A computational perspective of the role of Thalamus in cognition

Nima Dehghani[1,2,a)] and Ralf D. Wimmer[3]
[1)] *Department of Physics, Massachusetts Institute of Technology*
[2)] *Center for Brains, Minds and Machines (CBMM), Massachusetts Institute of Technology*
[3)] *Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology*

(Dated: March 5, 2018)

Thalamus has traditionally been considered as only a relay source of cortical inputs, with hierarchically organized cortical circuits serially transforming thalamic signals to cognitively-relevant representations. Given the absence of local excitatory connections within the thalamus, the notion of thalamic 'relay' seemed like a reasonable description over the last several decades. Recent advances in experimental approaches and theory provide a broader perspective on the role of the thalamus in cognitively-relevant cortical computations, and suggest that only a subset of thalamic circuit motifs fit the relay description. Here, we discuss this perspective and highlight the potential role for the thalamus in dynamic selection of cortical representations through a combination of intrinsic thalamic computations and output signals that change cortical network functional parameters. We suggest that through the contextual modulation of cortical computation, thalamus and cortex jointly optimize the information/cost tradeoff in an emergent fashion. We emphasize that coordinated experimental and theoretical efforts will provide a path to understanding the role of the thalamus in cognition, along with an understanding to augment cognitive capacity in health and disease.

Keywords: Thalamo-cortical system, Recurrent Neural Network, Reservoir Computing, Multi-objective Optimization, Cognitive Computing, Artificial Intelligence

Cortico-centric view of perceptual and cognitive processing

Until recently, cognition [in mammalian, bird and reptilian nervous system] has been viewed as a cortico-centric process, with thalamus considered to only play the mere role of a relay system. This classic view, much driven by the visual hierarchical model of the mammalian cortex[24], puts thalamus at the beginning of a feedforward hierarchy. The transmission of information from thalamus to early sensory cortex (V1 in the visual system for example), and the gradual increasing complex representations from V2 to MT/IT and eventually prefrontal cortex (PFC), constitute the core of the perceptual representation under the hierarchical model. A recent comparative study of biologically-plausible Convolutional Neural Networks (CNN) and the visual ventral stream, emphasizes on the feature enrichment through the network of hierarchically interconnected computational modules (layers in the neural network and areas in the ventral stream)[110].

The strictly static feedforward model has since morphed to a dynamic hierarchical model due to the discoveries of the role of feedback from higher cortical areas to lower cortical areas[35]. These dynamic hierarchies are even considered to be favorable for recursive optimizations, where the overall optimization can be achieved by breaking the problem into smaller ones in order to find the optimum for each of these smaller problems. However, it is possible for the recursive optimization not to be confined to just one area of the cortex. It may also be that these smaller optimizations may be solved differently at different cortical areas[58]. This view of cortical computation is also paralleled with the growing use of recurrent neural networks (RNN) that can capture the dynamics of single neurons or neural population in a variety of tasks. As such, RNNs can mimic (a) context-dependent prefrontal response[57] or (b) can reproduce the temporal scaling of neural responses in medial frontal cortex and caudate nucleus[103]. Nonetheless, the main attribute of perception/cognition remains cortico-centric under the umbrella of dynamic hierarchical models or RNN embodiment of cortical cognitive functions. Since the computation that is carried by the system should match the computing elements at the appropriate scale[17] a mismatch between these presumed computational systems and the underlying circuitry becomes vividly apparent. Specifically, let's note that a) associative cortex (and not just sensory cortex) receives thalamic input, b) certain thalamic territories receive their primary (driving) input from the cortex, rather than sensory periphery, some of which are likely to be highly convergent on a single thalamic cell level and c) cortex is demarcated by local excitatory connectivity (while the thalamus is devoid of this feature). These points should prompt us to reconsider this cortico-centric view of computation and indicates that it need to be extended to a thalamocortical one. As such, the thalamic contribution to cognitive computations can be twofold: First, they can modulate the ongoing cortical activity and second, they might be pure thalamic computations. The nature of such intrinsic thalamic computations would be determined by both the types of inputs a thalamic neuron receives as well as how these inputs interact (both linearly and non-linearly). As we will discuss in the next several sections, the combination of these features allows the thalamus to compute a variety of signals associated with ongoing cortical states, which can be

[a)] Electronic mail: nima.dehghani@mit.edu

more easily used to dynamically adjust cortical computations through real-time modulation of functional cortical architecture. Further, we propose that thalamus may provide a contextual signal that reorganizes functional connectivity in frontal cortices in response to contextual changes. We suggest that the unique cognitive capability of the thalamo-cortical system is tightly bound to parallel processing and contextual modulation that are enabled by the diversity of computing nodes (including thalamic and cortical structures) and complexity of the computing architecture (Fig. 1). We will start with a brief overview of the thalamic architecture, followed by experimental evidence and a computational perspective of the thalamic role in contextual cognitive computation .

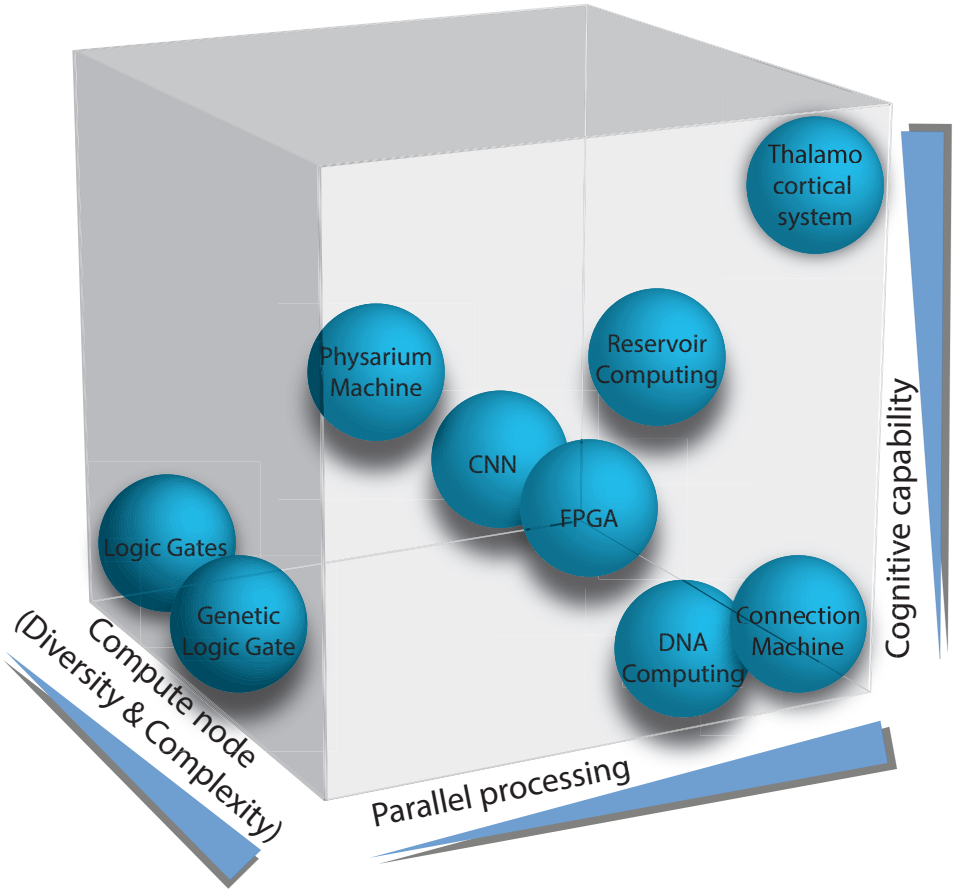

Figure 1. **Cognitive computing morphospace:** Morphospace of a few example biological and synthetic computing engines in a multidimensional layout. The thalamo-cortical system standout as a unique system with high cognitive capacity, massive parallel processing and extreme diversity of the computing nodes. Other computing systems occupy less desirable domains of this morphospace. *Logic Gates*: NAND, NOR; *Genetic Logic Gate*: Synthetic biology adaptation of logic gate; *FPGA*: field-programmable gate array (configurable integrated circuit); *CNN*: Convolutional Neural Network; *Physarium Machine*: programmable amorphous biological computer experimentally implemented in slime mould. *Reservoir Computing*: A reservoir of recurrent neurons that dynamically change their activity to nonlinearly map the input to a new space; *DNA computing*: A computing paradigm where many different DNA molecules are used to perform large number of logical computations in parallel; *Connection Machine*: the first commercial supercomputer designed for problems of Artificial Intelligence (AI) using hardware enabled parallel processing.

Thalamic architecture: anatomical and functional features

Among all brain structures, the forebrain is likely to be most related to cognitive capacity, given that forebrain expansion positively correlates (and perhaps defines) cognitive expansion throughout evolution[46,74,75,87]. Forebrain is composed of diencephalon (hypothalamus, epithalamus and subthalamus) along with the telencephalon

(cortex and basal ganglia)[10]. Analogous structures exist throughout the vertebrate lineage[10,15]. Both the cortex and basal ganglia are demarcated by local recurrent connections, with the main difference being that cortical excitatory recurrence gives rise to long-lasting attractor dynamics thought to underlie many cognitive processes such as memory and decision making[104,107]. The basal ganglia are inhibitory structures, with local recurrence thought to implement selection through lateral inhibition rather than long lasting attractor state[63,105,106]. Within this framework, what does the thalamus do? The answer depends on 'the thalamus' in question. Traditionally, thalamic nuclei (see Fig. 2) are defined as collections of neurons that are segregated by gross features such as white matter tracts and various types of tissue staining procedures[42]. This gross anatomical classification has been equated with a functional one, where individual thalamic nuclei giving rise to a set of defined functions[42,43]. More recent fine anatomical studies challenge this notion, showing that within individual nuclei, single cell input/output connectivity patterns are quite variable (see below).

This thalamic input diversity sets the bases for non-unitary computational role of thalamus in integrating information. Thalamus lacks lateral excitatory connections and rather receives inputs from other subcortical structures and/or cortex. In fact, a major feature of forebrain expansion across evolution is the invasion of the thalamus by cortical inputs[30,85]. Most (90-95%) afferents to the relay nuclei are not from the sensory organs[43,93]. The majority of these inputs originate from intrinsic (local GABAergic neurons, projections from reticular nucleus, thalamic interneurons–mostly absent in non-LGN relay nuclei) and extrinsic sources (feedback from layer 6 of the cortex as well as from brainstem reticular formation) (see Fig. 3 for an example view of cell and network architecture diversity). Recent anatomical studies have shown a great diversity of cortical input type, strength and inferred degree of convergence, even within individual thalamic nuclei[11].

