

NEUROSCIENCE

Thalamic modulation of cortical linearity across arousal states

Eli J. Müller^{1,2*}, Brandon R. Munn^{1,2}, Giulia Baracchini^{1,2}, Ben D. Fulcher³, Vicente Medel^{4,5}, Michelle J. Redinbaugh^{6,7,8}, Yuri B. Saalmann^{7,8}, Jordan J. Wehrman^{9,10}, Robert D. Sanders^{9,10}, Bingni W. Brunton¹¹, Steven L. Brunton¹², James M. Shine^{1,2}

The human brain must support both stable and flexible neural dynamics to adapt to changing contexts. The thalamus has been hypothesized to coordinate these opposing dynamical regimes in the cerebral cortex; however, this has not been directly tested at the systems level. Here, we apply a time-resolved measure of linearity to multimodal recordings in humans and macaques, demonstrating irregular fluctuations in both human functional magnetic resonance imaging and electrophysiological in the awake state. These fluctuations follow a key neuroanatomical axis of the thalamus, are diminished under pharmacological ablation of arousal, and recovered following arousal-inducing direct electrical stimulation of thalamus. Together, these findings identify a thalamic organizational axis that modulates cortical linearity across arousal states, linking thalamocortical architecture to the regulation of large-scale cortical dynamics.

INTRODUCTION

The human brain supports a wide behavioral repertoire through coordinated interactions among distributed neuronal populations. For this activity to be adaptively useful, the dynamics that unfold over time must balance stability, supporting reliable execution, with flexibility, supporting adaptation to changing demands. How neuroanatomical organization regulates this trade-off remains incompletely understood.

Few existing analyses are capable of providing time-resolved estimates of the stability (and flexibility) of brain dynamics that are sensitive to the waxing and waning nature of the awake brain (1). Most current statistical approaches (2–9) have heavily relied on assumptions of stationarity, ignoring information in moment-to-moment dynamic shifts. While some approaches relax stationarity by analyzing shorter, sliding windows (10, 11) or by tracking rapid changes in activity (12–14), these methods primarily quantify changes in summary statistics (e.g., covariance or amplitude) rather than the governing dynamics, limiting direct interpretation in terms of stability and flexibility.

Recent work has begun to track moment-to-moment changes in macroscale spatiotemporal brain dynamics (15, 16). Linear autoregressive models can outperform more complex nonlinear alternatives in explaining both resting-state human electrophysiological (EEG) and blood-oxygen-level-dependent (BOLD) functional magnetic resonance imaging (fMRI) activity (15), supporting previous evidence for approximately linear large-scale dynamics (17–19). In

contrast, estimates of turbulent dynamics, a fundamentally nonlinear process, also vary across modalities and cognitive states (16, 20). Together, these findings suggest that macroscale recordings can exhibit both linear and nonlinear structure, motivating methods that can reconcile these descriptions within a single framework.

Many natural dynamical systems express both linear and nonlinear behavior in a context-dependent manner. Rather than exhibiting their complete repertoire of behaviors at any given moment, a subset of their possible dynamics is typically dominant (21). For instance, many fluids exhibit linear dynamics near boundaries or walls, and turbulence at other locations. A simple pendulum is well approximated as linear for small angle amplitudes, but nonlinear for larger ones. Similarly, a single neuron's dynamics may be dominated by fast sodium currents for one period but by slower calcium currents at other times (21). An appreciation of this spatiotemporally varying behavior can be applied to a complex system (such as the brain) to unpack shifts in dominant dynamics, thus providing insight into the dominant physical processes at play.

From this perspective, we can naturally ask what role neuroanatomical constraints play in shaping the moment-to-moment expression of linear (and nonlinear) dynamics. The thalamus is a highly conserved subcortical structure that has been argued to play a central role in shaping and constraining the adaptive neural dynamics associated with wakefulness (22–25). At the microscale, specific thalamic cell types relay sensory information (26, 27), coordinate motor actions (28), and integrate subcortical input into distributed cortico-thalamic networks (29–31). At the macroscale, diffusely projecting subtypes of thalamocortical nuclei have been linked to increased variability of cortical dynamics (29–33, 34), supporting shifts in cognition (25, 35) and arousal (33, 36–39). Together, these perspectives place the thalamus in a privileged position to both drive the cerebral cortex with targeted inputs and modulate ongoing dynamics through diffuse projections (25, 29, 33, 40). Despite this, there are few direct empirical links between thalamocortical organization and moment-to-moment stability-flexibility trade-offs in large-scale cortical dynamics (22).

Here, motivated by the hypothesis that thalamocortical organization regulates moment-to-moment stability-flexibility trade-offs

¹School of Medical Sciences, Faculty of Medicine and Health, The University of Sydney, Sydney, NSW, Australia. ²Centre for Complex Systems, The University of Sydney, Sydney, NSW, Australia. ³School of Physics, The University of Sydney, Sydney, NSW, Australia. ⁴Latin American Brain Health Institute, Universidad Adolfo Ibáñez, Santiago, Chile. ⁵Facultad de Ciencias Biológicas, Pontificia Universidad Católica de Chile, Santiago, Chile. ⁶Department of Biology, Stanford University, Stanford, CA, USA. ⁷Department of Psychology, University of Wisconsin–Madison, Madison, WI, USA. ⁸Wisconsin National Primate Research Center, Madison, WI, USA. ⁹Central Clinical School, Faculty of Medicine and Health, The University of Sydney, NSW, Australia. ¹⁰Department of Anaesthetics, Royal Prince Alfred Hospital, Sydney Local Health District, NSW, Australia. ¹¹Department of Biology, University of Washington, Seattle, WA, USA. ¹²Department of Mechanical Engineering, University of Washington, Seattle, WA, USA.

*Corresponding author. Email: eli.muller@sydney.edu.au

in large-scale cortical dynamics that we introduce a time-resolved measure of linear predictability. Specifically, we quantify how well a global linear autoregressive model forecasts successive time-points in multimodal neural recordings from humans and macaques. We show that neural activity fluctuates irregularly through high and low linear predictability and that these fluctuations are not explained by

matched surrogate data that separately preserve covariance structure and the Fourier power spectrum (Fig. 1). We then decompose the fitted linear dynamics into a spectrum of dynamic modes, revealing coherent spatiotemporal structure underlying predictability. Next, we show that propofol-induced reduction of arousal is associated with increased linear predictability in both human BOLD and




Fig. 1. Dynamic linear predictability in large-scale brain recordings. (A) Natural systems can transition between locally linear and nonlinear regimes: We use this natural analogy to motivate a time-resolved assay of cortical predictability. (B) Linear forecast analysis (LFA): A global linear propagator is estimated in a low-dimensional state space and used to forecast future brain states from each time point. Prediction error [mean squared error (MSE)] is computed between the linear model forecast and observed states across a range of forecast horizons (τ), yielding a horizon plot (time $\times$ τ) that summarizes when ongoing dynamics are well captured by a global linear model versus when they deviate. An example horizon plot from one subject (7-T resting-state fMRI) shows intermittent increases in prediction error, especially at longer horizons. (C) Empirical bounds using matched surrogates: time-averaged prediction errors across horizons (the linearity horizon) for empirical 7-T resting-state fMRI ($N = 59$, green) compared with: (i) a phase-randomized surrogate preserving covariance structure and Fourier power spectrum (orange); and (ii) a time-shuffled surrogate preserving covariance structure only (see Materials and Methods; figs. S1 and S2). Empirical dynamics are more predictable than the shuffled surrogate but less predictable than the phase-randomized surrogate, indicating that covariance and Fourier power spectrum are insufficient to explain observed predictability. (D) Dynamic mode decomposition (DMD) of the linear propagator expresses predicted dynamics as a superposition of spatiotemporal modes with characteristic timescales. (E to G) Example empirical dataset: (E) top principal component amplitude time series; (F) covariance and estimated linear propagator matrices, for example, subject; and (G) distribution of DMD timescales across all $n = 59$ subjects. (H to M) surrogate comparisons for the same example subject shown in (E): phase-randomized surrogate (H to J) and time-shuffled surrogate (K to M) with corresponding covariance/propagator matrices and modal timescale spectra. Note that axis limits differ across (G, J, and M), reflecting systematic changes in the modal timescales under different temporal constraints.

EEG data (Fig. 2). Last, we demonstrate that distinct thalamocortical projection systems differentially modulate cortical predictability and that causal thalamic stimulation can partially restore awake-like departures from linear predictability induced by anesthesia (Fig. 3). Together, these findings implicate thalamocortical organization as a system-level mechanism for regulating global linearity in large-scale cortical dynamics and propose testable predictions for its role in stability-flexibility trade-offs during behavioral task performance.

RESULTS

Linear dynamics in large-scale cortical recordings

To quantify the influence of the thalamus on cortical dynamics, we required a time-varying measure of the stability of brain state dynamics. To this end, we applied a simple linear autoregressive model estimation technique (41) (Fig. 1 and fig. S1; see Materials and Methods) to multimodal, multispecies neural time series. First, a global linear propagator tracking moment-to-moment shifts in a recording is generated for each time series (or epoch) by using the singular-value-decomposition (see fig. S1A). Then, for every time point in the recording, the linear propagator is used to forecast state dynamics at a set of future time points (τ).

