



Published in final edited form as:

Trends Neurosci. 2020 September ; 43(9): 725–737. doi:10.1016/j.tins.2020.07.001.

Principles of corticocortical communication: proposed schemes and design considerations

Adam Kohn^{1,2,3}, Anna I Jasper¹, Joao D Semedo⁴, Evren Gokcen⁴, Christian K Machens⁶, Byron M Yu^{4,5}

¹Dominik Purpura Department of Neuroscience, Albert Einstein College of Medicine, Bronx NY

²Department of Ophthalmology and Visual Sciences, Albert Einstein College of Medicine, Bronx NY

³Department of Systems and Computational Biology, Albert Einstein College of Medicine, Bronx NY

⁴Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh PA

⁵Department of Biomedical Engineering, Carnegie Mellon University, Pittsburgh PA

⁶Champalimaud Research, Champalimaud Centre for the Unknown, Lisbon, Portugal

Abstract

Nearly all brain functions involve routing neural activity among a distributed network of areas. Understanding this routing requires more than a description of inter-areal anatomical connectivity: it requires understanding what controls the flow of signals through inter-areal circuitry, and how this communication might be modulated to allow flexible behavior. Here we review proposals of how communication—particularly between visual cortical areas—is instantiated and modulated, highlighting recent work that offers new perspectives. We suggest transitioning from a focus on assessing changes in the strength of inter-areal interactions, as often seen in studies of inter-areal communication, to a broader consideration of how different signaling schemes might contribute to computation. To this end, we discuss a set of features that might be desirable for a communication scheme.

Keywords

inter-areal signaling; corticocortical communication; visual cortex; feedforward; feedback; pulvinar; synchrony; communication through coherence; communication subspace

Correspondence: adam.kohn@einsteinmed.org [Adam Kohn].

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Corticocortical communication is constrained but not explained by anatomy

Most perceptual, cognitive and motor functions rely on neuronal activity distributed across multiple brain areas [1,2]. These functions require not only the generation of relevant patterns of activity within each area, but also the appropriate communication of activity among areas.

Anatomy provides the basic scaffolding for inter-areal signaling: direct communication between areas requires anatomical connectivity. Yet anatomical connectivity does not fully specify communication. Consider a synaptic connection between two neurons. The existence of this connection makes communication possible, but how the presynaptic neuron will influence its postsynaptic partner will depend on multiple other factors: synaptic strength, the pattern of presynaptic firing (e.g. degree of bursting), the integrative properties of the postsynaptic cell, and so on. Similarly, signaling between neuronal populations could be strongly influenced by the pattern of activity in the source population, the properties of the circuit relaying activity between populations, and the properties of the postsynaptic target network.

The issue of how neuronal populations communicate is a general one, relevant for understanding signaling both within a brain area and between areas. However, there are at least two reasons why inter-areal communication merits special attention. The first is expedience: the physical separation and functional distinctiveness of the source and target networks can be leveraged to facilitate understanding of communication. Second, and more important, inter-areal signaling appears malleable on a moment-to-moment basis, and this flexibility is thought to explain flexible perceptual and cognitive functions [3–9] and motor behaviors [10–14]. Flexible inter-areal signaling is evident as a moment-to-moment change in the degree to which activity in two areas is related [3–8,15].

What, then, are the factors that influence inter-areal communication, and might allow its flexibility? A range of schemes have been suggested, including proposals that focus on the structure of activity in the source network, on ways in which signals may be gated between source and target areas, and on mechanisms within the target area that might modulate the efficacy of inputs received. Most schemes are proposed as general principles for how any two brain areas may communicate, though one proposal suggests that a particular structure—the pulvinar nucleus of the thalamus—may have a special role in regulating corticocortical communication. Here we review these proposals, discuss supporting evidence, and consider their mechanistic underpinnings. We then speculate on conceptual strengths and limitations of the different proposals for information processing, and propose that it is important to consider how different schemes might interact with network computations.

We focus on communication between cortical areas, particularly but not exclusively in the visual cortex, where these issues have been investigated extensively. However, we expect that many of the principles we discuss may be relevant for communication between other structures as well.

Temporal coordination of source population activity

One factor that might strongly influence the drive provided by a source area to a target area is the degree to which the spiking activity of source neurons is temporally coordinated (Figure 1A). Because neurons integrate synaptic inputs over a small time window, synchronous inputs from source neurons are more likely to generate a response in target neurons than asynchronous inputs (Figure 1A; [16–18]). There is extensive experimental evidence that the temporal coordination of spiking activity can be rapidly modulated—by stimulus properties [19,20], task conditions [21], and attentional engagement [22], among others.

Experimental work has provided some evidence that synchronous source activity is associated with a higher probability of spiking in a downstream target network. For instance, synchronous spikes in pairs of thalamic neurons have been shown to be more likely to generate a response in target primary visual cortex (V1) cells than predicted by the sum of input drive provided by each cell separately [23,24]. More recently, neuronal spiking in the input layers of V2 was shown to be associated with epochs of elevated neuronal population synchrony in the output layers of V1 [25]. This association was not evident for V2 neurons in layers that do not receive direct V1 input, suggesting that epochs of enhanced V1 synchrony do not cascade through multiple downstream networks.

The degree to which correlated synaptic input will provide potent drive to the downstream network depends on a multitude of factors, including the temporal precision of the coordination and the number of source neurons involved [26]; the inter-areal architecture, including the divergence and convergence of inter-areal connections [26]; the integrative properties of the downstream neurons [17]; and the degree to which rapid, feed-forward inhibition is recruited in downstream networks [16,25]. A rich computational literature has explored these factors and shown that changing the temporal coordination of activity can have widely varying efficacy in altering the drive provided to a target network [16,17,27]. Thus, the relevance of this scheme may depend on the particular source and target area and the circuitry that links them.

The relevance of synchrony-based schemes for corticocortical communication may also be limited because precise temporal synchrony between cortical neurons is typically weak. Pairs of V1 neurons, for instance, fire roughly 1–5% of their spikes within a few milliseconds of each other [20,28], the timeframe in which postsynaptic summation is likely to occur. In principle, weak pairwise synchrony could be overcome by involving larger pools of neurons—that is, generating synchronous population-wide activity. But empirical evidence to date suggests such events are rare [25,29,30]. Strong population synchrony can be induced by shared locking to stimulus drive [31,32]. But because this form of synchrony is determined by sensory input, it is unclear how it can be deployed to modulate communication in a goal-oriented manner. Thus, whether population temporal coordination plays an important role in modulating corticocortical signaling remains an open question.

Communication through coherence

Perhaps the most extensively studied proposal for modulating inter-areal signaling is ‘communication through coherence’ (Figure 1B; [33]). In this scheme, inter-areal communication is most effective when the phase of gamma oscillations (roughly 30-70 Hz) in the source and target areas are appropriately coordinated (though, conceptually, the proposal can be applied to other oscillation frequencies). Gamma oscillations involve rhythmic fluctuations in inhibition [34,35]. As a result, the efficacy of input to a target area will depend on whether it arrives at the phase of gamma when target neurons are more excitable (‘good’ gamma phase) or relatively unexcitable because they are receiving more inhibition (the ‘bad’ phase).

There is extensive evidence that altering task requirements—usually manipulations of attention—result in changes in gamma oscillations [36] and in their coherence between early visual cortical areas (e.g., [3–5,8,37,38]). In higher visual, motor, frontal and parietal areas, task demands alter inter-areal coherence in other (non-gamma) frequencies [3,5,13,39]. The extensive relevant literature has been reviewed by others [33,40–42].