In addition to the diversity of excitatory inputs, thalamic circuits receive a diversity of inhibitory inputs. The two major systems of inhibitory control are the thalamic reticular nucleus (TRN), a shell of inhibitory nucleus surrounding thalamic excitatory nuclei, and the extrathalamic inhibitory system; a group of inhibitory projects across the fore-, mid- and hindbrain (see[32] for a review on thalamic inhibition). Perhaps a major differentiating feature of these two systems (TRN and ETI) is temporal precision. One of the key characteristics of thalamus is lack of direct local loops. Only a very small group of inhibitory neurons with local connections exists in thalamus[98]. A mechanistic consequence of this architecture is the differential control of thalamic response gain and selectivity (Fig. 3 F and G), with the TRN controlling the first as observed in sensory systems[80], and ETI controlling the latter as observed in motor systems[102]. For example, basal ganglia control of thalamic responses,

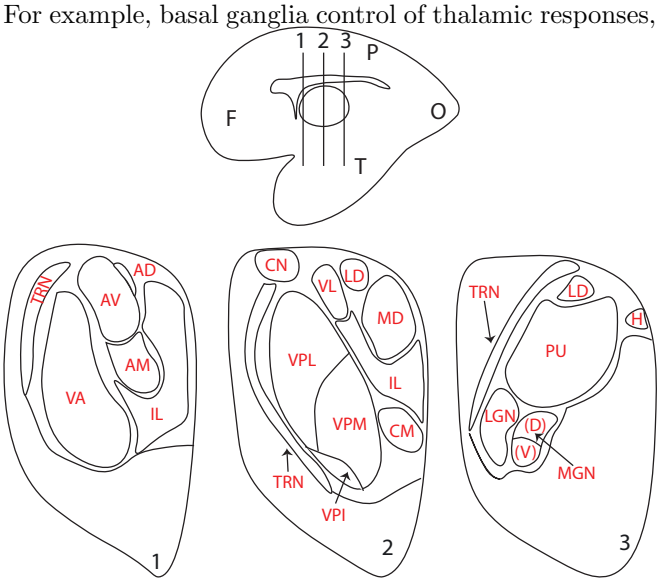

Figure 2. **Schematic layout of thalamic nuclei:** Three cross sections of monkey thalamus. *AD*: anterodorsal nucleus; *AM*: anteromedial nucleus; *AV*: anteroventral nucleus; *CM*: centromedian nucleus; *CN*: caudate nucleus; *H*: habenular nucleus; *IL*: intralaminar nuclei;*LD*: lateral dorsal nucleus; *LGN*: lateral geniculate nucleus;*MD*: mediodorsal nucleus; *MGN(D)*: medial geniculate nucleus (dorsal); *MGN(V)*: medial geniculate nucleus (ventral); *PU*: pulvinar; *TRN*: thalamic reticular nucleus (not a relay nucleus); *VA*: ventral anterior nucleus; *VL*: ventral lateral nucleus; *VPI*: ventral posterior nucleus (inferior); *VPL*: ventral posterior nucleus (lateral); *VPM*: ventral posterior nucleus medial. Redrawn from[93].

a form of ETI control, would be implemented through thalamic disinhibition, which is not only dependent on ETI input, but also a special type of thalamic conductance that enables high frequency 'bursting' upon release from inhibition[19,29].

Overall, the variety of thalamic inputs (both excitatory and inhibitory), combined with intrinsic thalamic features such as excitability and morphology will determine the type of intrinsic computations the thalamus performs. For example, if a thalamic neuron within a particular circuit motif received temporally offset cortical inputs, one direct and another through basal ganglia stations, then that neuron may be perfect for computing a prediction error based on a raw cortical signal (or state) at t_0 and a processed one at $t_0 + \Delta t$ (basal ganglia operation). This view appears to be consistent with recent observations of confidence encoding in both sensory[47] and motor systems[41,54].

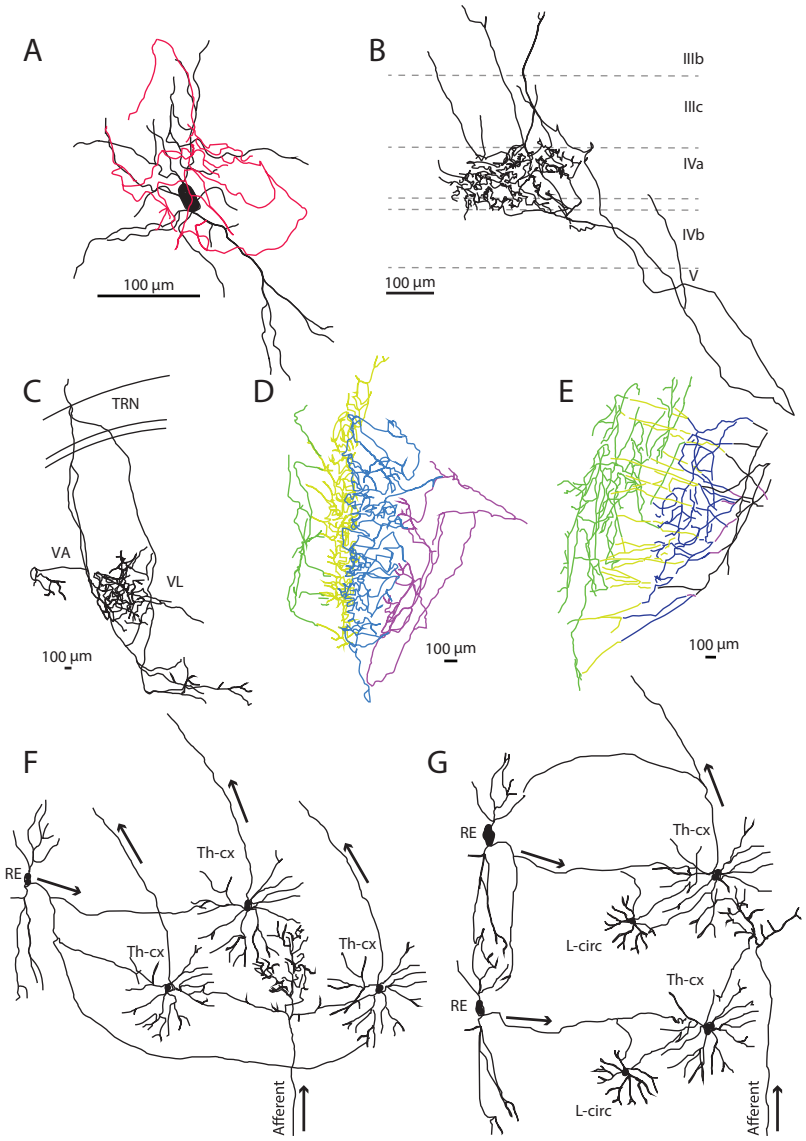

Figure 3. **Diversity and complexity of thalamo-cortical architectures:** *(A)* Comparative size of a single RGC (retinal ganglion cell) terminal (red) and an LGN neuron (black); Redrawn from[25]. *(B)* Complete terminal arbor of a single LGN neuron projecting to V1; Redrawn from[82]. *(C)* Projection of a TC (thalamo-cortical) neuron to cat motor cortex (with only 23 terminals in TRN versus 1632 terminals in VA/VL); Redrawn from[45,76] *(D,E)* Axonal arborization of a single MD neuron *(E)* and a single POm neuron *(E)*. Principal target layers are: layer I (green), layer II-IV (green-yellow), layer V (blue), layer VI (purple); Redrawn from[50]. Note the comparative size of panels *A-E* (scale bar at 100 µm). *(F,G)* Synaptic network of thalamo-cortical (Th-cx), Reticular thalamic cell (RE) and local circuit (L-circ) thalamic interneuron in rodents *(F)* and feline/primates *(G)*. Note that rodents do not have L-circ. Afferent axon excites Th-cx, which in return sends the signal to cortex. RE inhibitory effect on Th-cx cells varies depending on the excitatory drive to each Th-cx cell (*F*: compare the two neurons on the right versus the one on the left). Axonal collaterals of an RE cell could inhibit another RE cell (*G*: top RE inhibits the bottom RE), which releases the activity of L-circ leading to inhibition of weakly excited Th-cx (bottom) adjacent to the active Th-cx (top). Panels *F* and *G* are redrawn from[99] based on experiments from[97,100].

Thalamic output diversity sets thalamus to be poised as the relay as well as the modulator of cortical activity. Thalamic relay nuclei mostly project to the cortical middle layers in a topographic fashion. However, the majority of thalamic structures project more diffusely to the cortical superficial layers, such as mediodorsal (MD), posteriodmedial complex (POm) and pulvinar for example (see Fig. 3 for an example of thalamic cell and circuit diversity). These diffuse projections seem poorly suited to relay information in a piecewise manner. Rather,

they might have a modulatory role of cortical function. Further, a great degree of diversity can be observed at the level of thalamic axonal terminals within the cortex. While the idea of a thalamic relay was consolidated by observing that the main LGN neurons thought to be associated with form vision (M and P pathways) exhibit spatially compact cortical terminals, recent anatomical studies of individual neurons across the thalamus show a variety of terminal sizes and degree of spatial spread and intricate computational architecture (Fig. 3). This complexity of the architecture and diversity of the computing nodes are among the key factors that set apart the thalamo-cortical system from other conventional and unconventional computing engines (Fig. 1). Part of the complication in understanding how these anatomical types give rise to different functions is their potential for contacting different sets of excitatory and inhibitory cortical neurons. For example, activating the mediodorsal thalamus (MD) does not generate spikes across a population of prefrontal cortical neurons they project to, while activating LGN generates spikes in striate cortex[88]. Instead, MD activation results in overall enhancement of inhibitory tone, coupled with enhanced local recurrent connectivity within the PFC. This finding also argues against a relay function, because in this case the MD is changing the mode by which PFC neurons interact with one another, initiating and updating different attractor dynamics underlying distinct cognitive states. This idea of the thalamus controlling cortical state parameters is highlighted in Figs. 4, 5 and the next section.

Many facets of thalamic computation

It is commonly thought that processes like attention, decision making and working memory are implemented through distributed computations across multiple cortical regions[12,62,89]. However, it is unclear how these computations are coordinated to drive relevant behavioral outputs. From an anatomical standpoint, the thalamus is strategically positioned to perform this function, but relatively little is known about its broad functional engagement in cognition. The thalamic cellular composition and network structure constrain how cortex receives and processes information. The thalamus is the major input to the cortex and interactions between the two structures are critical for sensation, action and cognition[44,69,92]. Despite recent studies showing that the mammalian thalamus contains several circuit motifs, each with specific input/output characteristics, the thalamus is traditionally viewed as a relay to or between cortical regions[93].