The linear dynamics predicted by the propagator are compared to the observed dynamics using mean squared error (MSE), yielding a time-resolved measure of departure from adequacy of a single global linear propagator. Elevated error can reflect genuine nonlinearities, state-dependent regime switching, or unmodeled noise; as such, we interpret error operationally as reduced predictability under a fixed, global linear propagator and term this “linear predictability” in the following sections. While fitting separate models within shorted windows could improve local fit, we intentionally use a single global propagator to quantify when large-scale dynamics depart from a stationary linear description as this provides a consistent basis for testing system-level control principles.

We denote the time-resolved error as the forecast horizon of the linear model prediction (Fig. 1B), with the resultant graph denoted as a horizon plot. Periods of linearity will be well captured by the model, having a low error and a flat horizon, whereas periods that deviate from linearity will have a large error, and thus a steep horizon. A characteristic subject's horizon plot in Fig. 1B shows that human fMRI dynamics have time-varying fluctuations in linear predictability during rest. These results are observed for every individual in the dataset ($n = 59$; fig. S2), are robust across validation datasets (which had different temporal resolutions; see fig. S2), and do not covary with head motion [framework displacement showed no meaningful association; correlation coefficient (r) = 0.01 ± 0.03 across 10τ]. In addition, contrasting these findings against the time-varying displacement of the same signal shows that linear predictability in these data is not driven by autocorrelation effects (fig. S3). Later sections will show consistency in EEG recordings, mitigating against modality-specific sensitivities (Fig. 2).

To determine whether the observed fluctuations in global linearity were related to trivial features of the neural data, we applied the same approach to two separate surrogate time series that conserved distinct features of the empirical time series: a phase-randomized surrogate, which randomly shuffles the phases of all the oscillatory Fourier modes comprising the data while conserving the covariance structure and the Fourier power spectra; and a shuffled surrogate, which is simply the initial empirical fMRI data uniformly scrambled

in time conserving covariance structure only. From first principles, we expected that the phase-randomized surrogate should increase linear predictability, because the removal of phase-information ensures that the data are the summation of independent Fourier modes, whereas the shuffled surrogate should disrupt linear predictability, because all moment-to-moment information has been lost. This is precisely what we observed: Fig. 1C shows that these surrogate time series no longer match that seen in the empirical data. Specifically, the data are more linearly predictable in time than the shuffled surrogate but less predictable than the phased-randomized surrogate (Fig. 1C). Because both surrogates share the same covariance as the empirical data and the spectral surrogate additionally includes Fourier power spectral structure, we find that these time-averaged features are not sufficient to capture the observed level of global linearity seen in the neural recordings. This finding was consistent for all subjects ($N = 59$; fig. S2) and cross-validated in a 3-T resting-state fMRI dataset ($n = 14$; fig. S2).

These results support previous findings of strong linear dynamics in human neuroimaging (15, 17, 18) and large-scale EEG recordings (15, 19) but extend them by showing substantial yet irregular deviations from linear predictability in conscious states that is pervasive across subjects. We conclude that human brain dynamics show non-trivial yet constrained deviations from globally predictable linearity that may be indicative of nonlinear dynamics, nonstationarities or unmodeled noise.

Adaptive timescales and dynamic brain modes

We next asked what spatiotemporal structures support these fluctuations by leveraging dynamic mode decomposition (DMD). Similar to the way that singular value decomposition can be used to provide a compact representation of covariance as the weighted sum of independent spatial modes, DMD efficiently characterizes the dynamics of time-varying processes through eigendecomposition of the linear propagator matrix (41–44). This effectively reformulates the linear dynamics within a system as a function of a compressed set of spatiotemporal modes (42, 43, 45). Each dynamic mode contains four key characteristic properties: The real part of the eigenvector defines a spatial coherence map [$\text{real}(\mathbf{W})$; Fig. 1D, red/blue surface plot]; the imaginary part of the eigenvector is known as the dephasing map and represents the spatial pattern of delays relative to each of these coherence patterns [$\text{imag}(\mathbf{W})$; Fig. 1D, green surface plot]; additionally, each mode also has a corresponding eigenvalue defining its oscillatory frequency [$\text{imag}(\mathbf{\Lambda})$; Fig. 1D, purple line]; and a damping rate that captures the modes' temporal duration [$\text{real}(\mathbf{\Lambda})$; Fig. 1D, red line]. These features provide a concise representation of a systems' dynamic modes that can be used to effectively characterize its linear behavior.

Through the lens of DMD, we can identify the ways in which empirical fMRI data are distinct from the surrogates. Compared to both surrogates, fMRI data consist of a set of slowly fluctuating yet sustained spatiotemporal dynamic modes (Fig. 1G). That is, most of the dominant (highest power) dynamic modes across all subjects are concentrated in the lower left corner of a scatter plot comparing the frequency and damping rate (Fig. 1G). Due to its matched Fourier spectrum, the phase-randomized surrogate contains a similar distribution of modal frequencies (0 to 0.14 Hz) to that seen in the empirical data (Fig. 1J). However, these oscillations are more strongly damped as the spectral surrogate mostly excludes moment-to-moment information predictive of future states with only the time-averaged Fourier modes contributing to continuity of its temporal trajectories. In contrast, the shuffled surrogate displays a much wider range of high




Fig. 2. Linear predictability across states of arousal. (A) Schematic of fMRI time series. (B and C) Horizon plot of an example single subject resting-state human BOLD showing time-varying linear error during wake (B) and during propofol anesthesia (C). (D) Population average linear error of BOLD resting-state data during wake (black) and propofol (dark blue) conditions averaged over each recording: Shaded regions show SEM (paired *t* test used for significance). (E and F) Timescale spectra of the dynamic modes of the linear propagators for all $n = 14$ BOLD subject recordings. (G) Schematic of EEG time series. Modified from BioRender. E. J. Müller (2026), <https://BioRender.com/7dt8ar2>. (H and I) Horizon plot of an example single subject resting-state human EEG showing time-varying linear error during wake (H) and during propofol anesthesia (I). (J) Average linear error of human EEG during wake (black; six subjects with three separate 100-s recordings) and propofol (dark blue; three subjects with nine separate 100-s recordings) conditions: Shaded regions show SEM (paired *t* test used for significance). (K and L) Timescale spectra of the dynamic modes of the linear propagators for all EEG recordings. *** $p \leq 0.001$.




Fig. 3. Thalamic influence on cortical linear predictability. (A) Schematic of characteristic projection profiles of targeted (core) and diffuse (matrix) thalamocortical projection profiles [adapted from (56)]. (B) Core-matrix thalamocortical expression gradient [adapted from (30)]. (C) Population average cortical map showing regions of dominant linear dynamics 10 time points into the future. (D) Thalamic time series for each subject are projected along the core-matrix gradient and epochs showing that one SD above the mean (high matrix) and below the mean (high core) is collated. (E) The difference of the MSE of the linear model following high matrix compared to high core thalamic activity. (F) Experimental setup, in which a macaque monkey had 50-Hz electrical stimulation applied to the matrix-rich central lateral thalamus while local field potentials (LFPs) were recorded from the frontal eye fields (FEFs) and lateral intraparietal (LIP) cortex. Modified from BioRender. E. J. Müller (2026), <https://BioRender.com/9ppa41h>. (G) Examination of the empirical recordings from macaque cortex recovered the predicted relationships: Namely, central lateral thalamic stimulation recovered a drop in cortical linear predictability that matched the awake brain but was lost during propofol anesthesia. * $P \leq 0.05$; ** $P \leq 0.01$; *** $P \leq 0.001$.

frequency (0 to 0.8 Hz), strongly damped modes than either the empirical data or the phase-randomized surrogate (Fig. 1M). This is due to the shuffled surrogate's time-point scrambling favoring fast and short-lived shifts in moment-to-moment activity.

Together, these spatial and temporal properties constitute the globally linear dynamics inherent in these human fMRI data and demonstrate how decomposition of the dynamic modes can provide insight into ongoing linear dynamics. Having characterized these linear cortical dynamics during conscious states and to determine the role of the thalamus in supporting them, we next examined how causal pharmacological manipulations via propofol anesthesia affect linear predictability.

Anesthesia promotes linear predictability in the cerebral cortex

Given the highly susceptible, flexible, and context-dependent nature of consciousness (46–48), we reasoned that the conscious awake

state should be associated with irregular deviations from linear predictability, which would then be diminished during anesthesia. To test this hypothesis, we extended our approach to neural recordings from two distinct measurement modalities: human neuroimaging data and human EEG data. Neuroimaging and EEG recordings are sensitive to distinct biophysical processes, namely, hemodynamics and extracellular electric field potentials, respectively. This dual approach ensures that our findings generalize beyond the nuances of each unique measurement modality, each of which captures unique features of neural activity (49).

First, we tracked the linear predictability of an existing human neuroimaging resting-state fMRI dataset of 14 healthy volunteers [four women; mean age, 24 years, SD = 5 (50)] who were administered propofol that induced conscious anesthesia. As predicted, we found that anesthesia disrupted deviations from linearity compared to the awake state, with the linear forecasts performing substantially better across all subjects following propofol administration (Fig. 2D;

$P < 0.001$). This population-level observation can be readily appreciated when viewing the Horizon plot of a characteristic subject's recording, where the wake dynamics show periods of increased linear predictability (Fig. 2B), which are equivalent to those observed for the same subject under propofol anesthesia (Fig. 2C). But in contrast, the awake data are punctuated by periods of sharp deviations from linearity that are not observed under anesthesia (Fig. 2B). Furthermore, despite anesthesia being more linearly predictable than wake, we find that the linear timescales of the BOLD signal are similar between these two states (Fig. 2, E and F). This observation supports previous findings of reduced variability and complexity of brain dynamics during anesthesia (51–53) but explains their spatio-temporal structure as a shift toward more globally linear dynamics.