There are also well-established mechanisms for generating gamma oscillations, involving a rhythmic interplay between excitation and inhibition [34,35]. Establishing the appropriate phase delay between the source and target networks could be accomplished by the oscillatory spiking activity in the source network entraining the target [33].

Despite this evidence in favor of gamma coherence as a modulator of inter-areal communication, a number of groups have argued that the properties of gamma make it unsuitable to serve as the principal mechanism for regulating inter-areal signaling (see [43] for review). Gamma oscillations build slowly in strength [44] and are unstable [45], suggesting they may not be able to modulate inter-areal communication quickly or reliably (but see [46]). Many visual stimuli—like small, low-contrast stimuli [44,47]—induce minimal gamma yet are clearly perceptually distinguishable, suggesting gamma is not needed to relay signals (though it might still modulate the efficacy of signaling in a task-dependent manner). Experimental work has also shown that at least in some cases, the efficacy of input received in a target network is not strongly modulated by local gamma phase [48].

The communication subspace

Recent work has provided a new hypothesis of how inter-areal communication might be modulated: by changing the degree to which population activity patterns in a source area match a communication ‘channel’, which relays those signals to a downstream network [49–51]. Much like a lock-and-key mechanism, signals that match the communication channel are effectively communicated; those that do not remain confined to the source area (Figure 1C).

The core idea for this scheme was developed by Kaufman et al. (2014) to explain the relationship between activity in the motor cortex and the muscles [49]. Motor cortex activity is relayed through the spinal cord to the muscles, where it causes contraction. Yet there is

robust activity in motor cortex during the preparation for movement, which does not generate muscle activity. One longstanding hypothesis is that preparatory activity in motor cortex is gated by either local inhibition or inhibition in spinal circuits, preventing its propagation to the muscles; movement onset, then, would involve opening the gate. However, the evidence for such gating is limited [52].

An alternative explanation for the absence of muscle activity during preparatory periods is provided by a consideration of the mapping from neuronal to muscle activity. Suppose there is a linear mapping such that muscle activity, m , is equal to the summed activity of two neurons (r_1, r_2):

$$m = r_1 + r_2$$

Different combinations of neuronal activity will generate different degrees of muscle activity. Critically, if r_1 and r_2 change in equal and opposite directions, then there is no change in muscle activity, m . Such combinations of neuronal activity can be viewed as falling in the ‘null space’ of the readout relating neuronal to muscle activity. It has been shown that preparatory population activity in motor cortex resides in the null space of a linear model relating motor cortex activity to muscle activity [12,48,50], perhaps explaining why preparatory activity does not cause muscle contraction.

Building on these findings, Semedo et al. (2019) explored which population activity patterns in a cortical source area (V1) were related to population activity in target area V2 [51]. Inter-areal interactions were found to occur through a “communication subspace”, meaning that a small subset of V1 population activity patterns was related to V2 activity. That is, many V1 activity patterns fell in the null space of the mapping from V1 to V2 population activity. Further, the most prominent V1 population patterns were not those most strongly associated with V2 activity, as these prominent patterns were not matched to the communication subspace.

How could a communication subspace be implemented and used to regulate inter-areal communication? In principle, implementing a communication subspace simply requires an appropriate set of synaptic weights between projection and target neurons. Specifically, the weights between networks need to be describable as a linear combination of “basis” or canonical weights. Importantly, the communication subspace does not require that only a subset of source neurons project downstream—a type of trivial anatomical communication subspace, discussed further in Box 1.

With a communication subspace in place, inter-areal communication can be modulated by changing the degree to which population activity patterns in the source area match the communication subspace [51]. There is strong evidence that population activity patterns are highly malleable, by stimulus properties (e.g. [20,53–55]), attention (e.g. [9,56,57]), learning (e.g. [57–59]) and task requirements (e.g. [60–64])—though whether these changes alter inter-areal signaling through a communication subspace is not yet established.

Finally, a concept related to the communication subspace was proposed in a study that explored neuronal population representations in the prefrontal cortex of monkeys performing a color/motion discrimination task [65]. This study showed that the relevant sensory input (e.g. color signals) drove prefrontal representations whereas irrelevant sensory input (motion signals) did not, a form of gating or input selection. Modeling suggests this gating could occur via a context signal that serves to align relevant inputs (but not irrelevant ones) with low-dimensional local network dynamics. Thus, the alignment of inputs with target population dynamics could be a mechanism by which a target network selects inputs from different source areas or networks, instantiating flexible inter-areal signaling.

The pulvinar

All of the proposals described above focus on how the efficacy of direct communication between cortical areas might be modulated. But not all communication between cortical areas is direct. At least some signals between areas may be relayed through the pulvinar nucleus of the thalamus (Figure 1D). The cortical projections to the pulvinar arise in the deep layers of a source area; pulvinar projections to target areas terminate in their superficial layers. In addition, the pulvinar sends a projection back to the source area from which it receives input [66,67].

The pulvinar has been proposed to function as an intermediary, relaying signals received from a source area onto a target. One possibility is that the cortico-pulvino-cortical pathway operates in parallel with the direct corticocortical pathways, but is specialized in providing target cortical areas with an efference copy of the signals relayed to subcortical structures by the source layer 5 neurons [66,68]. The gain of the pulvinar relay might be modulated by inhibitory circuits in the thalamus (e.g. thalamic reticular nucleus; [69]). Others have argued, however, that the cortico-pulvinar-cortical pathway is unlikely to function as an adjustable relay because: (1) there is broad convergence of cortical inputs onto each pulvinar cell, so that cortical activity from the source area is likely reorganized and transformed in the thalamus; and (2) the projection to the target area is diffuse, so signals are relayed there in a non-specific manner [67].

An alternative proposal is that the pulvinar regulates the efficacy of direct corticocortical communication [67,70,71]. The pulvinar might do so by modulating the excitability of projection neurons in a source cortical area. Several studies have shown that pulvinar inactivation can have a dramatic effect on activity in the source area with which it is reciprocally connected [8,71; see also [72]). The pulvinar might also modulate inter-areal communication by affecting cortical oscillations or synchrony in a source or target area (depending on attentional state, [5]) which, as discussed previously, has been associated with altered corticocortical signaling.

In summary, the pulvinar might play an important role in regulating corticocortical signaling, either by functioning as a flexible relay or by modulating direct corticocortical signal transmission. Given that pulvinar lesions can have profound behavioral effects (e.g. visuospatial hemineglect; [67]), more work on the role of this structure in inter-areal signaling is needed.

Feedforward vs feedback signaling

We have discussed inter-areal signaling in a generic manner, without distinguishing whether signaling is occurring in a feedforward (e.g. from lower to higher visual cortex) or feedback (from higher to lower) manner. In the visual system, feedforward connections are thought to generate new receptive field properties in downstream networks (e.g., [73–75]); feedback connections have been linked to a more diverse set of functions, including providing spatial contextual information [76], contributing to perceptual learning [77], and relaying beliefs and top-down predictions about the state of the sensory world [78,79].

