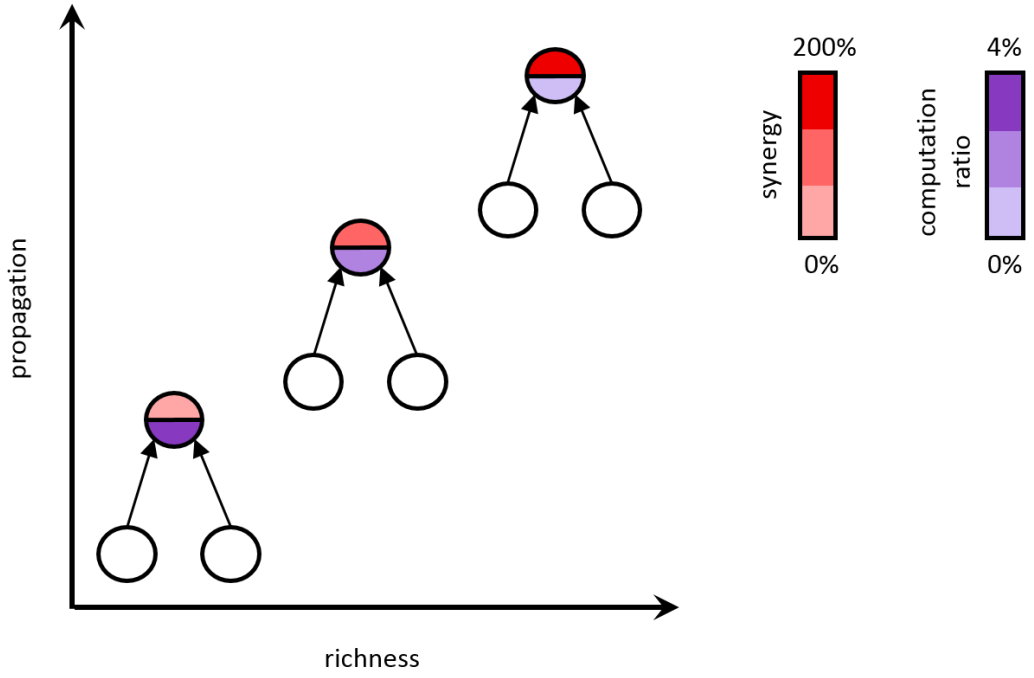

Figure 10. The synergy (computation) of triad receiver neurons increases significantly with propagation (arrows) and node richness. The computation ratio (computation relative to propagation) decreases modestly as node richness increases. Thus, rich club neurons compute much more than non-rich club neurons, although they compute somewhat less relative to how much they propagate.

Finding computation in cortical networks

What does it mean for organotypic cultures to process information? In a neural circuit that has no clear sensory inputs, it is easy to imagine that all spiking is either spontaneous or in response to upstream spontaneous spiking. Spontaneous spikes contribute to what, in an information theoretic framework, is considered entropy. When the spiking of upstream neurons allows us to predict future spiking, one colloquially says that those neurons carry information about the future state of the circuit. This is technically accurate because that prediction effectively reduces the uncertainty of the future state of the system. The cause of the upstream spiking, whether spontaneous or sensory driven, does not change this logic. As in systems with intact sensory inputs, the neurons in organotypic cultures process information by integrating received synaptic inputs. Though we argue that, in this respect, information processing by organotypic cultures is like the cortical circuits of intact animals, we recognize that important differences accompany the lack of true sensory input. For example, the spatiotemporal structure of the spontaneous spiking that drives activity in cultures is likely fundamentally different from that driven by sensory experience. Understanding how that structure influences information processing will be important to investigate as technologies enabling high temporal resolution recordings of hundreds of neurons *in vivo* mature.

With regard to building an understanding of the drivers of neural computation, the present work substantially builds on previous work that showed that computation was positively correlated with the number of outgoing connections (i.e., out degree) of the upstream neurons (Timme et al., 2016). Though the rich club membership of those neurons was not analyzed, it is reasonable to assume that neurons with high degree would be included in rich clubs. In that respect, our findings are consistent with the previous report. The work of Timme et al. (2016), however, looked only at degree and did not analyze

edge weights. Here, we found a strong correlation between synergy and information propagation (i.e., summed edge weights contributing to computation), indicating that computation is strongly dependent upon weight. This substantially alters our understanding of computation as it shows that, beyond the pattern of connections constituting a network, computation is sensitive to the quantity of the information relayed across individual edges.