Next, we tracked the linear predictability of a human EEG dataset during the resting-state ($n = 6$; 10× epochs of 10 s) and under moderate propofol anesthesia ($n = 3$; 10× epochs of 10 s). Despite substantial differences in the biophysical and temporal sensitivity of EEG versus fMRI, we observed a similar reduction in deviations from linear predictability under propofol anesthesia (Fig. 2H), albeit within a particular time horizon range. This difference between wake and propofol is seen at ~80 to 160 ms ($P < 0.001$; paired t test), which corresponds to a timescale within the alpha frequency band (8 to 12 Hz), a known spectral feature prominent in EEG during conscious arousal (2, 54). To address the low ($n = 3$) subject count for the propofol dataset, we cross-validated this finding in an independent EEG propofol dataset (55) with similar anesthetic concentration (fig. S4), which showed linear predictability differences over a broader time horizon (~80 to 400 ms; statistics were conducted at the subject level). Again, an example subjects' time-varying predictability shows that the conscious awake state is punctuated by periods of sharp deviation from global linear predictability not seen under propofol anesthesia (Fig. 2, H and I). The linear models for the fMRI and EEG data are necessarily distinct due to unique information each contains; however, the key finding is that both modalities show a consistent directionality of arousal-state effects.

Last, a key strength of the EEG signal is its temporal resolution, which allows tracking of broadband frequency content unavailable from the narrowband BOLD signal. This added resolution reveals dynamic timescales of the linear spatiotemporal dynamics concentrate within canonical alpha frequency ranges (~8 to 12 Hz; Fig. 2, K to L), showing how previous insights of covariance and frequency spectrum can be combined to differentiate wake and propofol states.

Together, these empirical recordings from distinct modalities demonstrate convergent evidence that the awake state is concomitant with fluctuations in global linear predictability across the cerebral cortex, which are diminished by propofol anesthesia. That is, tracking moment-to-moment shifts in the organization of brain dynamics can differentiate cognitive states. However, the question remains as to whether the thalamus itself plays a causal role in modulating these dynamics.

Thalamic axis of cortical linear predictability

Armed with this simple yet effective means for tracking fluctuations in brain states in neural data, we were now able to test our original hypothesis: namely, that axes of anatomical variation in the thalamus facilitate shifts in dynamical modes of the cerebral cortex. The thalamus itself comprises a rich diversity of cell types that have unique projection profiles to both the cerebral cortex and other subcortical structures (40, 56, 57). Although there is no single schema that effectively

captures the diversity of thalamocortical interactions (31, 40, 58–60), recent empirical studies in rodents have demonstrated the utility of considering a nonbinary continuous spectrum of thalamic cells that differ in terms of the topography and targets of their axonal arbors (61). At one end of this continuum sit the calbindin-rich matrix nuclei, which have thalamocortical projections spanning multiple cortical areas, typically targeting superficial layers of agranular association cortices (Fig. 3, A and B). This organization contrasts with the parvalbumin-rich core nuclei, which typically target granular layers of the cerebral cortex, providing feed-forward input to cortical pyramidal neurons (Fig. 3, A and B) (30, 31).

While there is evidence for individual thalamic cells that contain axonal features of both core and matrix types [even in the same axon; (56)], there are good reasons to consider a spectrum that extends between the two extreme phenotypes (24, 30). For one, simplifying assumptions at the microscale are required to conduct tractable work at the macroscale of large-scale recordings (22, 30, 34, 36). Second, there is now robust evidence that core and matrix cell types also differ in their functionality (24): Recent empirical and theoretical work has shown that electrical stimulation of the central lateral thalamus, which itself is abundant with matrix nuclei, but not the ventral lateral thalamus (which is a predominantly core-rich nucleus) can drive recovery of consciousness in propofol-anesthetized macaque monkeys by leveraging the nonspecific diffuse projection profiles of these matrix cells (34, 36–39). From these empirical vantage points, we thus ask: What type of cortical dynamics is facilitated by the thalamus, and how do these dynamics relate to the expression of core and matrix thalamic cell subtypes?

To answer this question, we applied a multitiered approach. First, we used the difference between the ground-truth dynamics of the 7-T resting-state human neuroimaging dataset and its predicted linear dynamics, calculated for each region, each time point, and each subject ($n = 59$). We then sum these errors across regions at a forecast horizon $\tau = 10$ (ensuring sufficient time for autocorrelation to decay) across all time points and all subjects. This defines a cortical map revealing a spatial pattern of dominant globally linear dynamics (Fig. 3C). Figure 3C shows that linear predictability was predominant in sensory cortices and that these dynamics were strongly negatively correlated with the core-matrix expression gradient ($r = -0.56$; $P < 10^{-10}$; fig. S5A) (30). That is, cortical brain regions with strong core thalamocortical projections also showed strong linear predictability. However, both of these spatial maps also covary with the first principal component of the imaging data (fig. S5B), which itself captures a substantial proportion of the empirical time series' overall covariance, thus making strong inferences about the thalamic involvement difficult.

To test thalamic involvement more precisely, we asked whether fluctuations in cortical predictability were temporally aligned with a thalamic organizational architecture consistent with core-matrix architecture. Specifically, we derived a time-resolved core-matrix proxy by projecting each subject's thalamic fMRI time series onto a CALB-PVALB expression gradient approximated from Allen Human Brain gene-expression data, which captures a continuum of thalamic molecular organization and corresponds at its extremes to canonical matrix and core thalamic cell phenotypes (see Materials and Methods; Fig. 3D) (30).

This projection provides a system-level readout of thalamic activity biased along the core-matrix axis, rather than a direct measurement of cell-type firing: Positive values indicate relatively greater

expression-weighted activity in regions aligned with the CALB (matrix-associated) end of the gradient, whereas negative values indicate relatively greater activity aligned with the PVALB (core-associated) end (and vice versa). For each subject, we identified time points at which this matrix-core proxy exceeded ± 1 SD from its mean (high-matrix-biased versus high-core-biased epochs) and then compared cortical prediction error time-locked to these events. Averaging across subjects, the difference between matrix- and core-biased epochs revealed a systematic shift toward increased cortical linear predictability following core-biased thalamic epochs (Fig. 3E), consistent with a temporal relationship in which core-aligned thalamic activity precedes more stable, globally predictable cortical dynamics. The matrix-core projection time series was not correlated with BOLD activity from the adrenergic locus coeruleus ($r = -0.18$, $P < 0.005$), cholinergic basal nucleus of Meynert ($r = -0.16$, $P < 0.005$), or their difference ($r = -0.01$, $P < 0.01$), suggesting that this effect was not driven by broader arousal fluctuations that drive increased cortical variability (62–64). Last, collating the linear model errors at the regional level for the high-matrix- versus high-core-biased epochs (using a time horizon $\tau \approx 9$ s, i.e., largest difference in Fig. 3E), shows errors concentrate in unimodal sensory cortices following matrix-biased activity (fig. S8).

These results provide empirical evidence for the hypothesis that targeted feed-forward projections from the core thalamus act as a feed-forward propagator of cortical activity and more diffusely projecting axons from matrix thalamic nuclei as modulators of ongoing cortical activity (37–39, 65). This integrates with a previous fMRI study (29) showing that cortical regions receiving more projections from the matrix thalamus had a greater flexibility in their coordination with other cortical regions (30). Together, these findings demonstrate that the thalamus can provide a stabilization of ongoing cortical dynamics in the form of stable linearity, which may support robust behaviors while also mediating the flexibility to depart from these linear dynamics, consistent with a stability-flexibility balance theorized as a quintessential signature of adaptive cognition.

Thalamic stimulation recovers cortical unpredictability across arousal states

Last, to confirm the role of the thalamus in modulating cortical linear predictability, we repurposed an existing dataset comprising multi-electrode recordings from the frontal eye fields (FEFs) and lateral intraparietal (LIP) cortical areas in macaque (Fig. 3F) (36, 37). Following propofol-mediated anesthesia, targeted 50-Hz electrical stimulation of the matrix-rich central lateral thalamus drove the recovery of conscious arousal (36, 37). Despite the fact that these data were collected in a different species and through a different measurement modality, i.e., chronically implanted cortical electrodes, we found the same shift between wake and propofol as in the human fMRI and EEG datasets. Namely, the induction of propofol increased linear predictability in the anesthetized monkeys (Fig. 3G). However, unlike the human fMRI and EEG recordings that spatially sample whole-brain scales, these macaque findings are isolated to these two specific regions.

Last, we calculated this same linear predictability before (pre), during, and after (post) electrical stimulation of the central thalamus (which was effective in recovering the monkey's arousal state; Fig. 3G). As predicted and consistent in humans, we observed substantial propofol-mediated linearity in the pre-stimulation window, a substantial decrease in linear predictability during stimulation (comparable to that seen in the wake condition), and a subsequent

return to relatively predictable, linear dynamics following stimulus cessation (Fig. 3G). This effect was consistent following increasing temporal downsampling (see fig. S6).

Together, our findings suggest that the matrix thalamus can drive deviations from linearly predictable cortical flow across different states of arousal and, in doing so, may provide an essential feature of flexible dynamics indicative of versatile and adaptive cognition.