Feedforward and feedback connections arise and terminate in distinct cortical layers [80–82]. In the visual cortex, these connections also differ in their precision. Feedforward connections project to a spatially-circumscribed portion of the target area, linking neurons that represent similar regions of visual space [80]. Feedback projections are more spatially diffuse within the target area, and these axons also branch extensively on their way from the source area (so that a source neuron will project to multiple lower cortical areas; [81]). There is also a marked asymmetry in the efficacy of feedforward and feedback pathways. Silencing lower areas usually strongly reduces activity in higher areas, suggesting feedforward inputs are ‘driving’ inputs [83], at least between V1 and higher visual areas [84,85]. In contrast, silencing higher cortex has more subtle effects on responses in lower cortex, reducing responses for some stimuli but not others [86–88]. As a result, feedback connections are considered ‘modulatory’ [83], though they equal feedforward connections in number [80,89].

Given these marked differences in their properties, inter-areal feedforward and feedback communication may operate differently. For instance, feedforward and feedback interactions may involve different oscillation frequencies [90,91]. Gamma oscillations appear to propagate from lower cortical areas to higher (i.e. the oscillations in the lower area lead those in the higher area), whereas those in low frequency bands, like alpha and beta, have the opposite phase lag relationship. However, given the complexity of population codes, it seems improbable that the efficacy of an entire corticocortical pathway can be summarized by a single summary statistic—the power in a certain frequency range. One alternative scenario is that feedforward (e.g., V1 leading V2) and feedback (V2 leading V1) interactions operate through distinct communication subspaces [92,93].

Design considerations

Work in the preceding two decades has given rise to several alternative (and not mutually exclusive) views of how inter-areal communication might be instantiated and flexibly modulated. What might be a fruitful approach for navigating these proposals?

We propose that rather than simply characterizing interaction strength under different task conditions, as is sometimes the case in studies of inter-areal communication, it might be productive to develop a theoretically- or computationally-driven perspective of how different schemes might contribute to function. Part of such an effort should involve assessing the strengths and limitations of different proposed schemes from a functional perspective. A set

of functional considerations is articulated in Box 2 and Figure 2. The list provided is not meant to be exhaustive, nor do we presume any insight into which considerations are most relevant for brain function. However, we suggest that considering issues such as those provided in Box 2, and discussed further below, may help elucidate the functional relevance of different proposed schemes. These considerations may also guide the search for alternative schemes not yet discovered.

A central, often overlooked consideration is that target networks transform the signals they receive from a source area; that is, they perform computations on those inputs. The focus of studies on inter-areal signaling has often been on how the output of area A might be successfully (or unsuccessfully) relayed to B [16,17,94], without much consideration of how the signaling scheme might constrain or influence downstream computation [95]. This shortcoming is perhaps most easily illustrated for synchrony-based schemes for propagating activity through multiple layers of a hierarchical network. Successful communication in this literature is typically defined as the ability to propagate population activity patterns introduced in the first layer effectively to the deep layers of the network [17,26,94]. But if the deep layers of the network produce an output that is identical to its input, it has performed no computation. Though this issue is easily illustrated in the context of synchrony-based schemes, it applies broadly. Networks may perform a range of different computations, such as transforming sensory representations (i.e. creating new receptive field structure; [73–75]), integrating sensory evidence [96], performing predictions [97], maintaining signals for working memory [98,99], marginalizing over nuisance variables [100], and many others. It will be critical to understand how different schemes interact with networks designed to perform these different computations.

A second widely-neglected consideration is the biological feasibility of implementation for flexible, task-directed communication. All proposed schemes have plausible mechanistic underpinnings, in the sense that the requisite phenomena have been observed experimentally (e.g. existence of gamma oscillations, population synchrony, or a communication subspace). But it is much less clear how the relevant modulation of activity might be recruited in a goal-directed manner. For instance, if communication is determined by the alignment of population activity pattern with a communication subspace, how will the structure of that activity be guided so that it yields the desired pattern of communication? Similarly, if communication is modulated by synchrony, how will the task-relevant source neurons be coordinated to enhance the drive they provide to the target network?

In addition to these shared issues, there are design considerations which some schemes may be better suited to address than others. For instance, allowing for high bandwidth communication between areas—the ability to relay a great deal of information in a short time—may be desirable. Oscillation-based schemes may be limited in their temporal bandwidth because oscillations modulate communication by establishing preferential epochs during which signals are effectively relayed and thus also establish epochs when communication is less effective (the ‘bad’ gamma phase; [33]). This may limit information flow, since communication can only occur in discrete epochs. In contrast, mechanisms like the communication subspace involve patterns of activity across neurons rather than time, and so may not be limited in this way [51].

A final example: inter-areal communication is often considered as a problem of relaying activity from area A to B, a formulation we have relied on throughout this review. But most brain functions involve activity distributed across many distinct areas and subnetworks. The control of signaling in such a distributed network may require schemes that can selectively modulate the efficacy of many distinct inputs to a target network. Such scalability might be challenging for oscillation-based schemes because of the difficulty of establishing the correct phase relationship among multiple areas, when communication between each pairing of areas involves a different temporal conduction delay (due to physical proximity; [43]; but see [46]). Scalability may also be challenging if different inputs to the target area need to be modulated independently, since this would require maintaining distinct oscillations in the target network that don't interfere with each other (but see [101]).

Concluding Remarks

We have reviewed a number of proposals for how inter-areal communication might be instantiated and flexibly modulated. We summarized key relevant supporting evidence for each of the proposals, and articulated some of their potential strengths and limitations for supporting computation.

We note that existing proposals need not be mutually exclusive. For instance, the degree to which synchronous activity is effective in driving downstream activity might depend on its timing relative to ongoing oscillations [27]. Or, the efficacy with which signals are routed through a communication subspace may depend on their fine temporal structure or on the phase of ongoing oscillations in the target network. More generally, inter-areal communication might involve a mixture of mechanisms in the source area, the target area, and in intermediate structures like the pulvinar. And this mixture may depend on cortical regions (e.g. frontal vs. occipital), the particular pair of areas considered, or the direction of signal flow (feedforward vs feedback). The possibility of a mixture of inter-areal communication schemes is reinforced by the vast differences—up to 5 orders of magnitude—in the degree of anatomical connectivity between areas that are considered connected [102,103].

Clearly, our understanding of inter-areal communication has advanced, but much remains poorly understood or unknown (see Outstanding Questions). We would argue that a key need is to understand better the relationship between neuronal population spiking activity in different areas (rather than surrogate signals, like the local field potential), because these are the signals that encode information and are actively propagated between networks [104]. Understanding the inter-areal interactions of neuronal population spiking responses will, in turn, require the development of new analytic approaches [105,106]. In this regard, the study of inter-areal signaling will likely offer a fruitful way to advance understanding of neuronal population coding more generally. By elucidating how different aspects of population responses affect propagation and downstream computation, we stand to further our understanding of cortical function more generally.

Acknowledgements:

We thank members of the Batista, Chase, and Yu labs for helpful comments. This work was supported by Simons Collaboration on the Global Brain 543009 (C.K.M.), 542999 (A.K.), 543065 (B.M.Y.), 364994 (A.K., B.M.Y.), NIH U01 NS094288 (C.K.M.), NIH R01 EY028626 (A.K.), NIH R01 HD071686 (B.M.Y.), NIH CRCNS R01 NS105318 (B.M.Y.), NSF NCS BCS 1533672 and 1734916 (B.M.Y.), NIH CRCNS R01 MH118929 (B.M.Y.), and NIH R01 EB026953 (B.M.Y.).