It is worth mentioning that this view of bona fide thalamic computations is quite distinct from the one in which thalamic responses reflect their inputs, with only linear changes in response size. This property of reflecting an input (with only slight modification of amplitude) was initially observed in the lateral geniculate nucleus (LGN), which receives inputs from the retina. LGN responses to specific sensory inputs (their receptive fields, RF) are very similar to those in the retina itself, arguing that there is little intrinsic computation happening in the LGN itself outside of gain control. Success in early vision studies[36,37] might have inadvertently given rise to the LGN relay function being generalized across the thalamus. The strictly feedforward thalamic role in cognition requires reconsideration[33]; only a few thalamic territories receive peripheral sensory inputs and project cortically in a localized manner, as the LGN does[25,42,45,82,92]. Below we will discuss how the distinctive anatomical architecture and computational role of pulvinar and MD differ from the one just described for LGN.

The largest thalamic structures in mammals, the MD and pulvinar contain many neurons that receive convergent cortical inputs and project diffusely across multiple cortical layers and regions[11,86]. For example, the primate pulvinar has both territories that receive topographical, non-convergent inputs from the striate cortex[86] and others that receive convergent inputs from non-striate visual cortical (and frontal) areas[60]. This same thalamic nucleus also receives inputs from the superior colliculus[79], a subcortical region receiving retinal inputs. This suggests that the pulvinar contains multiple input 'motifs' solely based on the diversity of excitatory input. This input diversity is not limited to the pulvinar, but is seen within many thalamic nuclei across the mammalian forebrain[8]. Local inactivation of pulvinar neurons results in reduced neural activity in primary visual cortex[81] suggesting a feedforward role. In contrast, recent findings show that during perceptual decision making pulvinar neurons encode choice confidence, rather than stimulus category[47], strongly arguing for more pulvinar functions beyond relaying information.

In the case of MD, direct sensory input is limited[64] and the diffuse, distributed projections to cortex[50] are poorly suited for information relay; this input/output connectivity suggests different functions. Recent studies[9,88] have begun to shed light on the type of computation that MD performs. Studying attentional control in the mouse[88] has revealed that MD coordinates task-relevant activity in the prefrontal cortex (PFC) in a manner analogous to a contextual signal regulating distinct attractor states within a computing reservoir. Specifically, in a task where animals had to keep a rule in mind over a brief delay period (Fig. 4A), PFC neurons show population-level persistent activity following the rule presentation, a sensory cue that instructs the animal to direct its attention to either vision or audition (Fig. 4B,C). MD neurons show responses that are devoid of categorical selectivity (Fig. 4D), yet are critical for selective PFC activity; optogenetic MD inhibition diminishes this activity, while MD activation augments it. The conclusion is that MD inputs enhance synaptic connectivity among PFC neurons or may adjust the activity of PFC neurons through selective filtering of the thalamic inputs. In other words, delay-period MD activity maintains rule-selective PFC representations by augmenting local excitatory recurrence[88].

In a related study, a delayed nonmatching-to-sample T-maze working memory task[9], it was shown that MD amplification and maintenance of higher PFC activity indicated correct performance during the subsequent choice phase of the task. Interestingly, MD-dependent increased PFC activity was much more pronounced during the later (in delay) rather than earlier part of the task. These findings indicate that PFC might have to recursively pull in MD to sustain cortical representations as working memory weakens with time. Together these studies indicate that PFC cognitive computation can not be dissociated from MD activity. Further evidence for the critical role of the MD-PFC interaction for cognition is the disrupted fronto-thalamic anatomical and functional connectivity seen in neurodevelopmental disorders[59,65,68,78,108].

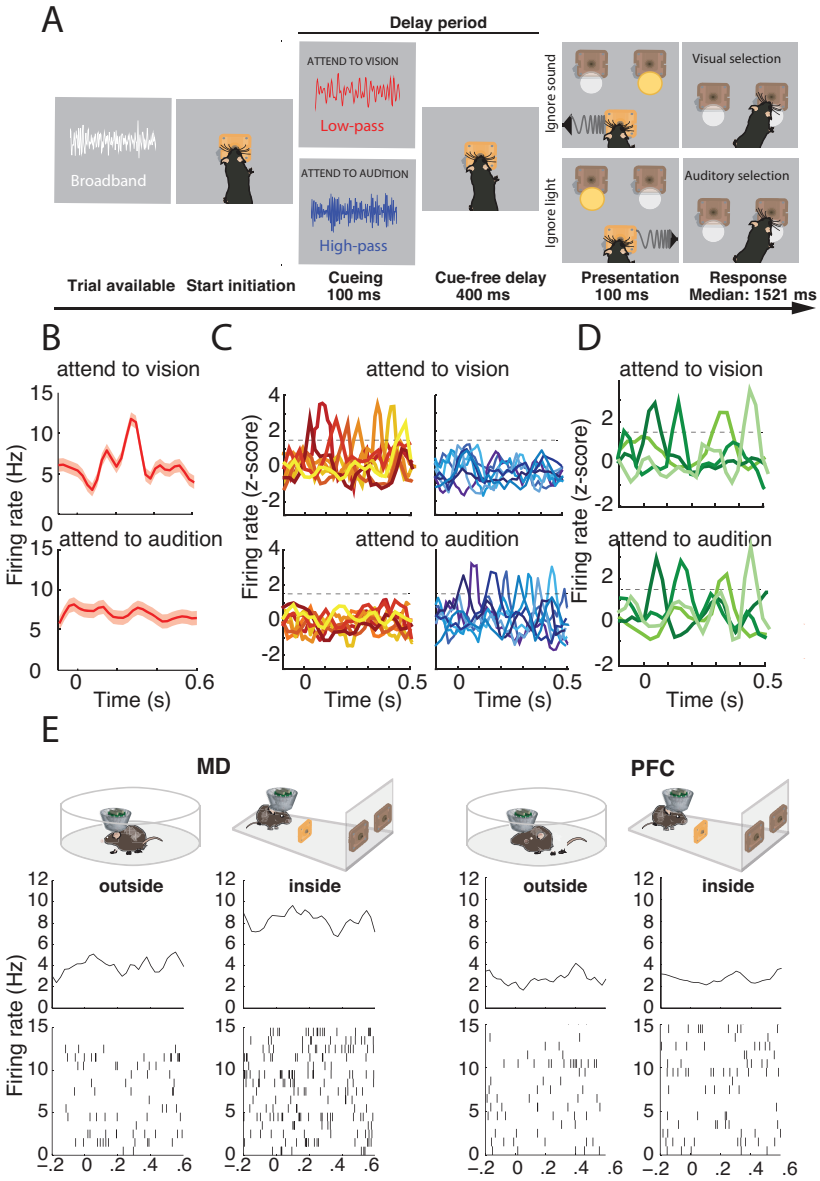

Figure 4. **MD-PFC interactions during sustained rule representations.** *(A)* attentional control task design. *(B)* Example peri-stimulus time histogram (PSTH) for a neuron tuned to attend to vision rule signaled through low pass noise cue *(C)* Examples showing that rule-specificity is maintained across distinct PFC rule-tuned populations. *(D)* PSTHs of four MD neurons showing consistent lack of rule specificity. *(E)* Example rasters and PSTHs of an MD and PFC neuron when the animal is engaged in the task and outside of the behavioral arena. In contrast to PFC, MD neurons show the contextual difference in a change in firing rate. Figure is redrawn from[88].

Can MD select cortical subnetworks based on contextual modulation?

Why would a recurrent network (PFC) computation depend on its interaction with a non-recurrent (MD) non-relay network? What computational advantage such system would have? Using a chemogenetic approach, a recent study suggested that information flow in the MD-PFC network can be unidirectional. While both inactivating PFC-to-thalamus and MD-to-cortex pathways impaired recognition of a change in reward value in rats performing a decision making task, only the inactivation of MD-to-cortex pathway had an impact on the behavioral response to a change in action-reward relationship[3]. Given that a sensory stimulus may require a different action depending on the context in which it occurs, the ability to flexibly re-route the active PFC subnetwork to a different output may be crucial. In an architecture like the PFC-MD network, where MD can modulate PFC functional connectivity, MD might well be suited to re-route the ongoing activity in a context dependent manner. In fact, in the mouse cognitive task described above (Fig. 4A), a subset of MD neurons showed substantial spike rate modulation during task engagement compared to when the animals is in its home in cage (see Fig. 4E)[88]. In contrast, PFC neurons show very little difference in spike rates when the animal gets engaged in the task. This suggests that perhaps different subsets of MD neurons are capable of encoding task 'contexts'. Subsequently, each given subset could unlock a distinct cortical association. These MD subsets have to be able to shift the cortical states dynamically while maintaining the selectivity based on the subset of cortical connections they target. This idea would also fit with the paradigm shift indicating that thalamic neurons exert dynamical control over information relay to cortex[6,77].

Overall, the anatomical and neurophysiological data show that the thalamic structure and cortico-thalamic network circuitry, and the interplay between thalamus and cortex, shape the frame within which thalamus plays the dual role of relay and modulator. Under this framework, different thalamic nuclei carry out multitude of computation including but not limited to information relay. A suggestion of this comparative computational role of LGN, pulvinar and MD is depicted in Fig 8.

Is thalamus a read-write medium for cortical parallel processing?

If the brain were to function as a simple pattern matching system without wiring and metabolic constrains, evolution would just expand the size and depth of the network to the point that it could potentially memorize a large number of possible patterns. Possibly, evolution would have achieved this approximation of arbitrary patterns by evolving a deep network. This would be a desirable solution since any system can be defined as a polynomial Hamiltonians of low order, which can be accurately approximated by neural networks[53]. But cognition is much more than template matching and classification achieved by a neural network. The limits of template matching methods in dealing with (rotation, translation and scale) invariance in object recognition quickly became known to neuroscientists and in early works on computer vision. One of the early pioneers of AI, Oliver Selfridge, proposed Pandemonium architecture to overcome this issue[90]. Selfridge envisioned serially connected distinct demons (an image demon, followed by a set of parallel feature demons, followed by a set of parallel cognitive demons and eventually a decision demon), that independently perceive parts of the input before reaching a consensus together through a mixture of serial culmination of evidence from parallel processing. This simple feedforward computational pattern recognition model is (in some ways) a predecessor to modern day connectionist feedforward neural networks, much like what we discussed earlier in the text. However, despite its simplicity, Pandemonium was a leap forward in understanding that the intensity of (independent parallel) activity along with a need to a summation inference are the keys to move from simple template matching to a system that has a concept about the processed input. A later extension of this idea was proposed by Allen Newell as the Blackboard model: *"Metaphorically, we can think of a set of workers, all looking at the same blackboard: each is able to read everything that is on it and to judge when he has something worthwhile to add to it. This conception is just that of Selfridge's Pandemonium: a set of demons independently looking at the total situation and shrieking in proportion to what they see that fits their natures"*[71]. Blackboard AI systems, adapted based on this model, have a common knowledge base (blackboard) that is iteratively updated (written to and read from) by a group of knowledge sources (specialist modules), and a control shell (organizing the updates by knowledge sources)[72].