DISCUSSION

Here, we showed that the thalamus modulates the manner in which cortical dynamics fluctuate in globally linear predictability across distinct states of consciousness. Using multimodal neuroimaging and EEG recordings in human and macaques, we demonstrated how a key principle of cellular organization within the thalamus modulates these ongoing linear dynamics and specifically how the diffusely projecting matrix thalamus destabilizes linear cortical flow. This approach provides a candidate neuroanatomical mechanism through which the thalamus may support adaptive cognitive benefits of the waking brain: namely, by facilitating brain dynamics that can support the formation of both stable and flexible brain states, which, in turn, can support reconfiguration as contingencies change across different behavioral contexts.

Previous work has shown how electrical stimulation of the thalamus can modulate conscious arousal (37, 38) and implicated the matrix thalamus as driving quasi-critical brain state formation (34, 36), thereby facilitating the recovery of awareness from anesthesia. In keeping with this, work in macaque monkeys has shown that electrical stimulation of the central thalamus is able to reverse the effect of propofol anesthesia and, thus, recover arousal (37–39). Specifically, the authors found that stimulation in the macaque thalamus led to the recovery of broad functional connectivity (38), alpha-band (8 to 14 Hz) coherence (37), and features of power spectra (39) in cortical recordings matching those observed in the wake state. Our present findings enrich these insights and suggest the matrix (diffuse)-core (targeted) cortical projections may play a key role in shifting the cortex toward or away from globally linear dynamics. A caveat of these findings is the utilization CALB-PVALB gene expression data to map thalamic regional BOLD activity to underlying neural activity of matrix-core cell-type activity, which characterizes a bias in activity across the thalamus as a whole system. While the macaque stimulation dataset allows us to explore the EEG impact of neural activity in the central lateral thalamus, a matrix-heavy structure, directly, further empirical work is needed to strengthen the explicit link to matrix-core cell-type activity.

Human neuroimaging studies have revealed a hierarchy of temporal and spatial autocorrelation scales across the cortex (66) and that these two features capture a large number of existing topological metrics (5). Our present findings expand these insights by quantifying intrinsic dynamic modal timescales that drive linear predictability in neural activity across imaging modalities. Our results also confirm the long-held notion that BOLD dynamics are predominantly constituted by slow and sustained patterns of activity [hence the utility of standard resting-state analyses (67, 68)] but extend our ability to characterize and explore the spatial and temporal properties of the modes shaping these dynamics. We also demonstrate that the same anatomical mechanism underpins similar observations in both BOLD and EEG data, which thus paves the way forward for future mechanistic work integrating findings in distinct measurement modalities. Furthermore,

our comparison of empirical imaging data to surrogates conserving covariance and Fourier spectra shows that these features are not sufficient to explain the observed linear predictability. An important limitation that we were unable to interrogate in these data is whether nonneural sources of noise, such as physiological activity or task contexts, alters the observed relationship between thalamic activity and cortical dynamics; however, this line of questioning is left for future, targeted experimental approaches.

It is worth noting that, in this implementation, a linear model is generated for each subject over the complete length of that recording. As with calculation of covariance and Fourier power spectrum, this may also be performed at a finer temporal granularity via windowing approaches: This will favor faster timescale effects over slower ones (69) and thus should depend on the investigators' question of interest. In this study, the imaging data were collected during the resting-state, but recordings with specific epochs of interest, such as during a task paradigm, could be served better by a piecewise approach defining the linear predictability instantiated by the brain within each epoch, particularly given the time-scale constraints imposed by the vast majority of task-based neuroimaging analysis protocols. The exploration of shifts in linear predictability during behavior then presents an exciting opportunity for future work to probe stability-flexibility trade-offs during task performance and its relationship to thalamocortical anatomical organization.

A previous study found linear models best capture resting-state macroscopic cortical dynamics (15), including several classical linear model types. In the present study, we did not compare different types of linear modeling approaches beyond autoregressive models, and future work is needed to determine whether unique linear modeling approaches can distinguish distinct aspects of the thalamocortical interaction. However, our findings are consistent with more a much more complicated biophysical corticothalamic model operating in a linear regime, where matrix thalamic activity drove departure from linearity toward quasi-critical states (34, 36), suggesting an effect independent of the linear modeling approach.

Last, our approach allows for the quantification of linear predictability within a given time series; however, it does not make explicit claims about the nature of the residual dynamics, which may be multifaceted. It is possible that deviations from linearity are driven by nonlinearities (in the mathematical sense) within the dynamical system becoming dominant, due to the presence of extrinsic or intrinsic drive (i.e., visual/auditory/olfactory stimuli in the brain), or through the instantiation of a different transient linear system on a shorter timescale relative to the epoch of time initially considered. These are exciting opportunities to be explored in future work.

In conclusion, we have demonstrated fluctuations in globally linear cortical dynamics in multimodal neuroimaging and EEG data across human and nonhuman primates differentiates arousal states and are modulated by the organization of the thalamus. In this way, we argue that this key feature of the brain's neurobiological architecture helps to shape the robust, yet flexible, adaptive neural dynamics required for effective cognitive function that define our waking lives.

MATERIALS AND METHODS

Human neuroimaging data

Sixty healthy adult participants (28 females; 18 to 33 years; right-handed) were recruited and the research was approved by The University of Queensland Human Research Ethics Committee. These

data were originally described in another study (70). A total of 1050 (~10 min) whole-brain 7-T resting-state fMRI echo planar images were acquired using a multiband sequence (acceleration factor, 5; voxels, 2 mm³; TR, 586 ms; TE, 23 ms; flip angle, 40°; FOV, 208 mm; and slices, 55). Structural images were also collected to assist functional data preprocessing (MP2RAGE sequence: voxels, 0.75 mm³; TR, 4300 ms; TE, 3.44 ms; and slices, 256). Each participant underwent a resting-state scan, followed by three 12-min runs of the behavioral performance of the Latin-squares task, although these task data were not used in the present study.

DICOM images were first converted to NIfTI format and realigned. T1 images were reoriented, skull stripped (FSL BET), and coregistered to the NIfTI functional images using statistical parametric mapping functions. Segmentation and the DARTEL algorithm were used to improve the estimation of nonneural signal in subject space and the spatial normalization. From each gray-matter voxel, the following signals were regressed: linear trends, signals from the six head-motion parameters (three translation and three rotation) and their temporal derivatives, white matter, and cerebrospinal fluid (CSF; estimated from single-subject masks of white matter and CSF). The aCompCor method (71) was used to regress out residual signal unrelated to neural activity (i.e., five principal components derived from noise regions of interest in which the time series data were unlikely to be modulated by neural activity). Participants with head displacement >3 mm in >5% of volumes in any one scan were excluded ($n = 5$). Following preprocessing, the mean time series was extracted from 400 predefined cortical parcels using the Schaefer atlas (72).

For the anesthesia BOLD fMRI dataset, we used data derived from previously published research works (50, 73) that have been openly shared on the OpenNeuro data repository (doi:10.18112/openneuro.ds003171.v2.0.0). The dataset comprises 14 healthy individuals (age, 24 ± 5 ; M/F, 13/4), all of whom were right-handed, were native English speakers, and had no recorded history of neurological disorders. Subjects were asked to relax and with their eyes open and not fall asleep. The original study obtained ethical approval from both the Health Sciences Research Ethics Board and Psychology Research Ethics Board of Western University (REB no. 104,755) and adhered to the principles outlined in the revised declaration of Helsinki (2000). We analyzed two conditions: awake (fully alert and communicative) and deep sedation, which was achieved with an initial target effect-site concentration of 0.6 µg/ml and oxygen titrated to maintain SpO₂ above 96%, with increments of 0.3 µg/ml with repeated assessments of responsiveness until deep sedation (Ramsey level 5).

Imaging was performed on a 3-T Siemens Tim Trio system with a 32-channel head coil. Subjects resting-state fMRI scans using BOLD EPI sequence (slices, 33; voxel size, 3 mm³; isotropic; interslice gap of 25%; TR, 2000 ms; TE, 30 ms; matrix size, 64 × 64; and FA, 75°). Resting-state scans had 256 volumes. Anatomical scans were also obtained using a T1-weighted three-dimensional (3D) Magnetization Prepared - Rapid Gradient Echo (MPRAGE) sequence (voxel size, 1 mm³; isotropic; TR, 2.3; TE, 4.25 ms; matrix size, 240 × 256 × 192; FA, 9°).

Preprocessing was completed using fMRIPrep standard pipeline, which involves the basic preprocessing steps (coregistration, normalization, unwarping, noise component extraction, segmentation, skull stripping, etc.). The extracted time series were denoised specifying motion and physiological signals from white matter and CSF

(66). For both the 7-T and 3-T fMRI datasets, a low-pass third-order Butterworth filter set at half the Nyquist frequency was applied to prevent linear model fits to frequencies on the timescale of one TR. Note that this only meaningfully affected the 3 T recording removing frame-to-frame noise driven effects but was integrated to both imaging datasets to ensure fair comparison.

Thalamic time series representing the “core” and “matrix” populations were generated by taking the dot-product between the voxelwise time series of thalamic BOLD signal (7-T dataset) and a standardized map of the relative mRNA expression levels for PVALB (core) and CALB1 (matrix) provided by the Allen Human Brain Atlas (30, 31, 74). Variogram modeling was applied before our obtaining the gene maps to correct for potential subject-level differences (75). Note that there are other calcium binding proteins with nontrivial expression in the thalamus (61, 76); however, these patterns were not considered here, as they have not been directly linked to the same core-matrix gradient in the thalamus (31). Note that we do not advocate for a strict dichotomy between core and matrix neurons in this work but rather that these two projection types represent anatomical extremes at either end of an approximately continuous spectrum (40, 56). See Müller *et al.* (30) for more details.