References cited

1. Heekeren HR et al. (2008) The neural systems that mediate human perceptual decision making. *Nat Rev Neurosci.* 9: 467–479. [PubMed: 18464792]
2. Steinmetz NA et al. (2019) Distributed coding of choice, action and engagement across the mouse brain. *Nature* 576: 266–273. [PubMed: 31776518]
3. Buschman TJ and Miller EK (2007) Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. *Science* 315: 1860–1862. [PubMed: 17395832]
4. Bosman CA et al. (2012) Attentional stimulus selection through selective synchronization between monkey visual areas. *Neuron* 75: 875–888. [PubMed: 22958827]
5. Saalman YB et al. (2012) The pulvinar regulates information transmission between cortical areas based on attention demands. *Science* 337: 753–756. [PubMed: 22879517]
6. Salazar RF et al. (2012) Content-specific fronto-parietal synchronization during visual working memory. *Science* 338: 1097–1100. [PubMed: 23118014]
7. Ruff DA and Cohen MR (2016) Attention increases spike count correlations between visual cortical areas. *J Neurosci.* 36: 7523–7534. [PubMed: 27413161]
8. Zhou H et al. (2016) Pulvinar-cortex interactions in vision and attention. *Neuron* 89: 209–220. [PubMed: 26748092]
9. Ruff DA and Cohen MR (2019) Simultaneous multi-area recordings suggest that attention improves performance by reshaping stimulus representations. *Nat Neurosci.* 22: 1669–1676. [PubMed: 31477898]
10. Chen JL et al. (2016) Long-range population dynamics of anatomically defined neocortical networks. *Elife* 5: e14679. [PubMed: 27218452]
11. Arce-McShane F.I. et al. (2016) Primary motor and sensory cortical areas communicate via spatiotemporally coordinated networks at multiple frequencies. *Proc Natl Acad Sci* 113: 5083–5088. [PubMed: 27091982]
12. Perich MG et al. (2018) A neural population mechanism for rapid learning. *Neuron* 100: 964–976. [PubMed: 30344047]
13. Wong YT et al. (2016) Coherent neuronal ensembles are rapidly recruited when making a look-reach decision. *Nat Neurosci.* 19: 327–334. [PubMed: 26752158]
14. Ames KC and Churchland MM (2019) Motor cortex signals for each arm are mixed across hemispheres and neurons yet partitioned within the population response. *eLife.* 8: e46159. [PubMed: 31596230]
15. Campo AT et al. (2015) Task-driven intra- and interarea communications in primate cerebral cortex. *Proc Natl Acad Sci* 112: 4761–4766. [PubMed: 25825731]
16. Salinas E and Sejnowski TJ (2001) Correlated neuronal activity and the flow of neural information. *Nat Rev Neurosci.* 2: 539–550. [PubMed: 11483997]
17. Kumar A et al. (2010) Spiking activity propagation in neuronal networks: reconciling different perspectives on neural coding. *Nat Rev Neurosci.* 11: 615–627. [PubMed: 20725095]
18. Wang HP et al. (2010) Synchrony of thalamocortical inputs maximizes cortical reliability. *Science* 328: 106–109. [PubMed: 20360111]
19. Gray CM et al. (1989) Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature* 338: 334–337. [PubMed: 2922061]
20. Kohn A and Smith MA (2005) Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *J Neurosci.* 25: 3661–3673. [PubMed: 15814797]

21. Riehle A et al. (1997) Spike synchronization and rate modulation differentially involved in motor cortical function. *Science* 278: 1950–1953. [PubMed: 9395398]
22. Engel AK et al. (2001) Dynamic predictions: oscillations and synchrony in top-down processing. *Nat Rev Neurosci.* 2: 704–716. [PubMed: 11584308]
23. Alonso JM et al. (1996) Precisely correlated firing in cells of the lateral geniculate nucleus. *Nature* 383: 815–819. [PubMed: 8893005]
24. Bruno RM and Sakmann B (2006) Cortex is driven by weak but synchronously active thalamocortical synapses. *Science* 312: 1622–1627. [PubMed: 16778049]
25. Zandvakili A and Kohn A (2015) Coordinated neuronal activity enhances corticocortical communication. *Neuron* 87: 827–839. [PubMed: 26291164]
26. Diesmann M et al. (1999) Stable propagation of synchronous spiking in cortical neural networks. *Nature* 402: 529–533. [PubMed: 10591212]
27. Hahn G et al. (2019) Portraits of communication in neuronal networks. *Nat Rev Neurosci.* 20: 117–127. [PubMed: 30552403]
28. Bair W et al. (2001) Correlated firing in macaque visual area MT: time scales and relationship to behavior. *J Neurosci.* 21: 1676–1697. [PubMed: 11222658]
29. Schneidman E et al. (2006) Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature* 440:1007–1012. [PubMed: 16625187]
30. Ohiorhenuan IE et al. (2010) Sparse coding and high-order correlations in fine-scale cortical networks. *Nature* 466: 617–621. [PubMed: 20601940]
31. Bair W and Koch C (1996) Temporal precision of spike trains in extrastriate cortex of the behaving macaque monkey. *Neural Comput.* 8: 1185–1202. [PubMed: 8768391]
32. Butts DA et al. (2007) Temporal precision in the neural code and the timescales of natural vision. *Nature* 449: 92–95. [PubMed: 17805296]
33. Fries P (2009) Neuronal gamma-band synchronization as a fundamental process in cortical computation. *Ann Rev Neurosci.* 32: 209–224. [PubMed: 19400723]
34. Tiesinga P and Sejnowski TJ (2009) Cortical enlightenment: are attentional gamma oscillations driven by ING or PING? *Neuron* 63: 727–732. [PubMed: 19778503]
35. Buzsaki G and Wang XJ (2012) Mechanisms of gamma oscillations. *Annu Rev Neurosci.* 35: 203–225. [PubMed: 22443509]
36. Fries P et al. (2001) Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* 291: 1560–1563. [PubMed: 11222864]
37. Gregoriou GG et al. (2009) High-frequency, long-range coupling between prefrontal and visual cortex during attention. *Science* 324: 1207–1210. [PubMed: 19478185]
38. Rohenkohl G et al. (2018) Gamma Synchronization between V1 and V4 Improves Behavioral Performance. *Neuron* 100: 953–963. [PubMed: 30318415]
39. Rubino D et al. (2006) Propagating waves mediate information transfer in the motor cortex. *Nat Neurosci.* 9: 1549–1557. [PubMed: 17115042]
40. Buzsaki G (2006) *Rhythms of the brain.* Oxford University Press.
41. Fries P (2015) Rhythms for Cognition: Communication through Coherence. *Neuron* 88: 220–235. [PubMed: 26447583]
42. Wang XJ (2010) Neurophysiological and computational principles of cortical rhythms in cognition. *Physiol Rev.* 90: 1195–1268. [PubMed: 20664082]
43. Ray S and Maunsell JH (2015) Do gamma oscillations play a role in cerebral cortex? *Trends Cogn Sci.* 19: 78–85. [PubMed: 25555444]
44. Jia X et al. (2011) Stimulus selectivity and spatial coherence of gamma components of the local field potential. *J Neurosci.* 31: 9390–9403. [PubMed: 21697389]
45. Burns SP et al. (2011) Is gamma-band activity in the local field potential of V1 cortex a "clock" or filtered noise? *J Neurosci.* 31: 9658–9664. [PubMed: 21715631]
46. Palmigiano A et al. (2017) Flexible information routing by transient synchrony. *Nat Neurosci.* 20: 1014–1022. [PubMed: 28530664]
47. Ray S and Maunsell JH (2010) Differences in gamma frequencies across visual cortex restrict their possible use in computation. *Neuron* 67: 885–896. [PubMed: 20826318]