Interestingly, this computational metaphor can also be extended to the interaction between thalamus and cortex, though thalamic blackboard is not a passive one as in the blackboard systems[34,66,67]. Although, initially the active blackboard was used as an analogy for LGN computation, the nature of MD connectivity and its communication with cortex seem much more suitable to the type of computations that is enabled by an active blackboard. Starting with an input, thalamus as the common blackboard visible to processing (cortical) modules, initially presents the problem (input) for parallel processing by modules. Here by module, we refer to a group of cortical neurons that form a functional assembly which may or may not be clustered together (in a column for example). By iteratively reading (via thalamo-cortical projections) from and writing (via cortico-thalamic projections) to this active blackboard, expert pattern recognition modules, gradually refine their initial guess based on their internal processing and the updates of the common knowledge. This process continues until the problem

is solved (Fig. 5). However, since we are dealing with a biological system with finite resources, this back and forth communication needs to have some characteristics to provide a viable computational solution. First and foremost, the control of interaction and its scheduling has to have a plausible biological component and should bind solutions as time evolves. Second, to avoid turning into an NP-hard (non-deterministic polynomial-time hardness) problem, there must exist a mechanism that stops this iterative computation once an approximation has been reached (Fig. 5). In this paper, we propose specific solution to the first problem and a plausible one for the later issue. We suggest that phase-dependent contextual modulation serves to deal with the first issue and a multi-objective optimization of efficiency (computational information gain) and economy (computational cost, i.e. metabolic needs and the required time for computation) handles the second issue (Fig. 6). In both cases, we suggest that thalamus plays an integral role in conjunction with cortex.

Core attributes of cognitive processing, attention and binding in time

To expand the idea of read-write further, let's revisit some core attributes of cognitive processing. First, we wish to point to two key features of cognitive processing, namely "searchlight of attention" and "binding in time". Then we examine how these attributes match the computational constrains that are met by thalamic circuitry and functions. The "searchlight of attention", first proposed by Crick as a function of thalamic reticular formation[14], proposes that through change in firing pattern (bursting vs tonic firing), thalamic nuclei dynamically switch between detection and perception. The bursting nonlinear response signals a change in the environment, while the tonic mode underlies perceptual processing. One biophysical mechanism responsible for this change between bursting and tonic modes is implemented by cortico-thalamic activation of glutamate metabotropic receptors[61]. This mechanism puts thalamus as the mediator between peripheral and cortical input, creating a closed-loop computation[1]. A necessary property of thalamic-driven change of perceptual processing, and hence cognition, is the existence of cortical re-entry. From both anatomical studies and electrophysiological investigations[31], we know that thalamus is at a prime position to modify the relay signal based on the cognitive processing that is happening in the cortex[9,88]. This thalamic-driven attention entails "binding in time" since how thalamus modifies its relay at a given time is itself influenced by what is perceived by the cortex in time prior.

But how can the "binding in time" avoid locking-in the thalamic function to a set of inputs at a given time? How can thalamus constantly be both ahead of cortex and yet keep track of the past information? The secret may be embedded in the non-recurrent intrinsic structure of thalamus and the recurrent structure of the higher cortical areas. As mentioned earlier, we know that hierarchical convolutional neural networks (HCNN), which can recapitulate certain properties of static hierarchical forward models, can not capture any processes that need to store prior states[109]. As a result, context-dependent processing can be extremely hard to implement in neural network models[84]. The most widely used ANNs (Feedforward nets , i.e. multilayer perceptrons/Deep Learning algorithms) face fundamental deficiencies: the ubiquitous training algorithms (such as back-propagation), i) have no biological plausibility, ii) have high computational cost (number of operations and speed), and iii) require millions of examples for proper adjustment of NN weights. These features renders feedforward NNs not suitable for temporal information processing. In contrast, recurrent neural networks (RNNs) can universally approximate the state of dynamical systems[26], and because of their dynamical memory are well suited for contextual computation. It has been suggested that RNNs are able to perform real-time computation at the edge of chaos, where the network harbors optimal readout and memory capacity[7]. The functional similarity of gated-recurrent neural networks and cortical microcircuitry continues to be at the forefront of AI and neuroscience interface[13]. If higher cortical areas were to show some features of RNN-like networks, as manifested by the dynamical response of single neurons[57], then we anticipate that the local computation (interaction between neighboring neurons) to be mostly driven by external biases. The thalamic projections could then play the role of bias where they seed the state of the network. While this may be a feasible portrait, the picture misses the existence of large-scale feedback loops which are neither feedforward nor locally recurrent[2,31]. To reconcile, we need our system of interest to show phase-sensitive detection that binds the locally recurrent activity in cortex, with large-scale feedback loops.

Computational constrains and the role of thalamus in phase-dependent contextual modulation

Based on the observations of behavior, higher cognition requires "efficient computation", "time delay feedback", the capacity to "retain information" and "contextual" computational properties. Such computational cognitive process surpass the computational capacity of simple RNN-like networks. The essential required properties of a complex cognitive system of such kind are: 1) input should be nonlinearly mapped onto the high-dimensional

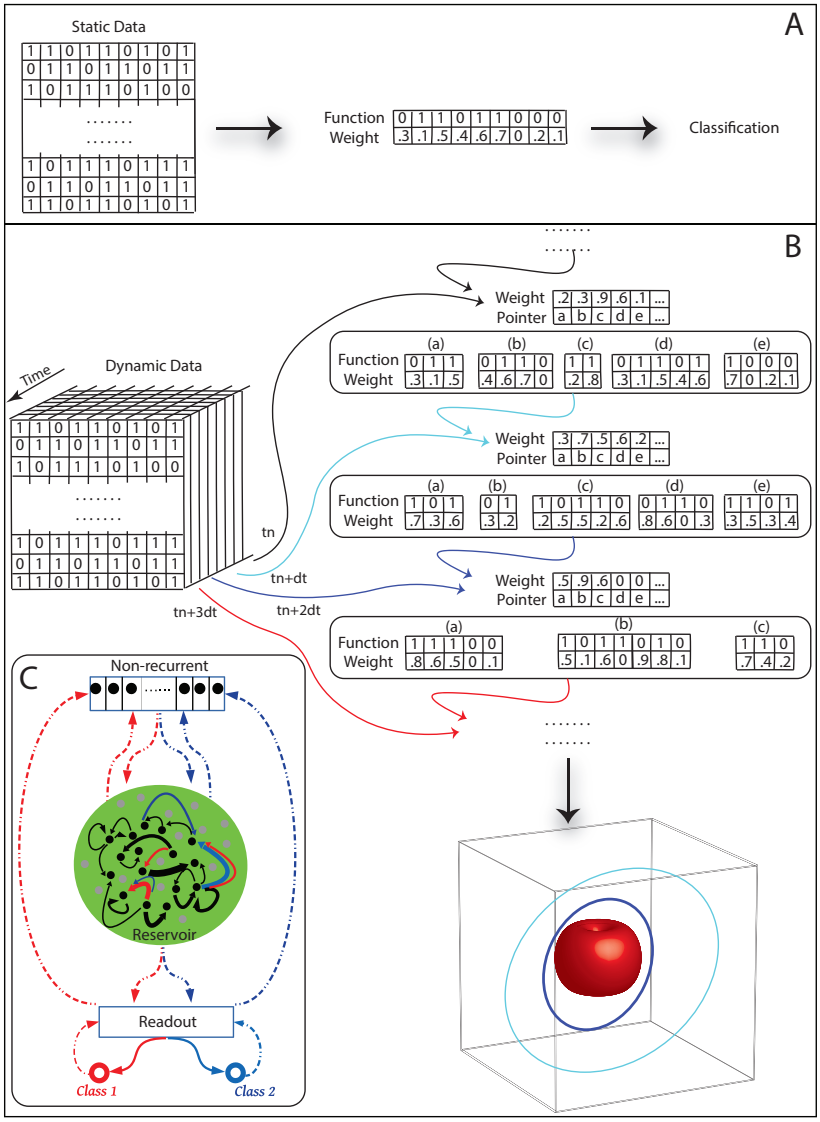

Figure 5. **Schematic representation of thalamic cognitive contextual computation.** *(A)*: In the case of static data, a set of function/weight modules can yield good classification. Function represents a polynomial (since any system that is known to be a polynomial Hamiltonians of low order can be accurately approximated by neural networks[53]) and the weights exemplify the connection matrix of an artificial neural network encapsulating this polynomial. Stacking multiple of such module can increase the accuracy of polynomial approximation (such as in the case of CNN). *(B)*: Thalamo-cortical computation for contextual processing of dynamic data. Each dataframe is processed by a weight/pointer module (thalamus MD-like structure) which like a blackboard is writable by different sets of neuronal assemblies in cortex. Thalamic pointers assign the assemblies; modules' weights adjust the influence of each assembly in further computational step [*inset C* shows a non-recurrent thalamic nuclei (MD-like) modulating the weights in the PFC (reservoir and readout). Here, depending on the context (blue or red), the interactions between MD and Reservoir, between Reservoir and Readout, and between Readout and MD could pursue one of the two possible outcomes. Specifically, MD changes the weights in the Reservoir to differentially set assemblies that produce two different attractor states, each leading to one of the two possible network outputs]. In *(B, C)*, each operation of the thalamic module is itself influenced, not only by the current frame (t), but also by the computation carried by cortex module on the prior frame ($t-1$). Cortical module is composed of multiple assemblies where each operate similar to the function/weight module of the static case. These assemblies are locally recurrent and each cell may be recruited to a different assembly during each operation. This mechanism could explain why prefrontal cells show mixed selectivity in their responses to stimuli (as reported in[27,83]). Through this recursive interaction between thalamus and cortex, cognition emerges not as just a pattern matching computation, but through contextual computation of dynamic data (bottom right schematic drawing).

state, while different inputs map onto different states, 2) slightly different states should map onto identical targets, 3) only recent past should influence the state and network is essentially unaware of remote past, 4) a phase-locked

loop should decode information that is already encoded in time and 5) the combination of 1-4, should optimize sensory processing based on the context. The first three attributes of such system have close relevance to constrains and computational properties of higher cortical areas (prefrontal). The same three are also the main features of reservoir computing, namely "separation property", "approximation property" and "fading memory". Reservoir Computing (RC), with Echo State Networks (ESN)[39,40] and Liquid State Machines (LSM)[55,56] as two early variants of what is now unified as RC, have emerged as powerful RNNs. Although a large enough random network could essentially represent all combinations of cortical network[84], the training of such system would be hard and changing the network from task to task will not be easily achievable. An advantage of a RC systems is to "non-linearly" map a lower dimensional system to a high-dimensional space facilitating classification of the elements of the low-dimensional space. The last two properties match the structure and computational constraints of non-relay thalamic system as a contextual modulator that is phasically changing the input to the RC system. In fact, in an RC model of prefrontal cortex, addition of a phase neuron significantly improved the networks performance in complex cognitive tasks. The phase neuron improves the performance by generating input driven attractor dynamics that best matched the input[23]. In a recent study, electronic implementation and numerical studies of a limited RC system of a single nonlinear node with delayed feedback showed efficient information processing[4]. Such reservoir's transient dynamical response follows delay-dynamical systems, and only a limited set of parameters set the rich dynamical properties of delay systems[38]. This system was able to effectively process time-dependent signals. The phase neuron[23] and delayed dynamical RC[4] both show properties that resemble the thalamic functions as discussed here. The collective system (thalamus and cortex together) is neither feedforward, nor locally recurrent, but it has a mixture of non-recurrent phase encoder that keeps copies of the past processing and modulates the sensory input relay and its next step processing (Fig. 5). These features emphasize that the perceptual and cognitive processing can not be solely cortico-centric operations.