Human EEG recordings

As previously reported (77–79), 256 channel EEG recordings (250 Hz) were made during wakefulness and during propofol sedation titrated using a serial awakening paradigm to achieve reports of unconsciousness (lack of subjective experience). Six participants' data were included. Propofol was administered using a RUGLOOP target-controlled infusion (Ghent University, Ghent, Belgium) according to the model of Schnider and colleagues. As previously described, most of reports of unconsciousness occurred at doses >1 mcg/ml.

For LFA analysis, three distinct resting-state recordings divided into 10 epochs of 10 s each were used for each of the six subjects. Two resting-state recordings were excluded from one subject due to an average z -scored variance across all channels an order of magnitude smaller than for all 160 wake and 90 propofol epochs included. For the propofol analysis, three subjects were excluded as they reported a subjective experience following the awakening paradigm suggesting insufficient anesthetic concentration was used. Of the remaining subjects, a total of nine distinct recordings with 10 epochs of 10 s satisfied these criteria and were included for analysis.

A supplementary human EEG dataset was also used consisting of 20 participants (9 males and 11 females) (mean age, 30.85; SD = 10.98). Each experimental run began with an awake baseline period lasting 25 to 30 min following which a target-controlled infusion of propofol was commenced via a computerized syringe driver (Alaris Asena PK, Carefusion, Berkshire, UK). Three blood plasma levels were targeted: 0.6 μ g/ml (mild sedation), 1.2 μ g/ml (moderate sedation), and recovery from sedation. The state of mild sedation was aimed to engender a relaxed but still responsive behavioral state. At each target level, a period of 10 min was which behavioral tests and EEG measurements were commenced. From each participant, ~ 7 min of 128-channel high-density EEG data were collected at each level of sedation. EEG was measured in microvolts, sampled at 250 Hz and referenced to the vertex, using the Net Amps 300 amplifier (Electrical Geodesics Inc., Eugene, Oregon, USA). Participants had their eyes closed in a resting state during data collection. Data from 91 channels over the scalp surface were retained for further analysis. Channels on the neck, cheeks, and forehead, which tended

to contribute most of the movement-related noise, were excluded. Retained channels were filtered between 0.5 and 45 Hz and segmented into 10-s-long epochs. Each epoch thus generated was baseline corrected relative to the mean voltage over the entire epoch. Data containing excessive eye movement or muscular artefact were rejected by a quasi-automated procedure: Abnormally noisy channels and epochs were identified by calculating their normalized variance and then manually rejected or retained by visual inspection.

Macaque EEG dataset

We compared modeling data to electrophysiology data from our previous study (37). This dataset consists of simultaneous local field potential (LFP) recordings from the FEF, LIP area, and central lateral thalamus in the right hemisphere of two macaques (*Macaca mulatta*, 4.3 to 5.5 years old, 7.63 to 10.30 kg body weight). All procedures conformed to the National Institutes of Health's *Guide for the Care and Use of Laboratory Animals* and were approved by the University of Wisconsin–Madison Institutional Animal Care and Use Committee (G005777). We low-pass filtered LFPs to 250 Hz and then linearly detrended and extracted artifacts. Bipolar derivations of the LFPs were then calculated to minimize any possible effects of a common reference and volume conduction. Recordings were performed during general anesthesia—either isoflurane (0.8 to 1.5% on O₂ flow of 1 liter/min; nine sessions) or propofol (0.17 to 0.33 mg/kg per minute, intravenous; nine sessions)—or wakefulness, using 16- or 24-contact linear micro-electrode arrays (LMEAs; MicroProbes). The electrode contacts had a diameter of 12.5 μ m and a spacing of 200 μ m between contacts. These resting-state recordings were captured in conjunction with behavioral recordings while performing an oddball auditory task, although the task data was not analyzed here.

In previous work, members of our team performed deep brain stimulation via simultaneous stimulation of 16 contacts of the LMEAs in the central thalamus. We applied 400-ms biphasic pulses of 200 μ A, at 50-Hz stimulation frequency, for a total of 60-s stimulation duration for any given stimulation event. To localize LMEAs, we averaged two 3D T1-weighted structural images of the MRI-compatible electrodes in situ [inversion-recovery prepared gradient echo sequence with: FOV, 128 mm²; matrix, 256 $\times$ 256; number of slices, 166; 0.5 mm isotropic; TR, 9.68 ms; TE, 4.192 ms; flip angle, 12°; and inversion time (TI), 450 ms] and cross-validated the electrode location with electrode depth measurements during recordings/stimulation. We used an arousal index (0 to 10) based on eye openings, body movements, and vital signs, as well as EEG and EMG, to measure stimulation-induced changes in the level of arousal. We defined effective stimulations, increasing the level of arousal, as having an arousal index of ≥ 3 , which all experimenters could differentiate from ineffective stimulations having an arousal index of 0 to 2. See (37, 80–82) for more details. The University of Wisconsin–Madison Institutional Animal Care and Use Committee approved all procedures, which conformed to the National Institutes of Health's *Guide for the Care and Use of Laboratory Animals*. Due to hardware constraints, the LFP time series were downsampled from 1 kHz to 100 Hz for linear model estimation, although we note that results in Fig. 3G and subsequent conclusions are not timescale sensitive as shown in fig. S6. Furthermore, the 50-Hz stimulation, and the 60-Hz ac mains and its subharmonic 30-Hz frequencies were removed using a second-order Butterworth filter with a half-power of 2 Hz. Last, to standardize with the human EEG analysis, we applied a band pass

(0.5 to 45 Hz) using a third-order Butterworth filter, which will also dampen the effects of the mains and stimulation frequencies.

Linear forecast analysis

Here, we outline linear forecast analysis (LFA) as applied to all datasets within this manuscript. First, each time series $\mathbf{Z} = [\mathbf{z}_1 \mathbf{z}_2 \mathbf{z}_3 \dots \mathbf{z}_n]$ with n observations of p variables is partitioned into a pre- $\mathbf{X}$ and post- $\mathbf{Y}$ time series, where

$$\mathbf{X} = [\mathbf{z}_1 \mathbf{z}_2 \mathbf{z}_3 \dots \mathbf{z}_{n-1}]$$

$$\mathbf{Y} = [\mathbf{z}_2 \mathbf{z}_3 \mathbf{z}_4 \dots \mathbf{z}_n]$$

with $\mathbf{X}$ lacking the final time point of $\mathbf{Z}$ and $\mathbf{Y}$ lacking the initial time point. Note that we have dropped the function-of-time notation for simplicity. Singular value decomposition is then applied, $\mathbf{X} = \mathbf{U} \times \mathbf{S} \times \mathbf{V}^T$, and a rank reduction of the data is made that captures 99% of the explained variance (results were stable to alterations of this threshold). Although not explicitly necessary for time series with substantially more time points than variables, some of the data used here violate this condition, and, thus, the results are standardized around this threshold; see fig. S7 showing truncation effects on BOLD data eigenspectrum. Furthermore, the EEG data showed substantial changes in dimensionality of covariance, and, thus, the highest number of components required to meet the 99% threshold was used to standardize across all data (60 components; results were stable to alterations of this threshold). A linear propagator matrix $\mathbf{A}$ can then be estimated from the data where $\mathbf{Y} = \mathbf{A}\mathbf{X}$ as follows

$$\mathbf{A} \approx \tilde{\mathbf{A}} = \mathbf{Y}\mathbf{X}^\dagger \triangleq \mathbf{Y}\mathbf{V}\mathbf{\Sigma}^{-1}\mathbf{U}^*$$

By stepping through each time point of $\mathbf{Z}$ projected into the SVD subspace, the linear propagator can be used to forecast the proceeding time points. MSE is used to compare between the predicted time series and the ground-truth time series

$$MSE_{t,\tau} = \langle |\mathbf{Z}_{t+\tau} - \tilde{\mathbf{Z}}_{t+\tau}|^2 \rangle_k$$

where $\mathbf{Z}$ is the ground-truth data, $\tilde{\mathbf{Z}}$ is the forecasted time series using the linear model, and k is the number of SVD modes. MATLAB code for linearity analysis can be found at <https://doi.org/10.5281/zenodo.19905547>.

Interpretation of horizon plots and summary linearity metrics

LFA quantifies the adequacy of a single, global linear propagator for forecasting the next brain state over multiple future horizons (τ). For each time point t , we apply the propagator to the current reduced state to generate forecasts $\hat{\mathbf{x}}_{t+\tau}$ and compute MSE against the observed state $\mathbf{x}_{t+\tau}$, producing a horizon plot (time $\times \tau$) that visualizes when ongoing dynamics are well captured by the global linear model. We use linear predictability to refer to low forecast error under this fixed global model; conversely, increased forecast error indicates departures from the global linear model's adequacy, which can arise from nonlinear dynamics, state-dependent regime switching, or unmodeled inputs/noise. To bound these interpretations, we compare empirical results against surrogate time series that preserve covariance structure (and, for phase-randomized surrogates, Fourier power spectrum) while disrupting higher-order temporal organization (see

the "Surrogate models" section). Unless otherwise stated, all LFA metrics are computed in the reduced SVD subspace used to estimate the propagator and are summarized either (i) as time-resolved horizon plots or (ii) by averaging error over time to obtain a horizon-dependent "linearity horizon" for group-level comparisons.