48. Jia X et al. (2013) Gamma and the coordination of spiking activity in early visual cortex. *Neuron* 77: 762–774. [PubMed: 23439127]
49. Kaufman MT et al. (2014) Cortical activity in the null space: permitting preparation without movement. *Nat Neurosci.* 17: 440–448. [PubMed: 24487233]
50. Elsayed GF et al. (2016) Reorganization between preparatory and movement population responses in motor cortex. *Nat Commun.* 7: 13239. [PubMed: 27807345]
51. Semedo JD et al. (2019) Cortical areas interact through a communication subspace. *Neuron* 102: 249–259. [PubMed: 30770252]
52. Kaufman MT et al. (2013) The roles of monkey M1 neuron classes in movement preparation and execution. *J Neurophysiol.* 110: 817–825. [PubMed: 23699057]
53. Churchland MM et al. (2010) Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat Neurosci.* 13: 369–378. [PubMed: 20173745]
54. Tan AY et al. (2014) Sensory stimulation shifts visual cortex from synchronous to asynchronous states. *Nature* 509: 226–229. [PubMed: 24695217]
55. Cowley BR et al. (2016) Stimulus-driven population activity patterns in macaque primary visual cortex. *PLoS Comput Biol.* 12(12): e1005185. [PubMed: 27935935]
56. Snyder AC et al. (2018) Distinct population codes for attention in the absence and presence of visual stimulation. *Nat Commun.* 9(1): 4382. [PubMed: 30348942]
57. Ni AM et al. (2018) Learning and attention reveal a general relationship between population activity and behavior. *Science* 359: 463–465. [PubMed: 29371470]
58. Jeanne JM et al. (2013) Associative learning enhances population coding by inverting interneuronal correlation patterns. *Neuron* 78: 352–363. [PubMed: 23622067]
59. Oby ER et al. (2019) New neural activity patterns emerge with long-term learning. *Proc Natl Acad Sci* 116: 15210–15215. [PubMed: 31182595]
60. Raposo D et al. (2014) A category-free neural population supports evolving demands during decision-making. *Nature Neurosci.* 17: 1784–1792. [PubMed: 25383902]
61. Kobak D et al. (2016) Demixed principal component analysis of neural population data. *Elife* 5: e10989. [PubMed: 27067378]
62. Bondy AG et al. (2018) Feedback determines the structure of correlated variability in primary visual cortex. *Nat Neurosci.* 21: 598–606. [PubMed: 29483663]
63. Bagur S et al. (2018) Go/No-Go task engagement enhances population representation of target stimuli in primary auditory cortex. *Nat Commun.* 9(1): 2529. [PubMed: 29955046]
64. Gallego JA et al. (2018) Cortical population activity within a preserved neural manifold underlies multiple motor behaviors. *Nat Commun.* 9: 1–3. [PubMed: 29317637]
65. Mante V et al. (2013) Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* 503: 78–84. [PubMed: 24201281]
66. Sherman SM (2016) Thalamus plays a central role in ongoing cortical functioning. *Nat Neurosci.* 19: 533–541. [PubMed: 27021938]
67. Halassa MM and Kastner S (2017) Thalamic functions in distributed cognitive control. *Nat Neurosci.* 20: 1669–1679. [PubMed: 29184210]
68. Theyel BB et al. (2010) The corticothalamocortical circuit drives higher-order cortex in the mouse. *Nat Neurosci.* 13: 84–88. [PubMed: 19966840]
69. Jaramillo J et al. (2019) Engagement of pulvino-cortical feedforward and feedback pathways in cognitive computations. *Neuron* 101: 321–336. [PubMed: 30553546]
70. Olshausen BA et al. (1993) A neurobiological model of visual attention and invariant pattern recognition based on dynamic routing of information. *J Neurosci.* 13: 4700–4719. [PubMed: 8229193]
71. Purushothaman G et al. (2012) Gating and control of primary visual cortex by pulvinar. *Nat Neurosci.* 15: 905–912. [PubMed: 22561455]
72. Guo ZV et al. (2017) Maintenance of persistent activity in a frontal thalamocortical loop. *Nature* 545: 181–186. [PubMed: 28467817]
73. Hubei DH and Wiesel TN (1962) Receptive fields, binocular interaction and functional architecture in the ca s visual cortex. *J Physiol.* 160: 106–154. [PubMed: 14449617]

74. Rust NC et al. (2006) How MT cells analyze the motion of visual patterns. *Nat Neurosci.* 9: 1421–1431. [PubMed: 17041595]
75. Yamins DL and DiCarlo JJ (2016) Using goal-driven deep learning models to understand sensory cortex. *Nat Neurosci.* 19: 356–365. [PubMed: 26906502]
76. Angelucci A and Bressloff PC (2006) Contribution of feedforward, lateral and feedback connections to the classical receptive field center and extra-classical receptive field surround of primate V1 neurons. *Prog Brain Res.* 154: 93–120. [PubMed: 17010705]
77. Gilbert CD and Li W (2013) Top-down influences on visual processing. *Nat Rev Neurosci.* 14: 350–63. [PubMed: 23595013]
78. Harris KD and Mrsic-Flogel TD (2013) Cortical connectivity and sensory coding. *Nature* 503: 51–58. [PubMed: 24201278]
79. Haefner RM et al. (2016) Perceptual decision-making as probabilistic inference by neural sampling. *Neuron* 90: 649–660. [PubMed: 27146267]
80. Felleman DJ and Van Essen DC (1991) Distributed hierarchical processing in the primate cerebral cortex. *Cereb Cortex* 1: 1–47. [PubMed: 1822724]
81. Salin PA and Bullier J (1995) Corticocortical connections in the visual system: structure and function. *Physiol Rev.* 75: 107–154. [PubMed: 7831395]
82. Harris JA et al. (2019) Hierarchical organization of cortical and thalamic connectivity. *Nature* 575: 195–202. [PubMed: 31666704]
83. Sherman SM and Guillery RW (1998) On the actions that one nerve cell can have on another: distinguishing "drivers" from "modulators". *Proc Natl Acad Sci.* 95: 7121–7126. [PubMed: 9618549]
84. Girard P and Bullier J (1989) Visual activity in area V2 during reversible inactivation of area 17 in the macaque monkey. *J Neurophysiol.* 62: 1287–1302. [PubMed: 2600626]
85. Girard P et al. (1991) Visual activity in areas V3A and V3 during reversible inactivation of area V1 in the macaque monkey. *J Neurophysiol.* 66: 1493–1503. [PubMed: 1765790]
86. Hupé JM et al. (1998) Cortical feedback improves discrimination between figure and background by V1, V2 and V3 neurons. *Nature* 394: 784–787. [PubMed: 9723617]
87. Nassi JJ et al. (2013) Corticocortical feedback contributes to surround suppression in V1 of the alert primate. *J Neurosci.* 33: 8504–8517. [PubMed: 23658187]
88. Pafundo DE et al. (2016) Top-Down-Mediated Facilitation in the Visual Cortex Is Gated by Subcortical Neuromodulation. *J Neurosci.* 36: 2904–2914. [PubMed: 26961946]
89. Markov NT et al. (2013) Cortical high-density counterstream architectures. *Science* 342: 1238406. [PubMed: 24179228]
90. van Kerkoerle T et al. (2014) Alpha and gamma oscillations characterize feedback and feedforward processing in monkey visual cortex. *Proc Natl Acad Sci.* 111: 14332–14341. [PubMed: 25205811]
91. Bastos AM et al. (2015) Visual areas exert feedforward and feedback influences through distinct frequency channels. *Neuron* 85: 390–401. [PubMed: 25556836]
92. Semedo JD et al. (2019b) Temporal dynamics of inter-area neuronal population interactions. *COSYNE*.
93. Gokcen E. Dissecting feedforward and feedback interactions between populations of neurons; *COSYNE*; 2020.
94. Litvak V et al. (2003) On the transmission of rate code in long feedforward networks with excitatory-inhibitory balance. *J. Neurosci* 23: 3006–3015. [PubMed: 12684488]
95. Shadlen MN and Newsome WT (1998) The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J. Neurosci* 18: 3870–3896. [PubMed: 9570816]
96. Wang XJ (2008) Decision making in recurrent neuronal circuits. *Neuron* 60: 215–234. [PubMed: 18957215]
97. Heeger DJ (2017) Theory of cortical function. *Proc Natl Acad Sci.* 114: 1773–1782. [PubMed: 28167793]
98. Machens CK et al. (2005) Flexible control of mutual inhibition: a neural model of two-interval discrimination. *Science* 307: 1121–1124. [PubMed: 15718474]