Rich clubs as a home for computation

Prior work on rich clubs convincingly argued that rich clubs play a significant role in the routing of network information (van den Heuvel et al., 2011; Harriger et al., 2012; van den Heuvel et al., 2013; Nigam et al., 2016). For example, Nigam et al. (2016) showed, using the same data, that the top 20% richest neurons transmit ~70% of network information. Here, we showed that information propagation is directly related to computation. Combining this discovery with the previous knowledge that rich clubs perform large amounts of information propagation accounts for the high densities of computation we observed in the cortical circuit rich clubs. Though the precise mechanism of how computation is derived from propagation is unknown, one possibility is that it is the result of what one might consider to be 'information collisions.' This idea is based on the finding of Lizier et al. (2010) who demonstrated that the dominant form of information modification (i.e., computation) in cellular automata are the result of collisions between the emergent particles (see also Adamatzky and Durand-Lose, 2012; Bhattacharjee et al., 2016; Sarkar, 2000). In the context of our circuits, the idea is that computation arises when packets of information embedded in the outgoing spike trains of sending neurons collide onto the same receiving neuron in sufficient temporal proximity to alter the way the receiver responds to those inputs. The density of propagating information in rich clubs would proportionately increase the likelihood of such collisions, and thereby increase the amount (and number) of computation(s) performed by the rich club (Flecker et al., 2011).

Detecting computation using partial information decomposition

Our primary analyses used synergy as a proxy for computation among triads consisting of a pair of transmitting neurons and a single receiving neuron, following the methods of Timme et al. (2016). Synergy, as a measure of the information gained when the pair of transmitters is considered jointly over the combined information carried by the neurons individually, provides an intuitively appealing measure of computation (see Timme et al., 2016 for a comprehensive discussion of this relationship). Synergy is a product of partial information decomposition (PID) (Williams & Beer, 2010). However, PID is not the only information theoretic tool available for quantifying neural computation. Our use of PID was motivated by several factors: 1) it can detect linear and nonlinear interactions; 2) it is capable of measuring the amount of information a neuron computes based on simultaneous inputs from other neurons; and 3) it is currently the only method capable of quantifying how much computation occurs in an interaction in which three variables predict a fourth as done here (the future state of the receiver is predicted from the present state of the receiver and two other senders).

Concerns have been raised about how PID calculates the redundancy term in that it results in an over-estimation of redundancy and consequently, synergy (Bertschinger et al., 2014; Pica et al., 2017). Here, we addressed this concern by demonstrating that an alternate implementation of PID that minimizes redundancy (and thus synergy) nonetheless yields the same pattern of results. Going further, we also used a non-information theoretic approach to identify neurons that likely perform substantial computation by finding those neurons with nonlinear input-output transfer functions, following the

methods of Chichilnisky (2001). This, like our other analyses, showed that the concentration of computation was greater inside of the rich clubs.

Organotypic cultures as a model system

The goal of the present work was to better understand information processing in local cortical networks. To do this, we used a high-density, 512-microelectrode array in combination with organotypic cortical cultures. This approach allowed us to record spiking activity at a temporal resolution (20 kHz; 50 microseconds) that matched typical synaptic delays in cortex (1-3 ms; Mason et al., 1991). The short inter-electrode spacing of 60 microns was within the range of most monosynaptic connections in cortex (Song et al., 2005). The large electrode count allowed us to simultaneously sample hundreds of neurons, revealing complex structures like the rich club. While the cortical layers in organotypic cultures can differ in some respects from those seen *in vivo* (Staal et al., 2011), organotypic cultures nevertheless exhibit very similar synaptic structure and electrophysiological activity to that found *in vivo* (Caeser et al., 1989; Bolz et al., 1990; Götz and Bolz, 1992; Plenz and Aertsen, 1996; Klostermann and Wahle, 1999; Ikegaya et al., 2004; Beggs and Plenz, 2004). The distribution of firing rates in these 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 (Song et al., 2005, reviewed in Buzsáki & Mizuseki, 2014). These features suggested that organotypic cortical cultures serve as a reasonable model system for exploring local cortical networks, while offering an unprecedented combination of large neuron count, high temporal resolution and dense recording sites that cannot currently be matched with *in vivo* preparations.