Dynamic mode decomposition

The eigendecomposition of the linear propagator $\mathbf{A}$ gives the dynamic modes of the system, $\mathbf{A}\mathbf{W} = \mathbf{W}\mathbf{\Lambda}$, where $\mathbf{W}$ is the matrix of eigenvectors and $\mathbf{\Lambda}$ is the diagonal matrix of eigenvalues λ_i . The DMD approximation of the data can then be written as the simple dynamic model

$$\hat{\mathbf{Z}}(t) = \Phi \exp(\Omega t) \mathbf{z}_i$$

where $\Omega \triangleq \mathbf{Y}\mathbf{V}\mathbf{\Sigma}^{-1}\mathbf{W}$ with $\Omega = \log(\mathbf{\Lambda}) / \Delta t$ and t is time (42). $\mathbf{z}_i$ is the initialization state and can be solved for each time point using the pseudoinverse.

Surrogate models

We used two surrogate models to test the validity of our approach. The first, referred to as the phase-randomized surrogate, randomly shuffles the phases of all the oscillatory Fourier modes comprising the data while conserving the covariance structure and the Fourier power spectra. This is a multivariate implementation of the adjusted-amplitude Fourier transform algorithm developed in signal analysis as a test for nonlinearities, whereby higher-order correlation structure is removed to linearize the data (83, 84). Crucially, this surrogate shares the empirical correlative structures between brain regions, which forms the basis of many approaches in the neuroimaging literature (9). In addition, this surrogate includes timescale information captured by the Fourier power spectrum, which is predictive of future dynamics, because the sinusoidal modes comprising this time averaged measure oscillate across time. But the relationship between these oscillations is randomized.

The second surrogate, referred to as the shuffled surrogate, is simply the initial empirical fMRI data uniformly scrambled in time, i.e., the regional activations relative to other regions at a given moment are conserved, matching empirical regional correlations, but, when they occur with respect to any other activations in time, is scrambled. As compared to the phase-randomized surrogate, the shuffled surrogate does not conserve the empirical data's Fourier power spectra. Because the empirical data and both surrogates have the same correlation structure, they also share the same principal components (i.e., covariance defining regional activation patterns). Despite these similarities, Fig. 1 (H and K) shows that the temporal expression of these patterns no longer matches that seen in the empirical data (Fig. 1E).

Supplementary Materials

This PDF file includes:

Figs. S1 to S8

REFERENCES

1. G. Deco, M. L. Kringelbach, V. K. Jirsa, P. Ritter, The dynamics of resting fluctuations in the brain: Metastability and its dynamical cortical core. *Sci. Rep.* **7**, 3095 (2017).
2. P. L. Purdon, E. T. Pierce, E. A. Mukamel, M. J. Preray, J. L. Walsh, K. F. K. Wong, A. F. Salazar-Gomez, P. G. Harrell, A. L. Sampson, A. Cimenser, S. Ching, N. J. Kopell, C. Tavares-Stoeckel, K. Habeeb, R. Merhar, E. N. Brown, Electroencephalogram signatures of loss and recovery of consciousness from propofol. *Proc. Natl. Acad. Sci. U.S.A.* **110**, E1142–E1151 (2013).