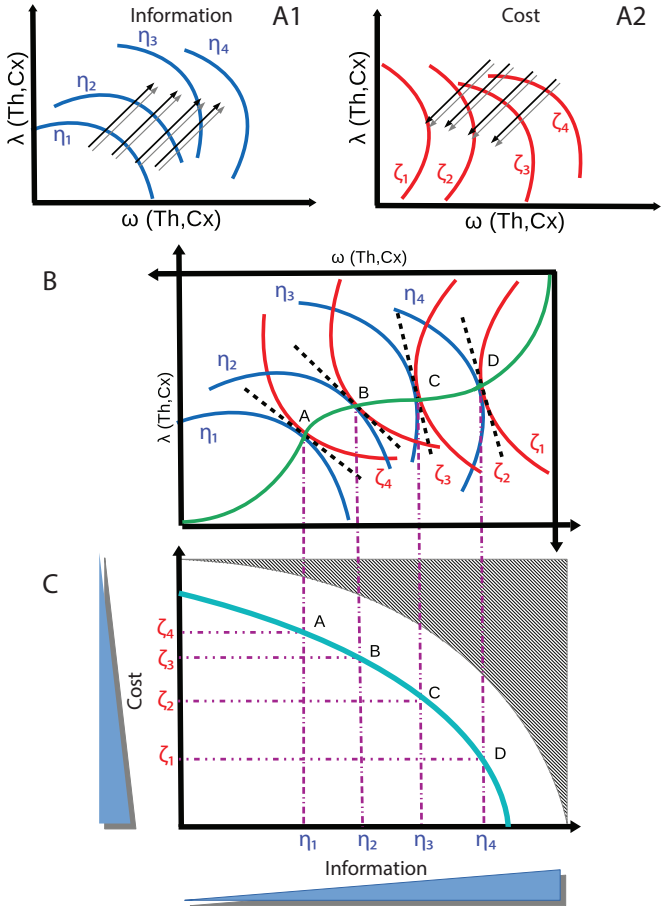

Figure 6. **Dynamic role of thalamo-cortical system in the information/cost optimization.** *(A)* Iso-maps of information *(A1)* and cost *(A2)* in the domain specified by ω and λ (functions of cortical and thalamic activity). Information across each Iso-quant curve (η_1 for example) is constant and is achieved at a certain mixture of ω and λ. Optimal information can be obtained by moving outward (arrows, *A1*). Cost optimization can be achieved by moving inward (*A2*). *(B)* since information and cost are both defined in the domain of ω and λ, thalamus and cortex jointly contribute to information and cost optimization. The points where the iso-quant curves' tangents are equal (black dashed line), provide the optimal combination of information/cost (green curve). In any cycle of cognitive operation, depending on the prior state of the system (ω and λ), the nearest points on the green curve are the optimal solutions for ending that cycle. *(C)* mapping of the optima curve to information/cost domain shows all pareto efficient allocations (cyan curve). The slope of the parto frontier shows how the system trades cost versus information: along the pareto curve, efficiency is constant but the exchange between information and cost is not. All allocations inside of this curve could be improved as thalamus and cortex interact. The grey zone shows the biophysically not-permissible allocation of computation and resource.

Biological constrains and the role of thalamus in computational optimization

Computation and optimization are two sides of the same coin. But how does the brain optimize the computations that would match its required objective, i.e. cognitive processing? There is a current trend of thinking that brain optimizes some arbitrary functions, with the hopes that the future discovery of these unknown functions may guide us to establish a link between brain's operations and deep learning[58]. This line of approach to optimizational (and computational) operations of the brain has few flaws. First, it avoids specifying what function the brain is supposed to optimize (and as a result it remains vague). Second, it refrains from addressing certain

limitations that brain has to cope with due to biological constrains. First of these limitations is the importance of using just enough resources to solve the current perceptual problem. Second is the necessity to come up with a solution just in (the needed) time. The importance of "just-enough" and "just-in-time" computation in cortical computation should not be overlooked[20]. If the first condition is not met, the organism can not sustain continued activity since the metabolic demand surpasses the dedicated energetic expenditure and the animal can not survive. In fact, the communication in neural networks are highly constrained by number of factors, specifically the energetic demands of the network operations[51]. From estimates of the cost of cortical computation[52], we know that the high cost of spiking forces the brain to rely on sparse communication and using only a small fraction of the available neurons[5,94]. While, theoretically, cortex can dedicate a large number of neurons (and very high dynamical space) to solve any cognitive task, metabolic demand of such high-energetic neural activity renders such mechanism highly inefficient. As a result, the "law of diminishing returns" dictates that increased energetic cost causing excessive pooling of active neurons to an assembly would be penalized[73]. The penalization for unnecessary high-energetic neural activity, in itself, should be driven by the nature of computation rather than being formulated as a fixed arbitrary threshold imposed by an external observer. On the other hand, a system can resort to low-cost computation at any given time but dedicate long enough time to solve the task on hand. Naturally, such system would not be very relevant to the biological systems since time is of essence. If an animal dedicates a long instance of its computational capacity to solve a problem, the environment has changed before it reaches a solution and the solution becomes obsolete. A deer would never have an advantage for its brain to have fully analyzed the visual scene instead of spotting the approaching wolf and shifting resources to the most-needed task, i.e. escape. As a result, many of the optimization techniques and concepts that may be relevant to artificial neural networks are irrelevant to embodied computational cognition of the brain. The optimization that the brain requires is not aiming for the best possible performance, but rather needs to reach a good mixture of economy and efficiency.

Not surprisingly, these constrains, i.e. efficiency and economy, are cornerstones of homeostasis and are observed across many scales in living systems[101]. The simple "Integral feedback" acts as the mainstay of control feedback in such homeostatic systems (such as E Coli heat-shock or DNA repair after exposure to gamma radiation)[18,21,22,48]. Change in input leads to change in the output and the proportional change in the controller aiming to reset the output to the desired regime. The constrains that we discussed above, remind us of a similar scenario. Instead of just trying to deal with one fitness function at a time (where the minima of the landscape would be deemed as "the" optima), the brain has to perform a multi-objective optimization, finding solutions to both metabolic cost (economy) and just-in-time (efficiency) computation. Thus we can infer that a unique solution does not exist for such a problem. Rather, any optimization for computational efficiency will cost us economy and any optimization for economy will cost us efficiency. In such case, a multi-objective optimization pareto frontier is desirable. Pareto frontier of information/cost will be the set of solutions where any other point in the space is objectively worse for both of the objectives[28,49,101]. As a result, the optimization mechanism should push the system to this frontier. The iterative dynamical interaction between thalamus and cortex seems to provide an elegant solution for this problem. We discuss this in more details below.

Consider a set of functions, ω and λ of f_{Th} (firing rate of thalamic cell) and f_{Cx} (firing rate of cortical cells). Uncertainty (or its opposite, information) and computational cost (a mixture of time and metabolic expense) can both be mapped to this functional space of $\omega\left(f_{Th}, f_{Cx}\right)$ and $\lambda\left(f_{Th}, f_{Cx}\right)$ (Fig. 6A1,2). Let's define computational cost and information as product and linear sum of cortical and thalamic activity ($\alpha f_{Th}^{n}.\beta f_{Cx}^{n}$, $\theta\frac{d_{fTh}^{n}}{d_t}+\psi\frac{d_{fCx}^{n}}{d_t}$; with α, β, θ, ψ as coefficients) to reflect the logarithmic nature of information (entropy) and the fact that biological cost is an accelerating function of the cost-inducing variables[18]. The hypothetical space of cost/information is depicted in Fig. 6, where top panels show indifference maps of information (A1) and cost (A2). The example simulations and parametric plots of the cost and information functions defined as above are shown in Fig. 7. In each indifference map, along each iso-quant curve, the total functional attribute is the same. For example, anywhere on the η_1 curve, the uncertainty (or information) in our computational engine is the same. However, different iso-quant curves represent different levels of the functional attribute. For example, moving outward increases information (reduces uncertainty) as $\eta_1 < \eta_2 < \eta_3 < \eta_4$ and thus if computational cost was not a constrain, the optimal solution would have existed on η_4 or further away (Fig.6 A1). In contrast, moving inward would preserve the cost ($\zeta_1 < \zeta_2 < \zeta_3 < \zeta_4$) if the computational engine did not have the objective of reducing uncertainty (Fig. 6A2). Since information and cost are interdependent and both depend on the interaction between thalamus and cortex, we suggest that information/cost optimization happens through an iterative interaction between thalamus and cortex (note the blackboard analogy and contextual modulation discussed above). Since we defined both information and cost as a set of iso-quant curves in the functional space of $\omega\left(f_{Th}, f_{Cx}\right)$ and $\lambda\left(f_{Th}, f_{Cx}\right)$, they can be co-represented in the same space (Fig. 6B). Optimal solutions for information/cost optimizations are simply the solutions to where the tangents of the iso-quant curves are equal (see the tangents [black dashed lines] and points A, B, C and D, in Fig. 6B). These points create a set of optimal solutions for the tradeoff between information and cost (green

curve). Mapping of the optimal solutions to the computational efficiency space E, gives us the pareto efficient curve (cyan curve, Fig. 6C). Anywhere inside the curve is not pareto efficient (i.e. information gain and computational cost can change in such a way that, collectively, the system can be in a better state (on the pareto curve). Points outside of the pareto efficient curve are not available to the current state of the system due to the coefficients of ω and λ. A change in these coefficients can potentially shape a different co-representation of information and cost (see Fig 7, top row for 3 different instances of ω and λ based on different α, β, θ, ψ values), and thus a different pareto efficient curve (see Fig 7, bottom row). These different possible pareto frontiers can be set based on the prior state of the system and the complexity of the computational problem on hand. Nonetheless, the computational efficiency of the system can not be infinitely pushed outward because of the system's intrinsic biophysical constrains (neurons and their wiring). The shaded region in Fig 7, bottom row, shows this non-permissible zone.