99. Barak O and Tsodyks M (2014) Working models of working memory. *Curr Opin Neurobiol.* 25: 20–24. [PubMed: 24709596]
100. Beck JM et al. (2011) Marginalization in neural circuits with divisive normalization. *J Neurosci.* 31: 15310–15319. [PubMed: 22031877]
101. Akam T and Kullmann DM (2010) Oscillations and filtering networks support flexible routing of information. *Neuron* 67: 308–320. [PubMed: 20670837]
102. Oh SW et al. (2014) A mesoscale connectome of the mouse brain. *Nature* 508: 207–214 [PubMed: 24695228]
103. Van Essen DC et al. (2019) Cerebral cortical folding, parcellation, and connectivity in humans, nonhuman primates, and mice. *Proc Natl Acad Sci.* 116: 26173–26180.
104. Kohn A et al. (2016) Correlations and neuronal population information. *Annu Rev Neurosci.* 39: 237–256. [PubMed: 27145916]
105. Zavitz E and Price NSC (2019) Understanding sensory information processing through simultaneous multi-area population recordings. *Front Neural Circuits.* 12:115. [PubMed: 30687020]
106. Semedo JD et al. (2020) Statistical methods for dissecting interactions between brain areas. *Curr. Opinion Neurobiology.*
107. Glickfeld LL et al. (2013) Cortico-cortical projections in mouse visual cortex are functionally target specific. *Nat Neurosci.* 16: 219–226. [PubMed: 23292681]
108. Marques T et al. (2018) The functional organization of cortical feedback inputs to primary visual cortex. *Nat Neurosci.* 21: 757–764. [PubMed: 29662217]
109. Yamashita T et al. (2018) Diverse long-range axonal projections of excitatory layer 2/3 neurons in mouse barrel cortex. *Front Neuroanat.* 12:33. doi: 10.3389/fnana.2018.00033. [PubMed: 29765308]
110. Economo MN et al. (2018) Distinct descending motor cortex pathways and their roles in movement. *Nature* 563: 79–84. [PubMed: 30382200]
111. Hwang EJ et al. (2019) Corticostriatal flow of action selection bias. *Neuron* 104: 1126–1140. [PubMed: 31706697]
112. Sincich LC and Horton JC (2003) Independent projection streams from macaque striate cortex to the second visual area and middle temporal area. *J Neurosci.* 23: 5684–5692. [PubMed: 12843271]
113. Rockland KS (2013) Collateral branching of long-distance cortical projections in monkey. *J Comp Neurol.* 521: 4112–4123. [PubMed: 23839719]
114. Han Y et al. (2018) The logic of single-cell projections from visual cortex. *Nature* 556: 51–56. [PubMed: 29590093]

Box 1**Anatomy 2.0: The rise of subnetworks**

Modern experimental techniques have renewed interest in how inter-areal circuitry might inform our understanding of inter-areal communication. More specifically, recent work has revealed that neighboring neurons in a source area can have strikingly different downstream targets. For instance, intermingled subsets of neurons in mouse primary visual cortex (V1) project either to downstream cortical areas PM or AL [107]. Feedback axons from LM to V1 also arise from distinct subsets of neurons, with different functional properties and different projection patterns [108]. Intermingled subnetworks relaying different information to distinct downstream network have also been described in mouse primary somatosensory cortex [109], motor cortex [110], and posterior parietal cortex [111], among others.

While this anatomical knowledge places important constraints on corticocortical communication, there are issues with relying on different functional subnetworks to instantiate flexible inter-areal signaling.

First, anatomical work has shown that individual cortical neurons often project to multiple downstream areas [112–114]. This limits the degree to which selective or flexible communication can be achieved by simply turning on or off different subsets of neurons. We note that selective or flexible communication can be achieved if the synaptic connections of source neurons to multiple downstream targets are arranged to form distinct communication subspaces [51]. That is, with communication subspaces, neurons in a source area can project to multiple downstream targets, and yet selectively communicate with one target or another by changing patterns of activity among those neurons [51]. Thus, the communication subspace is a functional concept, which does not rely on the presence or absence of connections to achieve selectivity, as is the case for anatomical subnetworks.

Second, neurons with distinct target projections often have different functional properties. For instance, V1 neurons projecting to area AL have different spatiotemporal selectivity from their neighbors who project to area PM [107]. Thus, shifting activity from one source subpopulation to another not only alters which target network is being communicated with, but also which information is being exported by the source area. This is distinct from the schemes considered in the main text, in which the same neurons are involved in the communication—and thus the same information can be conveyed—but their signaling efficacy is altered. In essence, instantiating flexible communication by shifting activity between different subpopulations of projection neurons is more akin to changing the areas involved in a function (where the “areas” are groups of neurons within an area) than to changing how a fixed constellation of areas communicates.

Box 2**Desiderata of communication schemes**

To further our understanding of inter-areal communication, a normative perspective may be helpful: What features might we desire from an inter-areal signaling scheme? Of course, these features might not all be accomplished by one scheme, might not be equally important for brain function, and might conflict (in that optimizing for one may compromise another). Yet considering these desired features can help guide understanding of the functional advantages and disadvantages—and perhaps the relevance—of different schemes. Some useful features might be:

Modulation strength:

The scheme should modulate inter-areal signaling strongly enough to modify the function of the downstream network in the desired manner.

Scalability:

The scheme should allow for the modulation of communication between multiple areas. Often schemes are considered as a way to alter signaling between two areas, although most functions will involve coordination among a larger network of areas.

Interference:

The scheme should allow for a downstream network to receive input from one area without contamination by signals provided by another.

Selectivity:

The scheme should allow for a source area to send signals selectively, communicating with one target network but not another.

Reconfiguration speed:

The scheme should allow for rapid reconfiguration of which areas are communicating. If the scheme takes time to instantiate or switch configurations, it may be useful only when demands are relatively constant in time.