Relevance to cognitive health

The increased computation in rich clubs described here may help to explain the functional importance of rich clubs as well as the role of rich clubs in healthy neural functioning. Prior work has shown that cognitively debilitating disorders, including Alzheimer's disease, epilepsy, and schizophrenia, are associated with diminished rich club organization (van den Heuvel & Sporns, 2011; van den Heuvel et al., 2013; Braun et al., 2015). An implication of our present findings is that such diminished rich club organization would lead to commensurately diminished neural computation. This could account for the impairments observed in such disorders which include loss of memory, consciousness, or mental cohesiveness. Future studies should make use of *in vivo* methods to explore the relationship between computation and behavior. For example, one could examine how neural computation alters or is altered by experience-dependent (synaptic) weight changes in networks of neurons.

Conclusions

The present work demonstrates, for the first time, that synergy is significantly greater inside, versus outside, of rich clubs. Given this, we conclude that rich clubs not only propagate a large percentage of information within cortical circuits, but are also home to a majority of the circuit-wide computation. We also showed that computation was robustly correlated with information propagation, from which, we infer that computation is driven by information availability. Finally, we found that the ratio of computation-to-propagation was slightly, though significantly reduced in rich clubs, suggesting that cortical circuits, like human-engineered distributed-computing architectures, may face a communication versus computation trade-off. These results substantially increase what is known regarding computation by cortical circuits.

METHODS

To answer the question of whether rich club neurons perform more computation than do non-rich club neurons 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 to 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 using a 512-channel array of 5 µm diameter electrode and 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 Timme et al (2014, 2016). Briefly, weighted effective connections between neurons were established using transfer entropy (TE; Schreiber, 2000). We computed TE at timescales spanning 1.6 – 14 ms to capture neuron interactions at timescales relevant to synaptic transmission. This was discretized into two logarithmically-spaced bins 1.6–6.4 ms and 3.5–14 ms and separate effective networks were constructed for each timescale resulting in two networks per recording (50 networks total). Only significant TE, determined through comparison to the TE values obtained with jittered spike trains (α = 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 percentage of the receiver neuron's capacity that can be accounted for by the transmitting neuron.

Quantifying computation

Computation was operationalized as synergy, as calculated by the partial information decomposition (PID) approach described by Williams and Beer (2010, 2011). PID compares the measured TE between neurons $TE(J \rightarrow I)$ and $TE(K \rightarrow I)$ with the measured multivariate TE between 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 the argMin of $TE(J \rightarrow I)$ and $TE(K \rightarrow I)$. Synergy was then computed via:

$$\begin{gathered}\text{Synergy}(\{J,K\} \rightarrow I) = TE(\{J,K\} \rightarrow I) - TE(J \rightarrow I) - \\ TE(K \rightarrow I) + \text{Redundancy}(\{J,K\} \rightarrow I)\end{gathered} \tag{1}$$

As with TE, synergy was normalized by the total entropy of the receiving neuron. Although there are other methods for calculating synergy (Bertschinger et al., 2014; Pica et al., 2017), we chose this measure because it is capable of detecting linear and nonlinear interactions and it is currently the only measure which can quantify how much synergy occurs in an interaction in which three variables (here, receiver past and pasts of the two transmitters) predict a fourth. Note, we chose not to consider higher order synergy terms, for systems with more than two transmitting neurons, due to the increased computational burden it presented (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.

Alternate methods of quantifying computation

To establish that our results are not unique to our approach for quantifying computation, we implemented two alternate methods. The first also uses PID but sets redundancy to be the smallest possible value. Effectively, in this approach synergy is computed as follows:

$$\text{Synergy}(\{J,K\} \rightarrow I) = \text{argMax}\,[\,TE(\{J,K\} \rightarrow I) - TE(J \rightarrow I) - TE(K \rightarrow I), 0\,] \qquad (2)$$

Consequently, synergy is minimized or set to zero when the sum of $TE(J \rightarrow I)$ and $TE(K \rightarrow I)$ is greater than $TE(\{J,K\} \rightarrow I)$. The second alternate method identifies neurons that perform the most computation as those with non-linear input-output transfer functions. The transfer functions are calculated following the methods of Chichilnisky (2001). Briefly, for each neuron, the pattern of inputs across neurons and time that drive the neuron to spike were estimated using a spike triggered average (STA) of the state of all neurons over a 14 ms window prior to each spike. The strength of input to that neuron over time was then estimated by convolving the STA with the time varying state of the network. Finally, the input-output transfer function was established by computing the probability that the neuron fired at each level of input. We then fit both a line and a sigmoid to the resulting transfer function to extract the summed squared error (SSE) that results from each. We categorized neurons with the lowest SSE from the sigmoid fit relative to the SSE from the linear fit as the neurons that perform the most computation.