In the defined computational efficiency space E, composed of the two variables information and cost (as the objective functions, shown in bottom panels of Fig. 6 and Fig. 7), solving a computational problem is represented by a decrease in uncertainty. However, any change in uncertainty has an associated cost. First derivative of the pareto frontier shows "marginal rate of substitution" as $\frac{\Delta_{info}}{\Delta_{cost}}$. This ratio varies among different points on the pareto efficient curve. If we take two points on the pareto curve in the computational efficiency space, such as A and C for example, computational efficiency of these two points are equal $E_{A(\eta_1,\lambda_4)} = E_{C(\eta_3,\lambda_2)}$. The change in efficiency of point A with respect to information and cost, are the partial derivatives $\frac{\partial E_A}{\partial info}$ and $\frac{\partial E_A}{\partial \cos t}$, respectively. As a result, $\frac{\partial E_A}{\partial \inf o} d_{\inf o} + \frac{\partial E_A}{\partial \cos t} d_{\cos t} = 0$, meaning that there is constant efficiency along the pareto curve, the tradeoff between information and cost is not constant. The optimization in this space is not based on some fixed built-in algorithm or arbitrary thresholds by an external observer. Rather, information/cost optimization is the result of back and forth interaction between thalamus and cortex. Based on the computational perspective that we have portrayed, thalamus seems to be poised to operate as an optimizer. Thalamus receives a copy of (sensory) input while relaying it, and receives an efferent copy from the processor (cortex), while trying to efficiently bind the information from past and present and sending it back to cortex. The outcome of such emergent optimization, is a pareto front in the economy-efficiency landscape (Fig. 6,7). If the cortex were to be the sole conductor of cognitive processing, the dynamics of the relay and cortical processing would meander in the parameter space and not yielding any optimization that can provide a feasible solution to economic and just-in-time computation. Such system is doomed to fail, either due to metabolic costs or due to computational freeze over time ; thus more or less be a useless cognitive engine. In contrast, with the help of an optimizer that acts as a contextual modulator, the acceptable parameters will be confined to a manifold within the parameter space. Such regime would be a sustainable and favorable domain for cognitive computing. This property shows another important facet of a thalamo-cortical computational cognitive system and the need to move passed the cortico-centric view of cognition.

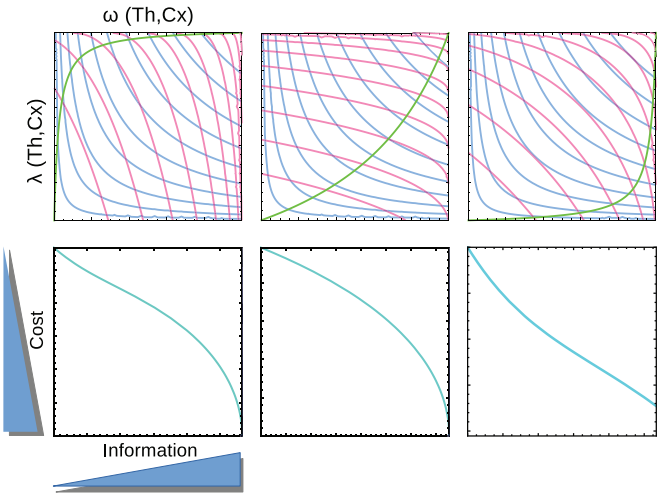

Figure 7. **Dynamic parameter space of thalamo-cortical joint optimization of information/cost.** Three different realization of information/cost interaction as a function of thalamic and cortical activity (ω, λ) and the corresponding pareto curves (see Fig. 6 for details of this optimization construct). Pareto curve shows the optimal set of both cost and information that can be obtained given the biophysical constrains of neurons and networks connecting them. Every point on the pareto frontier shows technically efficient levels for a given parameter set of ω, λ (see text for more details). All the points inside the curve are feasible but are not maximally efficient. The slope (marginal rate of transformation between cost and information) shows how in order to increase information, cost has to change. The dynamic nature of interaction between thalamus and cortex enables an emergent optimization of information/cost depending on the computational problem on hand and the prior state of the system.

Concluding remarks: Reframing Thalamic function above and beyond information relay

Lately, new evidence about the possible role of thalamus has started to challenged the cortico-centric view of perception/cognition. Anatomical studies and physiological measurements have begun to unravel the importance of the Cortico-Thalamo-Cortical loops in cognitive processes[6,77]. Under this emerging paradigm, thalamus plays two distinctive roles: a) information relay, b) modulation of cortical function[93], where the neocortex does not work in isolation but is largely dependent on tha-

lamus. In contrast to cortical networks which operate as specialized memory devices via their local recurrent excitatory connections, the thalamus is devoid of local connections, and is instead optimized for capturing state information that is distributed across multiple cortical nodes while animals are engaged in context-dependent task switching[88]. This allows the thalamus to explicitly represent task context (corresponding to different combinations of cortical states), and through its unique projection patterns to the cortex, different thalamic inputs modify the effective connections between cortical neurons[9,88].

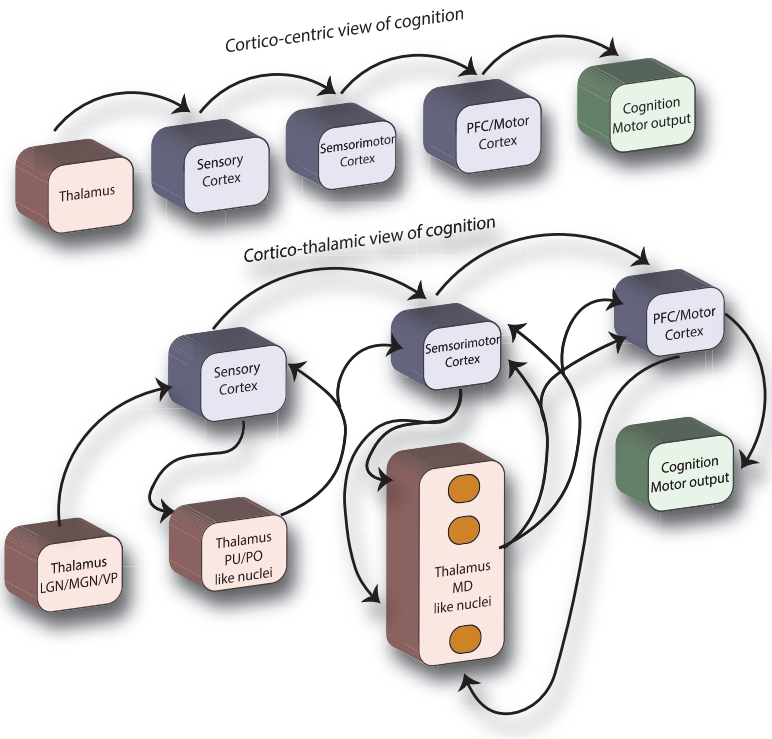

Figure 8. **The emergent view of thalamic role in cognition.** *(Top)* In the traditional view, serial processing of information confines the role of thalamus to only a relay station. *(Bottom)* the view that is discussed in this manuscript considers thalamus as a key player in cognition, above and beyond relay to sensory cortices. Through combining the efferent readout from cortex with sensory afferent, MD-like thalamic nuclei modulate further activity of the higher cortex. The contextual modulation enabled by MD is composed of distinctively parallel operations (individual circles represent the non-recurrent nature of these processes due to lack of local excitatory connections). Under this view, and the computational operatives discussed here, the thalamo-cortical system (and not just cortex) is in charge of contextual cognitive computing. The computation enabled by Pulvinar/PO like nuclei is different from LGN and also from MD-like nuclei.

Here, we started with a brief overview of the architecture of thalamus, the back and forth communication between thalamus and cortex, then we provided the electrophysiological evidence of thalamic modulatory function, and concluded with a computational frame that encapsulates the architectural and functional attributes of the thalamic role in cognition. In such frame, the computational efficiency of the cognitive computing machinery is achieved through iterative interactions between thalamus and cortex embedded in the hierarchical organization (Fig. 4, 5). Under this emergent view, thalamus serves not only as relay, but also as a read/write medium for cortical processing , playing a crucial role in contextual modulation of cognition (Fig. 8). Such multiscale organization of computational processes is a necessary requirement for design of the intelligent systems[16,95,96]. Distributed computing in biological systems in most cases operates without central control[70]. This is well reflected in the computational perspective that we discussed here. We suggest that through the continuous contextual modulation of cortical activity, thalamus (along with cortex) plays a significant role in emergent optimization of computational efficiency and computational cost. This phenomenon has a deep relation with phase transitions in complex networks. Different states (phases) of the network are associated with the connectivity of the com-

puting elements (see thalamic weight/pointer and cortical function/weight modules in Fig. 5). Interestingly, intrinsic properties of the complex networks do not define the phase transitions in system. Rather, the interplay of the system with its external environment shapes the landscape where phase transitions occur[91]. This parallel in well-studied physical systems and neuronal networks of thalamo-cortical system show the importance of the interplay between thalamus and cortex in cognitive computation and optimization. The proposed frame for contextual cognitive computation and the emergent information/cost optimization in thalamo-cortical system can guide us in designing novel AI architecture.

ACKNOWLEDGMENTS

We wish to thank Michael Halassa for helpful discussions.

REFERENCES

[1] Ahissar, E., and Kleinfeld, D. Closed-loop Neuronal Computations: Focus on Vibrissa Somatosensation in Rat. *Cerebral Cortex 13*, 1 (jan 2003), 53–62.

[2] Ahissar, E., and Oram, T. Thalamic Relay or Cortico-Thalamic Processing? Old Question New Answers. *Cerebral Cortex 25*, 4 (nov 2015), 845–848.

[3] Alcaraz, F., Fresno, V., Marchand, A. R., Kremer, E. J., Coutureau, E., and Wolff, M. Thalamocortical and corticothalamic pathways differentially contribute to goal-directed behaviors in the rat. *eLife 7* (feb 2018).

[4] Appeltant, L., Soriano, M. C., Van der Sande, G., Danckaert, J., Massar, S., Dambre, j., Schrauwen, B., Mirasso, C. R., and Fischer, I. Information processing using a single dynamical node as complex system. *Nature Communications 2* (sep 2011), 468.

[5] Baddeley, R., Abbott, L. F., Booth, M. C. A., Sengpiel, F., Freeman, T., Wakeman, E. A., and Rolls, E. T. Responses of neurons in primary and inferior temporal visual cortices to natural scenes. *Proceedings of the Royal Society B: Biological Sciences 264*, 1389 (dec 1997), 1775–1783.