3. R. Liégeois, T. O. Laumann, A. Z. Snyder, J. Zhou, B. T. T. Yeo, Interpreting temporal fluctuations in resting-state functional connectivity MRI. *Neuroimage* **163**, 437–455 (2017).
4. T. O. Laumann, A. Z. Snyder, A. Mitra, E. M. Gordon, C. Gratton, B. Adeyemo, A. W. Gilmore, S. M. Nelson, J. J. Berg, D. J. Greene, J. E. McCarthy, E. Tagliazucchi, H. Laufs, B. L. Schlaggar, N. U. F. Dosenbach, S. E. Petersen, On the stability of BOLD fMRI correlations. *Cereb. Cortex* **27**, 4719–4732 (2017).
5. M. Shinn, A. Hu, L. Turner, S. Noble, K. H. Preller, J. L. Ji, F. Moujaes, S. Achard, D. Scheinost, R. T. Constable, J. H. Krystal, F. X. Vollenweider, D. Lee, A. Anticevic, E. T. Bullmore, J. D. Murray, Functional brain networks reflect spatial and temporal autocorrelation. *Nat. Neurosci.* **26**, 867–878 (2023).
6. D. T. Jones, P. Vemuri, M. C. Murphy, J. L. Gunter, M. L. Senjem, M. M. Machulda, S. A. Przybelski, B. E. Gregg, K. Kantarci, D. S. Knopman, B. F. Boeve, R. C. Petersen, C. R. Jack, Non-stationarity in the “resting brain’s” modular architecture. *PLOS ONE* **7**, e39731 (2012).
7. A. Abbas, M. Belloy, A. Kashyap, J. Billings, M. Nezafati, E. Schumacher, S. Keilholz, Quasi-periodic patterns contribute to functional connectivity in the brain. *Neuroimage* **119**, 193–204 (2019).
8. D. Vidaurre, S. M. Smith, M. W. Woolrich, Brain network dynamics are hierarchically organized in time. *Proc. Natl. Acad. Sci. U.S.A.* **114**, 12827–12832 (2017).
9. E. Bullmore, O. Sporns, The economy of brain network organization. *Nat. Rev. Neurosci.* **13**, 336–349 (2012).
10. D. J. Lurie, D. Kessler, D. S. Bassett, R. F. Betzel, M. Breakspear, S. Kheilholz, A. Kucyi, R. Liégeois, M. A. Lindquist, A. R. McIntosh, R. A. Poldrack, J. M. Shine, W. H. Thompson, N. Z. Bielewicz, L. Douw, D. Kraft, R. L. Miller, M. Muthuraman, L. Pasquini, A. Razi, D. Vidaurre, H. Xie, V. D. Calhoun, Questions and controversies in the study of time-varying functional connectivity in resting fMRI. *Netw. Neurosci.* **4**, 30–69 (2020).
11. D. S. Bassett, N. F. Wymbs, M. A. Porter, P. J. Mucha, J. M. Carlson, S. T. Grafton, Dynamic reconfiguration of human brain networks during learning. *Proc. Natl. Acad. Sci. U.S.A.* **108**, 7641–7646 (2011).
12. J. Faskowitz, F. Z. Esfahani, Y. Jo, O. Sporns, R. F. Betzel, Edge-centric functional network representations of human cerebral cortex reveal overlapping system-level architecture. *Nat. Neurosci.* **23**, 1644–1654 (2020).
13. F. I. Karahanoğlu, D. Van De Ville, Transient brain activity disentangles fMRI resting-state dynamics in terms of spatially and temporally overlapping networks. *Nat. Commun.* **6**, 7751 (2015).
14. J. M. Shine, O. Koyejo, P. T. Bell, K. J. Gorgolewski, M. Gilat, R. A. Poldrack, Estimation of dynamic functional connectivity using multiplication of temporal derivatives. *Neuroimage* **122**, 399–407 (2015).
15. E. Nozari, M. A. Bertolero, J. Stiso, L. Caciagli, E. J. Cornblath, X. He, A. S. Mahadevan, G. J. Pappas, D. S. Bassett, Macroscopic resting-state brain dynamics are best described by linear models. *Nat. Biomed. Eng.* **8**, 68–84 (2024).
16. G. Deco, M. L. Kringelbach, Turbulent-like dynamics in the human brain. *Cell Rep.* **33**, 108471 (2020).
17. T. Gautama, D. P. Mandic, M. M. Van Hulle, Signal nonlinearity in fMRI: A comparison between BOLD and MION. *IEEE Trans. Med. Imaging* **22**, 636–644 (2003).
18. E. Gultepe, B. He, A linear/nonlinear characterization of resting state brain networks in fMRI time series. *Brain Topogr.* **26**, 39–49 (2013).
19. K. J. Blinowska, M. Malinowski, Non-linear and linear forecasting of the EEG time series. *Biol. Cybern.* **66**, 159–165 (1991).
20. G. Deco, S. Liebana Garcia, Y. Sanz Perl, O. Sporns, M. L. Kringelbach, The effect of turbulence in brain dynamics information transfer measured with magnetoencephalography. *Commun. Phys.* **6**, 74 (2023).
21. J. L. Callahan, J. V. Koch, B. W. Brunton, J. N. Kutz, S. L. Brunton, Learning dominant physical processes with data-driven balance models. *Nat. Commun.* **12**, 1016 (2021).
22. J. M. Shine, L. D. Lewis, D. D. Garrett, K. Hwang, The impact of the human thalamus on brain-wide information processing. *Nat. Rev. Neurosci.* **24**, 416–430 (2023).
23. S. Mofakham, J. Robertson, N. Lubin, N. A. Cleri, C. B. Mikell, An unpredictable brain is a conscious, responsive brain. *J. Cogn. Neurosci.* **36**, 1643–1652 (2024).
24. C. J. Whyte, M. J. Redinbaugh, J. M. Shine, Y. B. Saalman, Thalamic contributions to the state and contents of consciousness. *Neuron* **112**, 1611–1625 (2024).
25. M. M. Halassa, S. Kastner, Thalamic functions in distributed cognitive control. *Nat. Neurosci.* **20**, 1669–1679 (2017).
26. R. W. Guillery, S. M. Sherman, Review thalamic relay functions and their role in corticocortical communication: Generalizations from the visual system. *Neuron* **33**, 163–175 (2002).
27. J. N. MacLean, B. O. Watson, G. B. Aaron, R. Yuste, Internal dynamics determine the cortical response to thalamic stimulation. *Neuron* **48**, 811–823 (2005).
28. Z. Gao, C. Davis, A. M. Thomas, M. N. Economo, A. M. Abrego, K. Svoboda, C. I. De Zeeuw, N. Li, A cortico-cerebellar loop for motor planning. *Nature* **563**, 113–116 (2018).
29. J. M. Shine, The thalamus integrates the macrosystems of the brain to facilitate complex, adaptive brain network dynamics. *Prog. Neurobiol.* **199**, 101951 (2021).
30. E. J. Müller, B. Munn, L. J. Hearne, J. B. Smith, B. Fulcher, A. Arnatkevičiūtė, D. J. Lurie, L. Cocchi, J. M. Shine, Core and matrix thalamic sub-populations relate to spatio-temporal cortical connectivity gradients. *Neuroimage* **222**, 117224 (2020).
31. E. G. Jones, The thalamic matrix and thalamocortical synchrony. *Trends Neurosci.* **24**, 595–601 (2001).
32. D. D. Garrett, S. M. Epp, A. Perry, U. Lindenberger, Local temporal variability reflects functional integration in the human brain. *Neuroimage* **183**, 776–787 (2018).
33. M. Bonjean, T. Baker, M. Bazhenov, S. Cash, E. Halgren, T. Sejnowski, Interactions between core and matrix thalamocortical projections in human sleep spindle synchronization. *J. Neurosci.* **32**, 5250–5263 (2012).
34. E. Müller, B. Munn, J. Shine, Diffuse neural coupling mediates complex network dynamics through the formation of quasi-critical brain states. *Nat. Commun.* **11**, 6337 (2020).
35. J. Q. Kosciessa, U. Lindenberger, D. D. Garrett, Thalamocortical excitability modulation guides human perception under uncertainty. *Nat. Commun.* **12**, 2430 (2021).
36. E. J. Müller, B. R. Munn, M. J. Redinbaugh, J. Lizio, M. Breakspear, Y. B. Saalman, J. M. Shine, The non-specific matrix thalamus facilitates the cortical information processing modes relevant for conscious awareness. *Cell Rep.* **42**, 112844 (2023).
37. M. J. Redinbaugh, J. M. Phillips, N. A. Kambi, S. Mohanta, S. Andryk, G. L. Dooley, M. Afrasiabi, A. Raz, Y. B. Saalman, Thalamus modulates consciousness via layer-specific control of cortex. *Neuron* **106**, 66–75.e12 (2020).
38. J. Tasserie, L. Uhrig, J. D. Sitt, D. Manasova, M. Dupont, S. Dehaene, B. Jarraya, Deep brain stimulation of the thalamus restores signatures of consciousness in a nonhuman primate model. *Sci. Adv.* **8**, eabl5547 (2022).
39. A. M. Bastos, J. A. Donoghue, S. L. Brincat, M. Mahnke, J. Yanar, J. Correa, A. S. Waite, M. Lundqvist, J. Roy, E. N. Brown, E. K. Miller, Neural effects of propofol-induced unconsciousness and its reversal using thalamic stimulation. *eLife* **10**, e60824 (2021).
40. M. M. Halassa, S. M. Sherman, Thalamocortical circuit motifs: A general framework. *Neuron* **103**, 762–770 (2019).
41. J. N. Kutz, S. L. Brunton, B. W. Brunton, J. L. Proctor, *Dynamic Mode Decomposition*. (Society for Industrial and Applied Mathematics, 2016).
42. B. W. Brunton, L. A. Johnson, J. G. Ojemann, J. N. Kutz, Extracting spatial-temporal coherent patterns in large-scale neural recordings using dynamic mode decomposition. *J. Neurosci. Methods* **258**, 1–15 (2016).
43. J. M. Kunert-Graf, K. M. Eschenburg, D. J. Galas, J. N. Kutz, S. D. Rane, B. W. Brunton, Extracting reproducible time-resolved resting state networks using dynamic mode decomposition. *Front. Comput. Neurosci.* **13**, 75 (2019).
44. P. J. Schmid, Dynamic mode decomposition of numerical and experimental data. *J. Fluid Mech.* **656**, 5–28 (2010).
45. J. L. Proctor, S. L. Brunton, J. N. Kutz, Dynamic mode decomposition with control. *SIAM J. Appl. Dyn. Syst.* **15**, 142–161 (2016).
46. D. C. Dennett, *Consciousness Explained* (Penguin UK, 1993).
47. G. Tononi, M. Boly, M. Massimini, C. Koch, Integrated information theory: From consciousness to its physical substrate. *Nat. Rev. Neurosci.* **17**, 450–461 (2016).
48. G. A. Mashour, P. Roelfsema, J.-P. Changeux, S. Dehaene, Conscious processing and the global neuronal workspace hypothesis. *Neuron* **105**, 776–798 (2020).
49. G. Baracchini, E. Muller, J. Shine, Bridging the epistemological divide in neuroscience to improve ontological clarity. *Aperture Neuro* **5**, 10.52294/001c.138841 (2025).
50. S. Kandeepan, J. Rudas, F. Gomez, B. Stojanowski, S. Valluri, A. M. Owen, L. Naci, E. S. Nichols, A. Soddu, Modeling an auditory stimulated brain under altered states of consciousness using the generalized Ising model. *Neuroimage* **223**, 117367 (2020).
51. M. Schartner, A. Seth, Q. Noirhomme, M. Boly, M. A. Bruno, S. Laureys, A. Barrett, Complexity of multi-dimensional spontaneous EEG decreases during propofol induced general anaesthesia. *PLOS ONE* **10**, e0133532 (2015).
52. J. Wang, G. J. Noh, B. M. Choi, S. W. Ku, P. Joo, W. S. Jung, S. Kim, H. Lee, Suppressed neural complexity during ketamine- and propofol-induced unconsciousness. *Neurosci. Lett.* **653**, 320–325 (2017).
53. T. F. Varley, A. I. Luppi, I. Pappas, L. Naci, R. Adapa, A. M. Owen, D. K. Menon, E. A. Stamatakis, Consciousness & brain functional complexity in propofol anaesthesia. *Sci. Rep.* **10**, 1018 (2020).
54. V. S. Weiner, D. W. Zhou, P. Kahali, E. P. Stephen, R. A. Peterfreund, L. S. Aglio, M. D. Szabo, E. N. Eskandar, A. F. Salazar-Gomez, A. L. Sampson, S. S. Cash, E. N. Brown, P. L. Purdon, Propofol disrupts alpha dynamics in functionally distinct thalamocortical networks during loss of consciousness. *Proc. Natl. Acad. Sci. U.S.A.* **120**, e2207831120 (2023).
55. S. Chennu, S. O’Connor, R. Adapa, D. K. Menon, T. A. Bekinschtein, Brain connectivity dissociates responsiveness from drug exposure during propofol-induced transitions of consciousness. *PLOS Comput. Biol.* **12**, e1004669 (2016).
56. F. Clascá, P. Rubio-Garrido, D. Jabaudon, Unveiling the diversity of thalamocortical neuron subtypes. *Eur. J. Neurosci.* **35**, 1524–1532 (2012).
57. L. Acsády, The thalamic paradox. *Nat. Neurosci.* **20**, 901–902 (2017).
58. S. M. Sherman, W. M. Usrey, A reconsideration of the core and matrix classification of thalamocortical projections. *J. Neurosci.* **44**, e0163242024 (2024).