Rich club analyses

Weighted rich clubs were identified using a modified version of the rich_club_wd.m function from the Matlab Brain Connectivity toolbox (Rubinov and Sporns, 2010; van den Heuvel and Sporns, 2011), adapted according to Opsahl et al. (2008) to compute weighted rich clubs. To establish the significance of a rich club at a given threshold, we computed the ratio between the observed rich club coefficient and the distribution of those observed when the edges of the network were shuffled. Shuffling was performed according to the methods of Maslov and Sneppen (2002).

To test if synergy was greater for rich club neurons, our first approach randomly selected one of the thresholds–at which the rich club was identified as significant–separately for each network and considered all neurons above that threshold to be in the rich club. Our second approach swept across all possible thresholds, irrespective of the significance of the associated rich club, to assess the influence of varying the 'richness' of the neurons treated as if in the rich club. The third approach asked if the results of the second approach were correlated with the strength of the rich club, as quantified via the

normalized rich club coefficient, to assess if the strength of the observed effects varied with how much stronger the rich club was than expected by chance.

Statistics

All results are reported as medians followed by the 95% bootstrap confidence limits (computed using 10,000 iterations), reported inside of square brackets. Likewise, the median is indicated in figures and the vertical bar reflects 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.

ACKNOWLEDGEMENTS

We thank Olaf Sporns and Randy Beer for helpful comments and discussion. This work was sponsored in part by the Whitehall Foundation (PI:EN), MRI instrument grant (NSF number1429500, PI: JMB), and the Robust Intelligence grant (NSF number 1513779, PI: JMB).

REFERENCES

Adamatzky A, Durand-Lose J (2012) Collision-Based Computing. In: Handbook of Natural Computing (Rozenberg G, Bäck T, Kok JN, ed). pp. 1949-1978. Berlin, Heidelberg: Springer

Beggs JM, Plenz D (2004) Neuronal avalanches are diverse and precise activity patterns that are stable for many hours in cortical slice cultures. J Neurosci 24:5216 –5229

Bertschinger N, Rauh J, Olbrich E, Jost J, Ay N (2014) Quantifying unique information. Entropy, 16(4), 2161-2183.

Bhattacharjee K, Naskar N, Roy S, Das S (2016) A Survey of Cellular Automata: Types, Dynamics, Non-uniformity and Applications. arXiv:1607.02291

Bolz J, Novak N, Götz M, Bonhoeffer T (1990) Formation of target-specific neuronal projections in organotypic slice cultures from rat visual cortex. Nature 346:359 –362

Borst A, Theunissen F (1999) Information theory and neural coding. Nature Neuroscience 2:947–957

Braun U, Muldoon SF, Bassett DS (2015) On human brain networks in health and disease. eLS.

Buzsáki G, Mizuseki K (2014) The log-dynamic brain: how skewed distributions affect network operations. Nat Rev Neurosci 15:264–278.

Caeser M, Bonhoeffer T, Bolz J (1989) Cellular organization and development of slice cultures from rat visual cortex. Exp Brain Res 77:234 –244.

Chichilnisky EJ (2001) A simple white noise analysis of neuronal light responses. Network: Computation in Neural Systems, 12(2), 199-213.

Flecker B, Alford W, Beggs JM, Williams PL, Beer RD (2011) Partial information decomposition as a spatiotemporal filter. Chaos: An Interdisciplinary Journal of Nonlinear Science 21, 3:037104.

Götz M, Bolz J (1992) Formation and preservation of cortical layers in slice cultures. Journal of Neurobiology. 23: 783–802.

Harriger L, van den Heuvel MP, Sporns O (2012) Rich Club Organization of Macaque Cerebral Cortex and Its Role in Network Communication. PLoS ONE 7(9): e46497.

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 304:559 –564

Ikegaya Y, Sasaki T, Ishikawa D, Honma N, Tao K, et al. (2012) Interpyramid spike transmission stabilizes the sparseness of recurrent network activity. Cerebral Cortex 23: 293–304.

Ito S, Yeh FC, Hiolski E, Rydygier P, Gunning DE, Hottowy P, Timme N, Litke AM, Beggs JM (2014) Large-scale, high-resolution multielectrode- array recording depicts functional network differences of cortical and hippocampal cultures. PLoS One 9:e105324.