[6] Basso, M. A., Uhlrich, D., and Bickford, M. E. Cortical Function: A View from the Thalamus. *Neuron 45*, 4 (feb 2005), 485–488.

[7] Bertschinger, N., and Natschläger, T. Real-time computation at the edge of chaos in recurrent neural networks. *Neural Comput 16* (Jul 2004), 1413–36.

[8] Bickford, M. E. Thalamic Circuit Diversity: Modulation of the Driver/Modulator Framework. *Frontiers in Neural Circuits 9* (jan 2016).

[9] Bolkan, S. S., Stujenske, J. M., Parnaudeau, S., Spellman, T. J., Rauffenbart, C., Abbas, A. I., Harris, A. Z., Gordon, J. A., and Kellendonk, C. Thalamic projections sustain prefrontal activity during working memory maintenance. *Nature Neuroscience 20*, 7 (may 2017), 987–996.

[10] Butler, A. B., and Hodos, W. *Comparative Vertebrate Neuroanatomy.* John Wiley & Sons Inc., aug 2005.

[11] Clasca, F., Rubio-Garrido, P., and Jabaudon, D. Unveiling the diversity of thalamocortical neuron subtypes. *Eur J Neurosci 35* (May 2012), 1524–32.

[12] Corbetta, M. Frontoparietal cortical networks for directing attention and the eye to visual locations: identical, independent, or overlapping neural systems? *Proceedings of National Academy of Science 95* (Feb 1998), 831–8.

[13] Costa, R. P., Assael, Y. M., Shillingford, B., de Freitas, N., and Vogels, T. P. Cortical microcircuits as gated-recurrent neural networks. *arXiv* (2017).

[14] Crick, F. Function of the thalamic reticular complex: the searchlight hypothesis. *Proceedings of the National Academy of Sciences 81*, 14 (jul 1984), 4586–4590.

[15] Deacon, T. W. Rethinking Mammalian Brain Evolution. *American Zoologist 30*, 3 (aug 1990), 629–705.

[16] Dehghani, N. Design of the Artificial: lessons from the biological roots of general intelligence. *arXiv* (2017).

[17] Dehghani, N. Theoretical principles of multiscale spatiotemporal control of neuronal networks: a complex systems perspective. *BiorXiv* (jan 2017).

[18] Dekel, E., and Alon, U. Optimality and evolutionary tuning of the expression level of a protein. *Nature 436*, 7050 (jul 2005), 588–592.

[19] Deniau, J. M., and Chevalier, G. Disinhibition as a basic process in the expression of striatal functions. II. The striato-nigral influence on thalamocortical cells of the ventromedial thalamic nucleus. *Brain Research 334*, 2 (may 1985), 227–233.

[20] Douglas, R. J., and Martin, K. A. Mapping the Matrix: The Ways of Neocortex. *Neuron 56*, 2 (oct 2007), 226–238.

[21] El-Samad, H. J., Goff, J. P., and Khamash, M. H. Calcium Homeostasis and Parturient Hypocalcemia: An Integral Feedback Perspective. *Journal of Theoretical Biology 214*, 1 (jan 2002), 17–29.

[22] El-Samad, H. J., Kurata, H., Doyle, J. C., Gross, C. A., and Khamash, M. H. Surviving heat shock: Control strategies for robustness and performance. *Proceedings of the National Academy of Sciences 102*, 8 (jan 2005), 2736–2741.

[23] Enel, P., Procyk, E., Quilodran, R., and Dominey, P. F. Reservoir Computing Properties of Neural Dynamics in Prefrontal Cortex. *PLOS Computational Biology 12*, 6 (jun 2016), e1004967.

[24] Felleman, D. J., and Essen, D. C. V. Distributed Hierarchical Processing in the Primate Cerebral Cortex. *Cerebral Cortex 1*, 1 (jan 1991), 1–47.

[25] FitzGibbon, T., Eriköz, B., Grünert, U., and Martin, P. R. Analysis of the lateral geniculate nucleus in dichromatic and trichromatic marmosets. *Journal of Comparative Neurology 523*, 13 (jul 2015), 1948–1966.

[26] Funahashi, K., and Nakamura, Y. Approximation of dynamical systems by continuous time recurrent neural networks. *Neural Networks 6*, 6 (jan 1993), 801–806.

[27] Fusi, S., Miller, E. K., and Rigotti, M. Why neurons mix: high dimensionality for higher cognition. *Current Opinion in Neurobiology 37* (apr 2016), 66–74.

[28] Godfrey, P., Shipley, R., and Gryz, J. Algorithms and analyses for maximal vector computation. *The VLDB Journal 16*, 1 (sep 2006), 5–28.

[29] Goldberg, J. H., Farries, M. A., and Fee, M. S. Basal ganglia output to the thalamus: still a paradox. *Trends in Neurosciences 36*, 12 (dec 2013), 695–705.

[30] Grant, E., Hoerder-Suabedissen, A., and Molnár, Z. Development of the Corticothalamic Projections. *Frontiers in Neuroscience 6* (2012).

[31] Groh, A., Bokor, H., Mease, R. A., Plattner, V. M., Hangya, B., Stroh, A., Deschenes, M., and Acsády, L. Convergence of Cortical and Sensory Driver Inputs on Single Thalamocortical Cells. *Cerebral Cortex 24*, 12 (jul 2014), 3167–3179.

[32] Halassa, M. M., and Acsády, L. Thalamic Inhibition: Diverse Sources Diverse Scales. *Trends in Neurosciences 39*, 10 (oct 2016), 680–693.

[33] Halassa, M. M., and Kastner, S. Thalamic functions in distributed cognitive control. *Nature Neuroscience 20*, 12 (nov 2017), 1669–1679.

[34] Harth, E. M., Unnikrishnan, K. P., and Pandya, A. S. The

inversion of sensory processing by feedback pathways: a model of visual cognitive functions. *Science 237*, 4811 (jul 1987), 184–187.

[35] Heeger, D. J. Theory of cortical function. *Proceedings of the National Academy of Sciences 114*, 8 (feb 2017), 1773–1782.

[36] Hubel, D. H., and Wiesel, T. N. Receptive fields of single neurones in the cat's striate cortex. *The Journal of Physiology 148*, 3 (oct 1959), 574–591.

[37] Hubel, D. H., and Wiesel, T. N. Receptive fields binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology 160*, 1 (jan 1962), 106–154.

[38] Ikeda, K., and Matsumoto, K. High-dimensional chaotic behavior in systems with time-delayed feedback. *Physica D: Nonlinear Phenomena 29*, 1-2 (nov 1987), 223–235.

[39] Jaeger, H. The echo state approach to analysing and training recurrent neural networks. Tech. rep., 2001.

[40] Jaeger, H. Echo state network. *Scholarpedia 2*, 9 (2007), 2330.

[41] Jazayeri, M., and Shadlen, M. N. A Neural Mechanism for Sensing and Reproducing a Time Interval. *Current Biology 25*, 20 (oct 2015), 2599–2609.

[42] Jones, E. G. Functional subdivision and synaptic organization of the mammalian thalamus. *Int Rev Physiol 25* (1981), 173–245.

[43] Jones, E. G. Principles of Thalamic Organization. In *The Thalamus*. Springer US, 1985, pp. 85–149.

[44] Jones, E. G. Viewpoint: the core and matrix of thalamic organization. *Neuroscience 85*, 2 (apr 1998), 331–345.

[45] Kakei, S., Na, J., and Shinoda, Y. Thalamic terminal morphology and distribution of single corticothalamic axons originating from layers 5 and 6 of the cat motor cortex. *The Journal of Comparative Neurology 437*, 2 (Aug 2001), 170–185.

[46] Karten, H. J. Evolutionary developmental biology meets the brain: The origins of mammalian cortex. *Proceedings of the National Academy of Sciences 94*, 7 (apr 1997), 2800–2804.

[47] Komura, Y., Nikkuni, A., Hirashima, N., Uetake, T., and Miyamoto, A. Responses of pulvinar neurons reflect a subject's confidence in visual categorization. *Nature Neuroscience 16*, 6 (may 2013), 749–755.

[48] Krishna, S., Maslov, S., and Sneppen, K. UV-Induced Mutagenesis in Escherichia coli SOS Response: A Quantitative Model. *PLoS Computational Biology 3*, 3 (2007), e41.

[49] Kung, H.-T., Luccio, F., and Preparata, F. P. On Finding the Maxima of a Set of Vectors. *Journal of the ACM 22*, 4 (oct 1975), 469–476.

[50] Kuramoto, E., Pan, S., Furuta, T., Tanaka, Y. R., Iwai, H., Yamanaka, A., Ohno, S., Kaneko, T., Goto, T., and Hioki, H. Individual mediodorsal thalamic neurons project to multiple areas of the rat prefrontal cortex: A single neuron-tracing study using virus vectors. *Journal of Comparative Neurology 525*, 1 (jul 2017), 166–185.

[51] Laughlin, S. B., and Sejnowski, T. J. Communication in neuronal networks. *Science 301* (Sep 2003), 1870–4.

[52] Lennie, P. The Cost of Cortical Computation. *Current Biology 13*, 6 (mar 2003), 493–497.

[53] Lin, H. W., Tegmark, M., and Rolnick, D. Why Does Deep and Cheap Learning Work So Well? *Journal of Statistical Physics 168*, 6 (jul 2017), 1223–1247.

[54] Ma, W. J., and Jazayeri, M. Neural Coding of Uncertainty and Probability. *Annual Review of Neuroscience 37*, 1 (jul 2014), 205–220.

[55] Maass, W., Natschlaeger, T., and Markram, H. Computational Models for Generic Cortical Microcircuits. In *Computational Neuroscience*. Chapman and Hall/CRC, oct 2003.

[56] Maass, W., Natschläger, T., and Markram, H. Real-Time Computing Without Stable States: A New Framework for Neural Computation Based on Perturbations. *Neural Computation 14*, 11 (nov 2002), 2531–2560.

[57] Mante, V., Sussillo, D., Shenoy, K. V., and Newsome, W. T. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature 503*, 7474 (nov 2013), 78–84.

[58] Marblestone, A. H., Wayne, G., and Kording, K. P. Toward an Integration of Deep Learning and Neuroscience. *Frontiers in Computational Neuroscience 10* (sep 2016).

[59] Marenco, S., Stein, J. L., Savostyanova, A. A., Sambataro, F., Tan, H.-Y., Goldman, A. L., Verchinski, B. A., Barnett, A. S., Dickinson, D., Apud, J. A., Callicott, J. H., Meyer-Lindenberg, A., and Weinberger, D. R. Investigation of Anatomical Thalamo-Cortical Connectivity and fMRI Activation in Schizophrenia. *Neuropsychopharmacology 37*, 2 (sep 2012), 499–507.