59. E. G. Jones, Thalamic circuitry and thalamocortical synchrony. *Philos. Trans. R. Soc. B Biol. Sci.* **357**, 1659–1673 (2002).
60. E. G. Jones, Viewpoint: The core and matrix of thalamic organization. *Neuroscience* **85**, 331–345 (1998).
61. J. W. Phillips, A. Schulmann, E. Hara, J. Winnubst, C. Liu, V. Valakh, L. Wang, B. C. Shields, W. Korff, J. Chandrashekar, A. L. Lemire, B. Mensh, J. T. Dudman, S. B. Nelson, A. W. Hantman, A repeated molecular architecture across thalamic pathways. *Nat. Neurosci.* **22**, 1925–1935 (2019).
62. B. R. Munn, E. J. Müller, G. Wainstein, J. M. Shine, The ascending arousal system shapes neural dynamics to mediate awareness of cognitive states. *Nat. Commun.* **12**, 6016 (2021).
63. C. Grimm, S. N. Duss, M. Privitera, B. R. Munn, N. Karalis, S. Frässle, M. Wilhelm, T. Patriarchi, D. Razansky, N. Wenderoth, J. M. Shine, J. Bohacek, V. Zerbi, Tonic and burst-like locus coeruleus stimulation distinctly shift network activity across the cortical hierarchy. *Nat. Neurosci.* **27**, 2167–2177 (2024).
64. R. V. Raut, A. Z. Snyder, A. Mitra, D. Yellin, N. Fujii, R. Malach, M. E. Raichle, Global waves synchronize the brain's functional systems with fluctuating arousal. *Sci. Adv.* **7**, eabf2709 (2021).
65. M. Suzuki, M. E. Larkum, General anesthesia decouples cortical pyramidal neurons. *Cell* **180**, 666–676.e13 (2020).
66. J. Fallon, P. G. D. Ward, L. Parkes, S. Oldham, A. Arnatkevičiūtė, A. Fornito, B. D. Fulcher, Timescales of spontaneous fMRI fluctuations relate to structural connectivity in the brain. *Netw. Neurosci.* **4**, 788–806 (2020).
67. S. M. Smith, D. Vidaurre, C. F. Beckmann, M. F. Glasser, M. Jenkinson, K. L. Miller, T. E. Nichols, E. C. Robinson, G. Salimi-Khorshidi, M. W. Woolrich, D. M. Barch, K. Uğurbil, D. C. Van Essen, Functional connectomics from resting-state fMRI. *Trends Cogn. Sci.* **17**, 666–682 (2013).
68. S. M. Smith, P. T. Fox, K. L. Miller, D. C. Glahn, P. M. Fox, C. E. Mackay, N. Filippini, K. E. Watkins, R. Toro, A. R. Laird, C. F. Beckmann, Correspondence of the brain's functional architecture during activation and rest. *Proc. Natl. Acad. Sci. U.S.A.* **106**, 13040–13045 (2009).
69. J. N. Kutz, X. Fu, S. L. Brunton, Multiresolution dynamic mode decomposition. *SIAM J. Appl. Dyn. Syst.* **15**, 713–735 (2016).
70. L. J. Hearne, L. Cocchi, A. Zalesky, J. B. Mattingley, Reconfiguration of brain network architectures between resting-state and complexity-dependent cognitive reasoning. *J. Neurosci.* **37**, 8399–8411 (2017).
71. Y. Behzadi, K. Restom, J. Liao, T. T. Liu, A component based noise correction method (CompCor) for BOLD and perfusion based fMRI. *Neuroimage* **37**, 90–101 (2007).
72. A. Schaefer, R. Kong, E. M. Gordon, T. O. Laumann, X.-N. Zuo, A. J. Holmes, S. B. Eickhoff, B. T. T. Yeo, Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity MRI. *Cereb. Cortex* **28**, 3095–3114 (2018).
73. L. Naci, A. Haug, A. MacDonald, M. Anello, E. Houldin, S. Naqshbandi, L. E. Gonzalez-Lara, M. Arango, C. Harle, R. Cusack, A. M. Owen, Functional diversity of brain networks supports consciousness and verbal intelligence. *Sci. Rep.* **8**, 13259 (2018).
74. M. J. Hawrylycz, E. S. Lein, A. L. Guillozet-Bongaarts, E. H. Shen, L. Ng, J. A. Miller, L. N. Van De Lagemaat, K. A. Smith, A. Ebbert, Z. L. Riley, C. Abajian, C. F. Beckmann, A. Bernard, D. Bertagnoli, A. F. Boe, P. M. Cartagena, M. M. Chakravarty, M. Chapin, J. Chong, R. A. Dalley, B. D. Daly, C. Dang, S. Datta, N. Dee, T. A. Dolbeare, V. Faber, D. Feng, D. R. Fowler, J. Goldy, B. W. Gregor, Z. Haradon, D. R. Haynor, J. G. Hohmann, S. Horvath, R. E. Howard, A. Jeromin, J. M. Jochim, M. Kinnunen, C. Lau, E. T. Lazars, C. Lee, T. A. Lemon, L. Li, Y. Li, J. A. Morris, C. C. Overly, P. D. Parker, S. E. Parry, M. Reding, J. J. Royall, J. Schulkin, P. A. Sequeira, C. R. Slaughterbeck, S. C. Smith, A. J. Sodt, S. M. Sunkin, B. E. Swanson, M. P. Vawter, D. Williams, P. Wahnoutka, H. R. Zielke, D. H. Geschwind, P. R. Hof, S. M. Smith, C. Koch, S. G. N. Grant, A. R. Jones, An anatomically comprehensive atlas of the adult human brain transcriptome. *Nature* **489**, 391–399 (2012).
75. G. Gryglewski, R. Seiger, G. M. James, G. M. Godbersen, A. Komorowski, J. Unterholzner, P. Michenthaler, A. Hahn, W. Wadsak, M. Mitterhauser, S. Kasper, R. Lanzenberger, Spatial analysis and high resolution mapping of the human whole-brain transcriptome for integrative analysis in neuroimaging. *Neuroimage* **176**, 259–267 (2018).
76. M. C. Mü, H. J. Waldvogel, R. L. Faul, The distribution of calbindin, calretinin and parvalbumin immunoreactivity in the human thalamus. *J. Chem. Neuroanat.* **19**, 155–173 (2000).
77. C. Troyas, J. Sleight, Electroencephalographic signatures of consciousness: Uncovering the fake news. *Br. J. Anaesth.* **132**, 218–219 (2024).
78. C. P. Casey, S. Tanabe, Z. Farahbakhsh, M. Parker, A. Bo, M. White, T. Ballweg, A. McIntosh, W. Filbey, Y. Saalman, R. A. Pearce, R. D. Sanders, Distinct EEG signatures differentiate unconsciousness and disconnection during anaesthesia and sleep. *Br. J. Anaesth.* **128**, 1006–1018 (2022).
79. T. W. Schnider, C. F. Minto, P. L. Gambus, C. Andresen, D. B. Goodale, S. L. Shafer, E. J. Youngs, The influence of method of administration and covariates on the pharmacokinetics of propofol in adult volunteers. *Anesthesiology* **88**, 1170–1182 (1998).
80. M. J. Redinbaugh, M. Afrasiabi, J. M. Phillips, N. A. Kambi, S. Mohanta, A. Raz, Y. B. Saalman, Thalamic deep brain stimulation paradigm to reduce consciousness: Cortico-striatal dynamics implicated in mechanisms of consciousness. *PLOS Comput. Biol.* **18**, e1010294 (2022).
81. M. J. Redinbaugh, M. Afrasiabi, J. M. Phillips, N. A. Kambi, S. Mohanta, A. Raz, Y. B. Saalman, Thalamic deep brain stimulation as a paradigm to reduce consciousness: Implications for cortico-striatal dynamics, absence epilepsy and consciousness studies. *PLOS Comput. Biol.* **18**, e1010294 (2021).
82. M. J. Redinbaugh, M. Afrasiabi, J. M. Phillips, N. A. Kambi, S. Mohanta, A. Raz, A. M. Haun, Y. B. Saalman, Consciousness depends on integration between parietal cortex, striatum, and thalamus. *Cell Syst.* **12**, 363–373.e11 (2021).
83. J. Theuer, B. Galdrikian, A. Longtin, S. Eubank, J. Doynne Farmer, Testing for nonlinearity in time series: The method of surrogate data. *Phys. D Nonlinear Phenom.* **58**, 77–94 (1992).
84. C. Raeth, R. A. Monetti, "Surrogates with random Fourier phases," in *Topics on Chaotic Systems* (World Scientific Connect, 2008) pp. 274–285.

Acknowledgments

Funding: E.J.M. was supported by an ARC DECRA Fellowship (DE250100540). J.M.S. was supported by an NHMRC Investigator Fellowship (GNT2033162). Y.B.S. was supported by NIH grants R01MH110311 and P51OD011106 and US-Israel Binational Science Foundation grant 201732. B.R.M. was supported by an ARC DECRA Fellowship (DE260101935). V.M. was supported by ANID/FONDECYT de Iniciación (11251578), ANID/FONDECYT Exploración (13240170), and the Alzheimer's Association AARG-(25-1490306). **Author contributions:** Conceptualization: E.J.M., S.L.B., B.R.M., R.D.S., and J.M.S. Methodology: E.J.M. and J.M.S. Investigation: E.J.M., J.M.S., R.D.S., and M.J.R. Visualization: E.J.M. and J.M.S. Supervision: J.M.S. and R.D.S. Resources: E.J.M., R.D.S., and J.M.S. Funding acquisition: E.J.M., R.D.S., and J.M.S. Data curation: E.J.M., R.D.S., M.J.R., and J.M.S. Validation: E.J.M., R.D.S., and J.M.S. Visualization: E.J.M. and J.M.S. Project administration: E.J.M., R.D.S., and J.M.S. Writing—original draft: E.J.M. and J.M.S. Writing—review and editing: E.J.M., B.R.M., G.B., B.D.F., V.M., M.J.R., Y.B.S., J.J.W., R.D.S., B.W.B., S.L.B., and J.M.S. **Competing interests:** The authors declare that they have no competing interests. **Data, code, and materials availability:** All data and code needed to evaluate and reproduce the results in the paper are present in the paper and/or the Supplementary Materials. This study did not generate new materials. The human EEG data can be provided by R. A. Pearce (University of Wisconsin) pending scientific review and a completed material transfer agreement through University of Wisconsin. Requests for the EEG data should be submitted to r Pearce@wisc.edu.

Submitted 18 January 2026

Accepted 7 August 2026

Published 18 September 2026

10.1126/sciadv.aef5358

Thalamic modulation of cortical linearity across arousal states

Eli J. Müller, Brandon R. Munn, Giulia Baracchini, Ben D. Fulcher, Vicente Medel, Michelle J. Redinbaugh, Yuri B. Saalman, Jordan J. Wehrman, Robert D. Sanders, Bingni W. Brunton, Steven L. Brunton, and James M. Shine

Sci. Adv. **12** (38), eaef5358. DOI: 10.1126/sciadv.aef5358

View the article online

<https://www.science.org/doi/10.1126/sciadv.aef5358>

Permissions

<https://www.science.org/help/reprints-and-permissions>

Use of this article is subject to the [Terms of service](#)

Science Advances (ISSN 2375-2548) is published by the American Association for the Advancement of Science. 1200 New York Avenue NW, Washington, DC 20005. The title *Science Advances* is a registered trademark of AAAS.

Copyright © 2026 The Authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original U.S. Government Works. Distributed under a Creative Commons Attribution NonCommercial License 4.0 (CC BY-NC).