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

Lefort S, Tomm C, Floyd Sarria JC, Petersen CCH (2009) The excitatory neuronal network of the C2 barrel column in mouse primary somatosensory cortex. Neuron 61: 301–316.

Litke A, Bezayiff N, Chichilnisky E, Cunningham W, Dabrowski W, Grillo A, Grivich M, Grybos P, Hottowy P, Kachiguine S (2004) What does the eye tell the brain? Development of a system for the large-scale recording of retinal output activity. IEEE Trans Nucl Sci 51:1434 –1440.

Lizier JT, Prokopenko M, Zomaya AY (2010) Information modifications and particle collisions in distributed computation. Chaos: An Interdisciplinary Journal of Nonlinear Science 20(3)

Maslov S, Sneppen K (2002) Specificity and Stability in Topology of Protein Networks. Science 03:296, pp. 910-913.

Mason A, Nicoll A, Stratford K (1991) Synaptic transmission between individual pyramidal neurons of the rat visual cortex in vitro. J Neurosci 11:72– 84.

Nigam S, Shimono M, Ito S, Yeh FC, Timme N, Myroshnychenko M, Lapish CC, Tosi Z, Hottowy P, Smith WC, Masmanidis SC, Litke AM, Sporns O, Beggs, JM (2016) Rich-Club Organization in Effective Connectivity among Cortical Neurons. J Neurosci 36(4):670–684

Opsahl T, Colizza V, Panzarasa P, Ramasco JJ (2008) Prominence and control: the weighted rich-club effect. Phys Rev Lett 101:168702.
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 70:893–924.

Rubinov M, Sporns O (2010) Complex network measures of brain connectivity: Uses and interpretations. NeuroImage. 52:1059-69.

Sarkar P (2000) A brief history of cellular automata. ACM Comput Surv 32(1):80–107

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

Song S, Sjöström PJ, Reigl M, Nelson S, Chklovskii DB (2005) Highly non-random features of synaptic connectivity in local cortical circuits. PLoS Biol 3:e68.

Staal JA, Alexander SR, Liu Y, Dickson TD, Vickers JC (2011) Characterization of cortical neuronal and glial alterations during culture of organotypic whole brain slices from neonatal and mature mice. PLOS One 6: e22040

Strong SP, Koberle R, de Ruyter van Steveninck RR, Bialek W (1998) Entropy and information in neural spike trains. Physical Review Letters. 80(1):197–200.

Swadlow HA (1994) Efferent neurons and suspected interneurons in motor cortex of the awake rabbit: axonal properties, sensory receptive fields, and subthreshold synaptic inputs. J Neurophysiol 71:437–453

Tang A, Jackson D, Hobbs J, Chen W, Smith JL, Patel H, Prieto A, Petrusca D, Grivich MI, Sher A, Hottowy P, Dabrowski W, Litke AM, Beggs JM (2008) A maximum entropy model applied to spatial and temporal correlations from cortical networks in vitro. J Neurosci 28:505–518.

Timme NM, Ito S, Myroshnychenko M, Yeh FC, Hiolski E, Litke AM, Beggs JM (2014) Multiplex networks of cortical and hippocampal neurons revealed at different timescales. PLoS One 9: e115764.

Timme NM, Ito S, Myroshnychenko M, Nigam S, Shimono M, Yeh F-C, Hottowy P, Litke AM, Beggs JM (2016) High-Degree Neurons Feed Cortical Computations. PLoS Comput Biol 12(5): e1004858.

van den Heuvel MP, Sporns O (2011) Rich-club organization of the human connectome. J Neurosci 31(44):15775-86.

van den Heuvel MP, Sporns O, Collin G, Scheewe T, Mandl RC, Cahn W, ... & Kahn RS (2013) Abnormal rich club organization and functional brain dynamics in schizophrenia. JAMA psychiatry, 70(8), 783-792.

Wibral M, Priesemann V, Kay JW, Lizier JT, Phillips WA (2017) Partial information decomposition as a unified approach to the specification of neural goal functions. Brain and cognition, 112, 25-38.

Williams PL, Beer RD (2010) Nonnegative decomposition of multivariate information. arXiv: 1004.2515.

Williams PL, Beer RD (2011) Generalized measures of information transfer. arXiv: 1102.1507.