Bandwidth:

The scheme should not impose a bottleneck on the flow of information (e.g. bits/s) between the networks it aims to connect, both at any instant or on average across time.

Computation:

The scheme should contribute to, or at least not interfere with, network computation and function. The purpose of inter-areal communication is not simply to relay signals from one area to another. Rather, it is to relay signals such that desired computations, both within and across areas, can be instantiated.

Implementation:

The implementation of the scheme should be biologically plausible. Most proposed schemes have established mechanistic underpinnings. But it often less clear how a

scheme might be implemented to accomplish a specific desired routing of signals. If performing a task requires routing signals from area A to B and not from A to C, how will the necessary configuration of source or target activity be implemented?

Robustness:

The scheme should be robust, so that it is not derailed by unavoidable biological fluctuations, such as response variability (e.g. spike timing variability and synaptic transmission failure) or changes in network architecture (e.g. cell death, other forms of injury or aging, or plasticity in the source or target area). Robustness might also include an ability to modulate signaling in the presence of other communication schemes (i.e. lack of interference between schemes).

Learnable:

The scheme should be learnable during development, and updatable in the fully-developed brain.

Outstanding Questions

There is tremendous diversity in the strength of inter-areal connectivity. Given this diversity, to what degree are there canonical rules of communication? Are feedforward and feedback communication modulated using similar schemes or principles? Are the schemes for modulating inter-areal communication different from those that regulate signal flow between neuronal populations within a cortical area (e.g. different layers)? Are the mechanisms for modulating inter-areal signaling similar across species?

What is the role and relative importance of cortico-pulvino-cortical pathways compared to direct corticocortical pathways in inter-areal signaling?

To what degree does flexible behavior rely on changes in corticocortical communication as opposed to changes in the functions performed by individual brain areas?

Does modulating inter-areal signaling rely on recruiting distinct subsets of neurons for different functional purposes (i.e. distinct subnetworks), or instead rely on changing how a fixed pool of neurons interacts with each other?

Our understanding of corticocortical signaling is hamstrung by limited computational and theoretical frameworks. What are the strengths and limitations of different signaling schemes in supporting specific computations and functions in both source and target networks? What are the functional merits of having modular networks with flexible inter-areal communication?

Ultimately, inter-areal communication must be understood at the level of neuronal population spiking responses, as these are the signals that are actively relayed between areas. Yet we have limited analytical tools to relate sets of population spiking responses to each other. What analytic methods best summarize population activity and allow it to be related to activity in other areas and to behavior?

Highlights

- Corticocortical communication is a fundamental aspect of brain function. Flexible behavior suggests a need for modulating inter-areal signaling from moment to moment.
- Several schemes for modulating corticocortical communication have been proposed. These include altering the structure of activity within a source network, the sensitivity of a target network to the input it receives, or gating signals during the relay between areas.
- We review these schemes, and highlight new proposals which suggest communication may be determined by how source population signals align with inter-areal communication subspaces.
- We propose a set of design considerations for evaluating the relative merits of different communication schemes. When examining inter-areal communication, we suggest moving beyond merely characterizing changes in the strength of inter-areal interactions, to a wider consideration of the computational benefits and limitations of different communication schemes.