[60] Mathers, L. H. The synaptic organization of the cortical projection to the pulvinar of the squirrel monkey. *The Journal of Comparative Neurology 146*, 1 (sep 1972), 43–59.

[61] McCormick, D. A., and von Krosigk, M. Corticothalamic activation modulates thalamic firing through glutamate metabotropic receptors. *Proceedings of the National Academy of Sciences 89*, 7 (apr 1992), 2774–2778.

[62] Mesulam, M. M. Large-scale neurocognitive networks and distributed processing for attention, language, and memory. *Ann Neurol 28* (Nov 1990), 597–613.

[63] Middleton, F., and Strick, P. L. Anatomical evidence for cerebellar and basal ganglia involvement in higher cognitive function. *Science 266*, 5184 (oct 1994), 458–461.

[64] Mitchell, A. The mediodorsal thalamus as a higher order thalamic relay nucleus important for learning and decision-making. *Neurosci Biobehav Rev 54* (Jul 2015), 76–88.

[65] Mitelman, S. A., Byne, W., Kemether, E. M., Hazlett, E. A., and Buchsbaum, M. S. Metabolic Disconnection Between the Mediodorsal Nucleus of the Thalamus and Cortical Brodmann's Areas of the Left Hemisphere in Schizophrenia. *American Journal of Psychiatry 162*, 9 (sep 2005), 1733–1735.

[66] Mumford, D. On the computational architecture of the neocortex. I: The role of the thalamo-cortical loop. *Biological Cybernetics 65*, 2 (jun 1991), 135–145.

[67] Mumford, D. On the computational architecture of the neocortex. II The role of cortico-cortical loops. *Biological Cybernetics 66*, 3 (jan 1992), 241–251.

[68] Nair, A., Treiber, J. M., Shukla, D. K., Shih, P., and Müller, R.-A. Impaired thalamocortical connectivity in autism spectrum disorder: a study of functional and anatomical connectivity. *Brain 136*, 6 (jun 2013), 1942–1955.

[69] Nakajima, M., and Halassa, M. M. Thalamic control of functional cortical connectivity. *Current Opinion in Neurobiology 44* (jun 2017), 127–131.

[70] Navlakha, S., and Bar-Joseph, Z. Distributed information processing in biological and computational systems. *Communications of the ACM 58*, 1 (dec 2014), 94–102.

[71] Newell, A. Some problems of basic organization in problem solving programs, 1962.

[72] Nii, P. The Blackboard Model of Problem Solving and the Evolution of Blackboard Architectures. *The AI Magazine 7*, 2 (1986), 38–53.

[73] Niven, J. E., and Laughlin, S. B. Energy limitation as a selective pressure on the evolution of sensory systems. *Journal of Experimental Biology 211*, 11 (jun 2008), 1792–1804.

[74] Northcutt, R. G. Evolution of the Telencephalon in Nonmammals. *Annual Review of Neuroscience 4*, 1 (mar 1981), 301–350.

[75] Northcutt, R. G., and Kaas, J. H. The emergence and evolution of mammalian neocortex. *Trends in Neurosciences 18*, 9 (sep 1995), 373–379.

[76] Ohno, S., Kuramoto, E., Furuta, T., Hioki, H., Tanaka, Y., Fujiyama, F., Sonomura, T., Uemura, M., Sugiyama, K., and Kaneko, T. A morphological analysis of thalamocortical axon fibers of rat posterior thalamic nuclei: a single neuron tracing study with viral vectors. *Cereb Cortex 22* (Dec 2012), 2840–57.

[77] Parnaudeau, S., Bolkan, S. S., and Kellendonk, C. The Mediodorsal Thalamus: An Essential Partner of the Prefrontal Cortex for Cognition. *Biological Psychiatry* (nov 2017).

[78] Parnaudeau, S., Taylor, K., Bolkan, S. S., Ward, R. D.,

Balsam, P. D., and Kellendonk, C. Mediodorsal Thalamus Hypofunction Impairs Flexible Goal-Directed Behavior. *Biological Psychiatry 77*, 5 (mar 2015), 445–453.

[79] Partlow, G. D., Colonnier, M., and Szabo, J. Thalamic projections of the superior colliculus in the rhesus monkey,Macaca mulatta. A light and electron microscopic study. *The Journal of Comparative Neurology 171*, 3 (feb 1977), 285–317.

[80] Pinault, D. The thalamic reticular nucleus: structure function and concept. *Brain Research Reviews 46*, 1 (aug 2004), 1–31.

[81] Purushothaman, G., Marion, R., Li, K., and Casagrande, V. Gating and control of primary visual cortex by pulvinar. *Nat Neurosci 15* (Jun 2012), 905–12.

[82] Raczkowski, D., and Fitzpatrick, D. Terminal arbors of individual physiologically identified geniculocortical axons in the tree shrew's striate cortex. *The Journal of Comparative Neurology 302*, 3 (dec 1990), 500–514.

[83] Rigotti, M., Barak, O., Warden, M. R., Wang, X.-J., Daw, N. D., Miller, E. K., and Fusi, S. The importance of mixed selectivity in complex cognitive tasks. *Nature 497*, 7451 (may 2013), 585–590.

[84] Rigotti, M., Rubin, D. B. D., Wang, X.-J., and Fusi, S. Internal representation of task rules by recurrent dynamics: the importance of the diversity of neural responses. *Frontiers in Computational Neuroscience 4* (2010).

[85] Rouiller, E. M., and Welker, E. A comparative analysis of the morphology of corticothalamic projections in mammals. *Brain Research Bulletin 53*, 6 (dec 2000), 727–741.

[86] Rovo, Z., Ulbert, I., and Acsády, L. Drivers of the primate thalamus. *J Neurosci 32* (Dec 2012), 17894–908.

[87] Salas, C., Broglio, C., and Rodríguez, F. Evolution of Forebrain and Spatial Cognition in Vertebrates: Conservation across Diversity. *Brain Behavior and Evolution 62*, 2 (2003), 72–82.

[88] Schmitt, L. I., Wimmer, R. D., Nakajima, M., Happ, M., Mofakham, S., and Halassa, M. M. Thalamic amplification of cortical connectivity sustains attentional control. *Nature 545*, 7653 (may 2017), 219–223.

[89] Scott, B. B., Constantinople, C., Akrami, A., Hanks, T. D., Brody, C. D., and Tank, D. W. Fronto-parietal Cortical Circuits Encode Accumulated Evidence with a Diversity of Timescales. *Neuron 95* (Jul 2017), 385–398.e5.

[90] Selfridge, O. G. Pandemonium: a paradigm for learning in . In *Proceedings of the Symposium on Mechanisation of Thought Processes* (London, nov 1959), D. V. Blake and A. M. Uttley, Eds., National Physical Laboratory, pp. 513–526.

[91] Seoane, L. F., and Solé, R. Phase transitions in Pareto optimal complex networks. *Physical Review E 92*, 3 (sep 2015).

[92] Sherman, S. M. Thalamus plays a central role in ongoing cortical functioning. *Nature Neuroscience 16*, 4 (mar 2016), 533–541.

[93] Sherman, S. M., and Guillery, R. W. *Functional Connections of Cortical Areas*. The MIT Press, aug 2013.

[94] Shoham, S., O'Connor, D. H., and Segev, R. How silent is the brain: is there a "dark matter" problem in neuroscience? *Journal of Comparative Physiology A 192*, 8 (mar 2006), 777–784.

[95] Simon, H. A. The Architecture of Complexity. *Proceedings of the American Philosophical Society 106*, 6 (1962), 467–482.

[96] Simon, H. A. *The Sciences of the Artificial*. MIT Press, 1969, ch. The Architecture of Complexity.

[97] Steriade, M., Domich, L., and Oakson, G. Reticularis thalami neurons revisited: activity changes during shifts in states of vigilance. *J Neurosci 6* (Jan 1986), 68–81.

[98] Steriade, M., and Llinás, R. R. The functional states of the thalamus and the associated neuronal interplay. *Physiological Reviews 68*, 3 (jul 1988), 649–742.

[99] Steriade, M., and Pare, D. Morphology and electroresponsive properties of thalamic neurons. In *Gating in Cerebral Networks*. Cambridge University Press, 2007, pp. 1–26.

[100] Steriade, M., Parent, A., and Hada, J. Thalamic projections of nucleus reticularis thalami of cat: A study using retrograde transport of horseradish peroxidase and fluorescent tracers. *The Journal of Comparative Neurology 229*, 4 (nov 1984), 531–547.

[101] Szekely, P., Sheftel, H., Mayo, A., and Alon, U. Evolutionary Tradeoffs between Economy and Effectiveness in Biological Homeostasis Systems. *PLoS Computational Biology 9*, 8 (aug 2013), e1003163.

[102] Urbain, N., and Deschênes, M. Motor Cortex Gates Vibrissal Responses in a Thalamocortical Projection Pathway. *Neuron 56*, 4 (nov 2007), 714–725.

[103] Wang, J., Narain, D., Hosseini, E. A., and Jazayeri, M. Flexible timing by temporal scaling of cortical responses. *Nature Neuroscience* (dec 2018).

[104] Wang, X.-J. Neural dynamics and circuit mechanisms of decision-making. *Current Opinion in Neurobiology 22*, 6 (dec 2012), 1039–1046.

[105] Wei, W., and Wang, X.-J. Inhibitory Control in the Cortico-Basal Ganglia-Thalamocortical Loop: Complex Regulation and Interplay with Memory and Decision Processes. *Neuron 92*, 5 (dec 2016), 1093–1105.

[106] Wiecki, T. V., and Frank, M. J. A computational model of inhibitory control in frontal cortex and basal ganglia. *Psychological Review 120*, 2 (2013), 329–355.

[107] Wong, K.-F., and Wang, X.-J. A Recurrent Network Mechanism of Time Integration in Perceptual Decisions. *Journal of Neuroscience 26*, 4 (jan 2006), 1314–1328.

[108] Woodward, N. D., Giraldo-Chica, M., Rogers, B., and Cascio, C. J. Thalamocortical Dysconnectivity in Autism Spectrum Disorder: An Analysis of the Autism Brain Imaging Data Exchange. *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging 2*, 1 (jan 2017), 76–84.

[109] Yamins, D. L. K., and DiCarlo, J. J. Using goal-driven deep learning models to understand sensory cortex. *Nature Neuroscience 19*, 3 (feb 2016), 356–365.

[110] Yamins, D. L. K., Hong, H., Cadieu, C. F., Solomon, E. A., Seibert, D., and DiCarlo, J. J. Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proceedings of the National Academy of Sciences 111*, 23 (may 2014), 8619–8624.