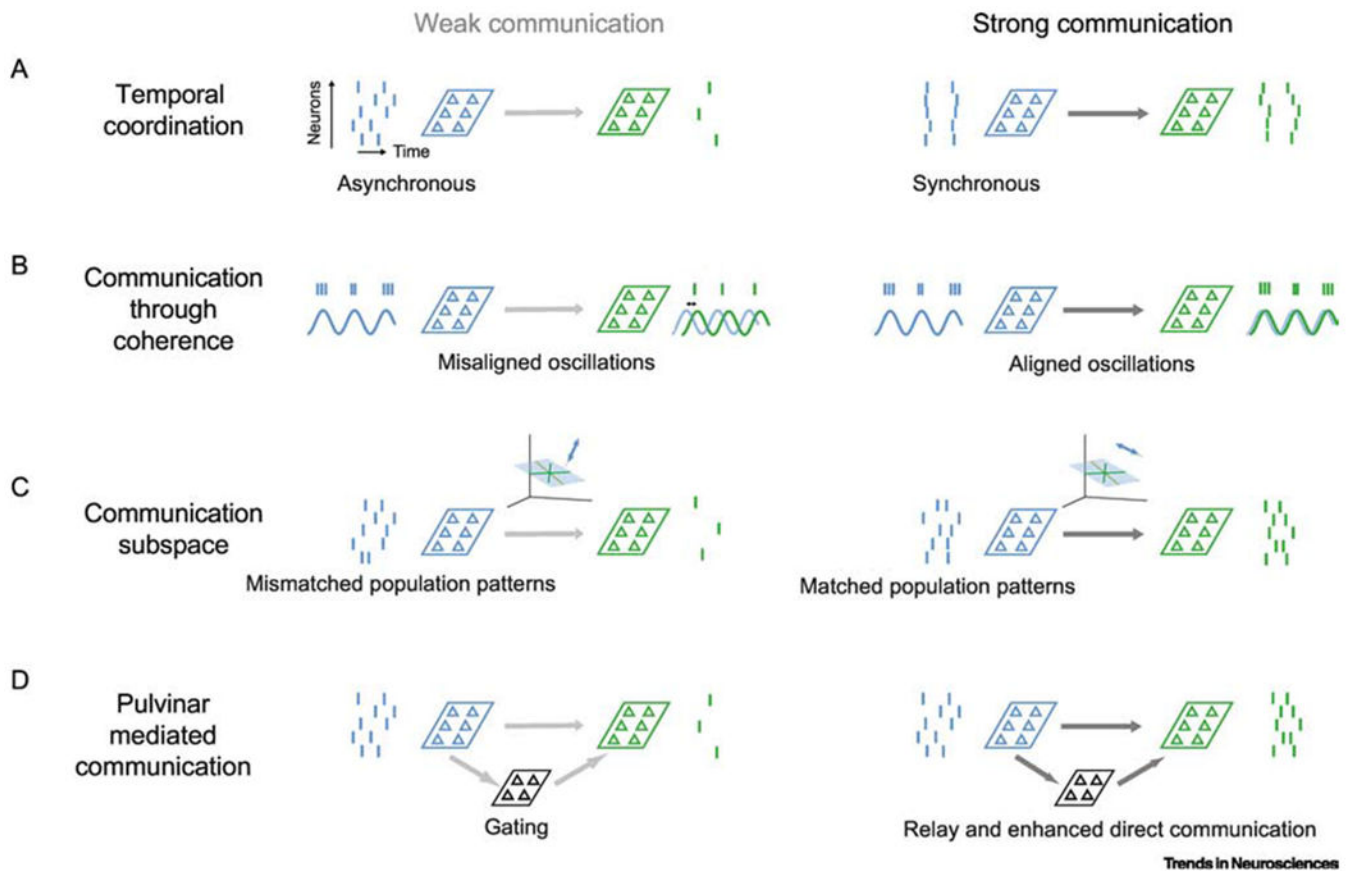
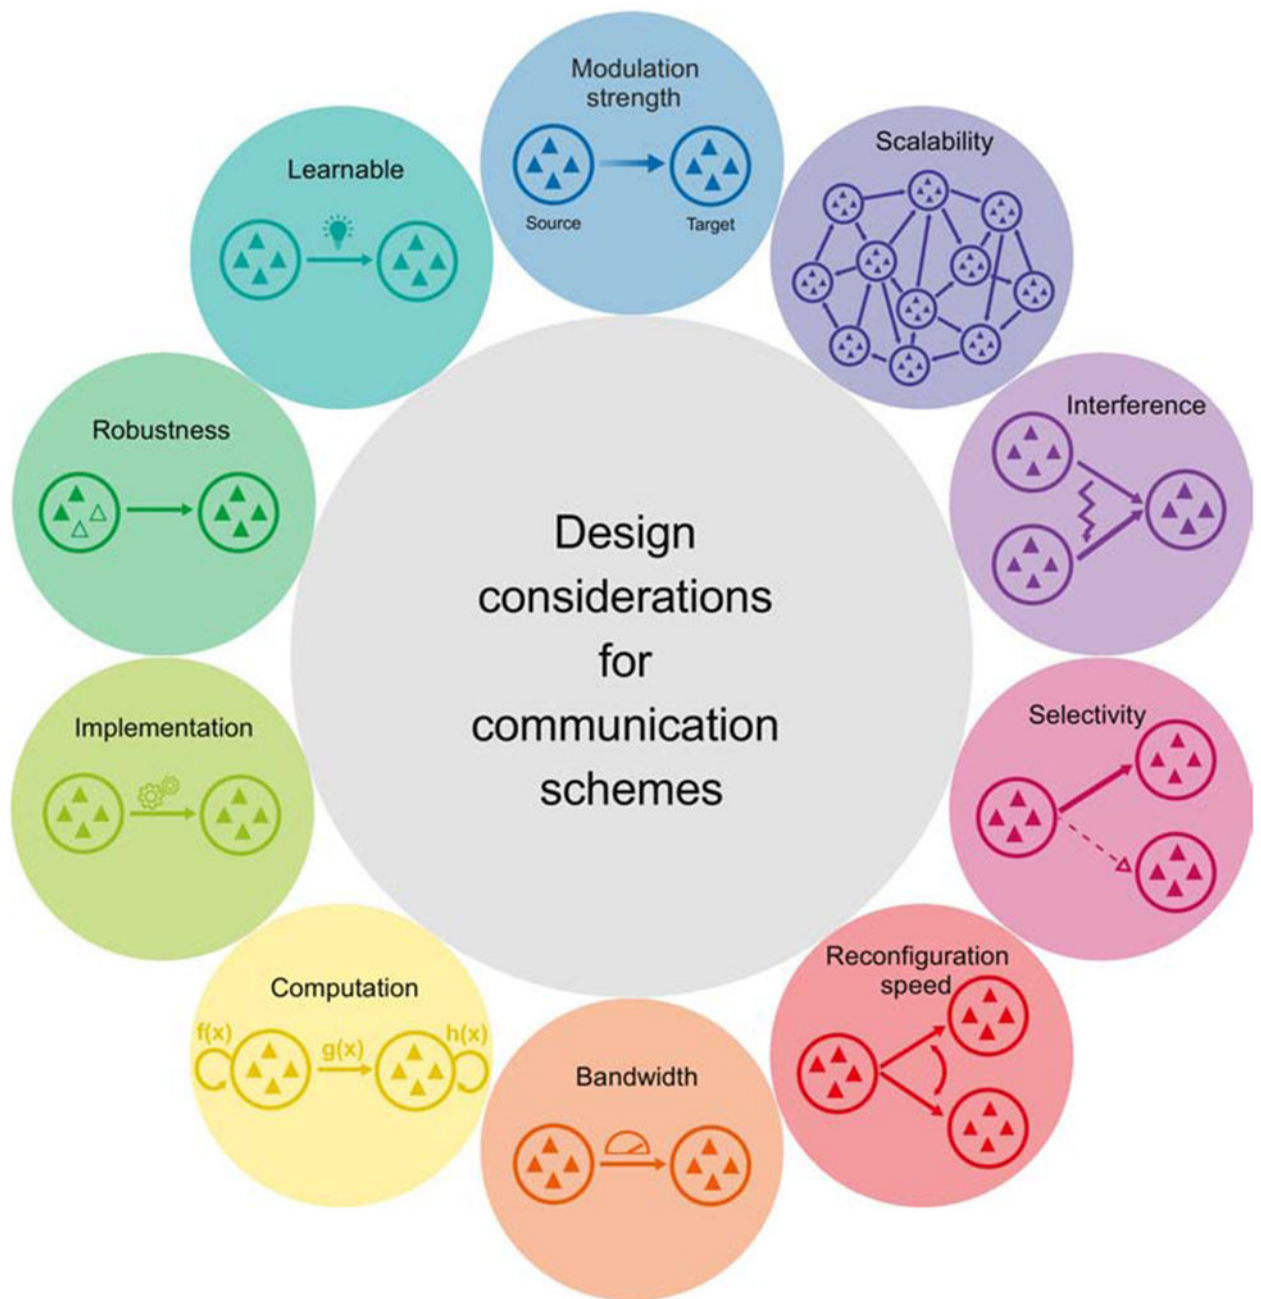


Figure 1:

Schemes for modulating corticocortical communication. (A) Temporal coordination. Left: Weak communication occurs when source population activity (blue) is asynchronous, resulting in little activity in the target network (green). Strong communication (right) occurs when source activity is synchronous. In the schematic on the right, resultant downstream activity is depicted as synchronous as well, as in [26], though it remains unclear whether in the cortex downstream activity driven by input synchrony is also synchronous or simply elevated. (B) Communication through coherence. Left: Weak communication occurs when oscillations in the source and target area are not appropriately phase aligned. Activity in the source area occurs at the peak of a local oscillation. If the phase of this oscillation is misaligned with the oscillation in the target network, there is little activity generated in the target (faint blue line, source oscillation; green line, oscillation in the target network). Right: Strong communication occurs when the two oscillations have appropriate phase offset. (C) Communication subspace. Left: Weak communication occurs when activity in the source population is mismatched to the communication subspace that relates source activity to target activity. In this illustration, the subspace is indicated as a plane in the source population activity space; fluctuations in the source area are orthogonal to the subspace and thus generate no activity in the target network. Right: Strong communication occurs when fluctuations in source activity fall in the communication subspace. (D) Pulvinar mediated communication. Left: Weak communication through both the direct corticocortical pathway and indirect pathway through the pulvinar. Weak communication may involve gating

mechanisms in the thalamus. Right: Strong communication involves better relaying of activity through the pulvinar, and enhanced direct corticocortical communication. Text alignment under the source/targets areas in each panel indicate whether the schemes involves changes in the source network, target network, or in between.



Trends in Neurosciences

Figure 2:

Design principles for inter-areal communication schemes. There are many inter-related design considerations that may distinguish the utility of different proposed communication schemes. These include modulation strength—how strongly the scheme alters communication; scalability—how well the scheme could coordinate signal flow in a distributed network; interference—how well the scheme separates inputs from different upstream areas; selectivity—how well the scheme allows a source area to communicate with one downstream area versus another; reconfiguration speed—how quickly the scheme can

switch which areas are communicating with which; bandwidth—how much information the scheme allows one area to send to another per unit time; computation—how the scheme contributes to and interacts with computations performed within and between areas; implementation—how the scheme could be instantiated to achieve a particular routing of signals; robustness—how well the scheme tolerates disruptions like neuron loss; and learnability—how easily the scheme might be learned from experience). See also Box 2.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript