



Time-evolving controllability of effective connectivity networks during seizure progression

Brittany H. Scheid^{a,b}, Arian Ashourvan^{a,b}, Jennifer Stiso^{a,c}, Kathryn A. Davis^{a,b,d}, Fadi Mikhail^{a,b,d}, Fabio Pasqualetti^e, Brian Litt^{a,b,d}, and Danielle S. Bassett^{a,b,f,g,h,i,1}

^aDepartment of Bioengineering, University of Pennsylvania, Philadelphia, PA 19104; ^bCenter for Neuroengineering and Therapeutics, University of Pennsylvania, Philadelphia, PA 19104; ^cDepartment of Neuroscience, University of Pennsylvania, Philadelphia, PA 19104; ^dDepartment of Neurology, Hospital of the University of Pennsylvania, Philadelphia, PA 19104; ^eDepartment of Mechanical Engineering, University of California, Riverside, CA 92521; ^fDepartment of Electrical & Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104; ^gDepartment of Physics & Astronomy, University of Pennsylvania, Philadelphia, PA 19104; ^hDepartment of Psychiatry, Perelman School of Medicine, University of Pennsylvania, Philadelphia, PA 19104; and ⁱSanta Fe Institute, Santa Fe, NM 87501

Edited by György Buzsáki, New York University Langone Medical Center, New York, NY, and approved December 24, 2020 (received for review April 7, 2020)

Over one third of the estimated 3 million people with epilepsy in the United States are medication resistant. Responsive neurostimulation from chronically implanted electrodes provides a promising treatment alternative to resective surgery. However, determining optimal personalized stimulation parameters, including when and where to intervene to guarantee a positive patient outcome, is a major open challenge. Network neuroscience and control theory offer useful tools that may guide improvements in parameter selection for control of anomalous neural activity. Here we use a method to characterize dynamic controllability across consecutive effective connectivity (EC) networks based on regularized partial correlations between implanted electrodes during the onset, propagation, and termination regimes of 34 seizures. We estimate regularized partial correlation adjacency matrices from 1-s time windows of intracranial electrocorticography recordings using the Graphical Least Absolute Shrinkage and Selection Operator (GLASSO). Average and modal controllability metrics calculated from each resulting EC network track the time-varying controllability of the brain on an evolving landscape of conditionally dependent network interactions. We show that average controllability increases throughout a seizure and is negatively correlated with modal controllability throughout. Our results support the hypothesis that the energy required to drive the brain to a seizure-free state from an ictal state is smallest during seizure onset, yet we find that applying control energy at electrodes in the seizure onset zone may not always be energetically favorable. Our work suggests that a low-complexity model of time-evolving controllability may offer insights for developing and improving control strategies targeting seizure suppression.

effective connectivity | network topology | controllability | epilepsy | GLASSO

Responsive neurostimulation (RNS) is a nonresective treatment for medication-resistant epilepsy that aims to suppress seizures by delivering an electrical stimulus through intracranial electrodes in response to abnormal electrographic activity. Despite the fact that the first implantable RNS device was approved in 2013, the mechanism of action and optimal patient-specific stimulation settings remain unknown (1, 2). A better understanding of the cortical locations and patterns of brain activity that are most responsive to stimulation would improve treatment efficacy and efficiency. By modeling the brain as a network, with distinct brain regions representing nodes and with statistical dependencies representing measures of influence between regions as edges, the challenge of finding the best strategy for seizure suppression can be recast as a problem of network control (3, 4).

Evidence from prior empirical and computational studies suggests that a linear network control framework can be used to predict how neural activity will respond to an external stimulus, where the dynamics are constrained by a static, structural

brain network built from the pattern of white matter tract connections between large-scale cortical and subcortical anatomical regions (5–8). The theory also provides predictions of the optimal stimulation parameters required to drive the brain state, a time-dependent measure of activity at each brain region in the network, from an initial state to a desired target state through a series of state transitions along an optimal trajectory (9, 10). These predictions have been quantified for direct electrical stimulation targeted at improving memory (7) but have not yet been defined or tested for stimulation applied to controlling seizures for two main reasons. First, seizures are characterized by rapid state transitions that violate the assumption of linearity over the course of the whole seizure (11, 12). Second, it is theorized and observed that information propagation through the brain may depend on the underlying state of coherence between brain regions and that exogenous inputs to brain regions may bias the direction of information flow (13, 14). In light of these data, it is clear that informing linear estimates of controllability with static white matter networks fails to capture dynamic corticocortical influences over changing patterns of network coherence (15).

A number of nonlinear models exist that faithfully describe the intricacies of seizure progression and stimulation (16, 17). A linear control framework can approximate such models over short time horizons and can offer a greater repertoire of accessible tools for full network control (18). Therefore, we address the above limitations in an extension to the linear control framework; instead of assuming that a single, structural connectivity

Significance

Responsive neurostimulation is an increasingly accessible treatment for medication-resistant epilepsy that aims to suppress seizures using electrical stimulation from implanted intracranial electrodes. However, the optimal cortical location and time point for intervening once a seizure begins are not well understood. Here we represent a seizure as a series of effective connectivity networks over time and compute metrics of network controllability and optimal control energy. Our results allow us to characterize when and where the brain network may be the most responsive to an external stimulus.

Author contributions: B.H.S. and D.S.B. designed research; B.H.S. performed research; B.H.S., A.A., K.A.D., F.M., B.L., and D.S.B. analyzed data; and B.H.S., A.A., J.S., F.P., B.L., and D.S.B. wrote the paper.

The authors declare no competing interest.


This article is a PNAS Direct Submission.

This open access article is distributed under Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 (CC BY-NC-ND).

¹ To whom correspondence should be addressed. Email: dsb@seas.upenn.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2006436118/-/DCSupplemental>.

Published January 25, 2021.



network governs the transitions between brain states, we estimate a sequence of effective connectivity (EC) networks from intracranial electroencephalography (iEEG) seizure recordings to capture the evolving constraints on neural dynamics throughout a seizure. Our networks are constructed such that nodes represent iEEG electrodes and edges represent a measure of EC derived from the electrographic time series. Unlike functional connectivity measures, which describe only how activity between network nodes is correlated, EC implies underlying causality between nodes in a network (19). A number of data-driven methods including Granger causality, directed transfer function, and partial directed coherence have been used to create EC brain networks (20–22). In the context of this paper, EC describes the undirected conditional dependence of nodes in a network, given the activity of all other nodes, and is therefore a powerful approach for understanding the impact of an external stimulus on seizure dynamics (20) (Fig. 1). We define brain state at each second to be high- γ band power (30 to 150 Hz) at each electrode; this measure of brain activity is sensitive to frequency and amplitude changes accompanying ictal transitions in neural recordings (7, 23, 24). Our model supports the assumption that structural connectivity only partially dictates our ability to control rapidly evolving neural processes via external stimulation and that the transient coupling between brain regions also impacts neural controllability.

We use our model to build on prior work in network control theory to address several specific hypotheses regarding seizure dynamics in humans. We focus on two measures of controllability that characterize the ease of driving the brain into new states as time evolves (3, 25). Average controllability reflects the average input energy required at a node to transition the activity from some initial brain state to all nearby states (brain state vectors with a unit magnitude) (9, 26). The EC networks used in our model can be decomposed into sets of eigenmodes, representing invariant axes in state space upon which neural activity can be measured over time, and eigenvalues dictating the rate of activity decay along each mode. Modal controllability reflects the ease of moving a system to states dominated by small eigenmodes

that are intuitively difficult to reach because they contribute less to the dynamics of the system and are far from the dominant energy minimum (6, 9, 27). Specifically, the metric quantifies the extent to which input at a given node controls activity along all eigenmodes of the system, weighted toward smaller modes, thus summarizing how the underlying network will respond to control inputs. Modal controllability can be separated into the measures of persistent modal controllability and transient modal controllability, which describe the extent to which control input at a node perturbs the slow- or fast-decaying modes of the system, respectively (4, 7, 25). Collectively, these metrics allow a broad assessment of a system's accessible control strategies. We hypothesize that seizure propagation will be accompanied by a heightened average controllability and decremented modal controllability with respect to values at seizure onset, consistent with an enhanced strengthening of the epileptic network (9, 27). We extend this evaluation by using a model of seizure-regime-dependent linear control to find the optimal control function required to drive the brain to a seizure-free state from a subset of control nodes. We predict that controlling the nodes that define the observed seizure onset zone (SOZ) as epileptiform activity emerges will be the most effective. Although the true epileptogenic zone, defined as the brain region supporting seizure generation, may extend beyond the SOZ (28), our prediction is built upon prior observations that SOZ nodes precede the rest of the nodes in the constructed network in showing increased functional connectivity (29). We reason that this may facilitate local control before the seizure spreads. Broadly, our study builds on the principles of network control theory to study time-evolving coherence-mediated control, which has the potential to inform practical efforts to improve neurostimulation-based therapies.

Results

Characterizing Seizure Regimes. Prior studies of functional connectivity networks built with electrocorticography (ECoG) data suggest that seizures can be segmented into three main regimes—onset, propagation, and termination (12, 30, 31)—and that linear modeling is applicable within 1-s time windows

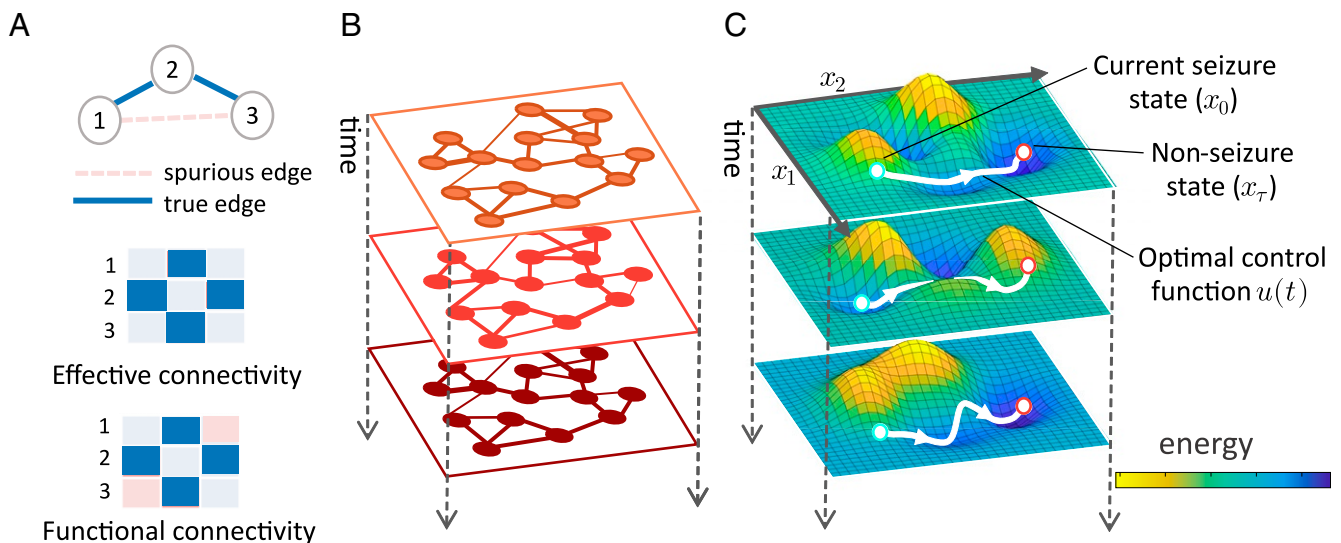


Fig. 1. Time-evolving EC as a tool in seizure control strategies. (A) In a toy network with two empirical functional connections across three nodes, correlative functional connectivity may estimate a spurious edge between nodes that interact indirectly. Conversely, EC reflects the direct influence between nodes in a network, setting the edge weight of conditionally independent nodes to zero, and thus provides a scaffold to model control strategies over time. Throughout a seizure, we can build EC networks to represent distinct regimes of neural connectivity (B) and infer a dynamic control energy landscape from those networks, shown superimposed on the principal axes of the system's state space (C). We can then use tools from control theory to determine the optimal control input required to drive neural activity to seizure freedom given the state of the control energy landscape.

in each regime given the stationarity of the time series (22, 32). We sought to determine whether such a regime separation existed in EC; if so, dynamic controllability could be meaningfully compared in seizure regimes. To address this question, we began by constructing EC networks from consecutive 1-s time windows of preictal and ictal ECoG recordings from 34 partial or secondarily generalized seizures across 14 subjects undergoing EEG monitoring prior to surgical treatment for medication-resistant epilepsy. EC networks were obtained using the Graphical Least Absolute Shrinkage and Selection Operator (GLASSO) method to estimate the regularized partial correlation between pairs of recorded time series in each time window as this method has been shown to estimate the presence of true network connections with high sensitivity compared to other EC methods (see *Materials and Methods* for more details) (11, 33–35). For each seizure of variable length T seconds recorded with a variable N channels, this procedure provided a total of T regularized partial correlation matrices of size $N \times N$ (see Fig. 2A and B and *Materials and Methods*).

For the purposes of our analysis, we next used a retrospective community detection method to detect seizure regimes; real-time seizure regime classification remains an interesting challenge that lies outside the scope of this study. We calculated what has previously been termed a functional connectivity dynamics matrix (36) or a configuration similarity matrix (30). This $T \times T$ similarity matrix contains ij th elements that represent the Pearson correlation coefficient between the upper triangles of the i th and j th EC networks. In addition, entries were inversely weighted by a constant β times the duration between time points to promote state contiguity (*Materials and Methods*).

Using a Louvain-like locally greedy community-detection algorithm to maximize a modularity quality function, we partitioned the columns of the similarity matrix thereby separating each seizure into three temporal regimes characterized by distinguishable patterns of EC (Fig. 2C). The largest temporally contiguous network assignments in each community were selected for subsequent analyses and were labeled chronologically as seizure regimes 1 through 3 with respect to the median time point of the community cluster (Fig. 2D). On average, $96 \pm 5\%$ of the EC networks from a given seizure were assigned to one of the three contiguous regimes, compared to only $30 \pm 12\%$ assigned during an equal-length preictal period where distinct regimes would not be expected (*SI Appendix, Fig. S2*). Thus, we found that our data-driven method could be used to demarcate onset, generalization, and termination regimes across seizures, as verified by a board certified epileptologist (BL), and that the pattern of region-to-region influences were sustained within each seizure regime.

Dynamic Controllability of Epileptic Networks. We now turn to an investigation of how the EC network representing each seizure regime impacts regime-dependent controllability. In the context of our model, analysis of the average and modal controllability metrics can answer the question of whether energy input during a given regime will be characterized by diffuse propagation throughout the network, pushing the brain state into other nearby states (high average controllability), or whether control input is able to drive the brain to states that are hard to reach by maintaining influence over the small modes of the system (high modal controllability). Additional analysis of persistent

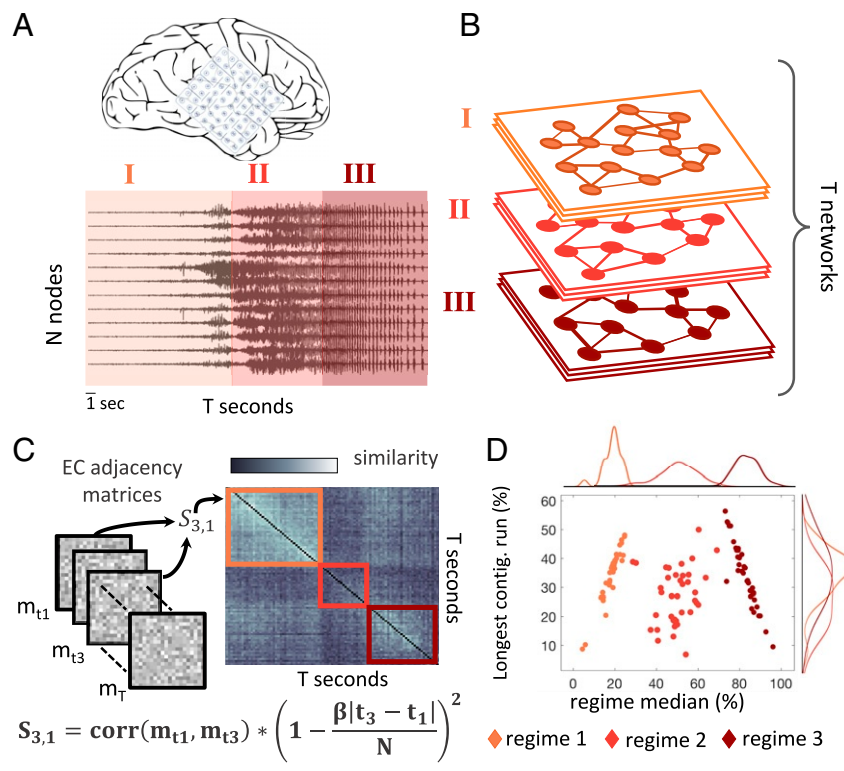


Fig. 2. Time evolving controllability through EC. (A) For each seizure, we extracted consecutive 1-s time windows of cortical ECoG recordings from N electrode channels. (B) From these data, we estimated T distinct EC networks. (C) We used community detection to determine three seizure regimes based on the similarity of the regularized partial correlation adjacency matrices and selected a single EC network to represent each regime for our optimal control analysis. (D) The distribution of the three largest regimes found in all 34 seizures. Regimes are organized chronologically and are plotted by their normalized temporal median versus the longest consecutive run of time windows, as a percentage of total seizure length. Each data point represents a regime in a single seizure.

modal and transient modal controllability metrics can describe the extent to which input energy into the brain network is either sustained while controlling slowly decaying modes or attenuated while controlling quickly decaying modes, respectively. We investigate metrics in each seizure regime under the assumption that the seizure progressed naturally (without prior stimulation) up to the seizure regime of interest.

For all four controllability metrics, we observed a significant regime dependence at the group level. We obtained 12 metric values for each seizure—four single values to summarize the average, modal, persistent modal, and transient modal controllability metrics in three regimes—and then compared the values in each regime across the 34 seizures. Within a regime, a given metric was calculated for each node in each EC network assigned to the regime. Then the median value for each node was obtained across networks, and the average of the resulting $N \times 1$ vector was used to summarize the regime. Across 34 seizures in 14 epilepsy patients, we observed a significant effect of seizure regime on average controllability ($\chi^2(2, 34) = 9.59, P < 0.01$), modal controllability ($\chi^2(2, 34) = 9.59, P < 0.01$), transient modal controllability ($\chi^2(2, 34) = 8.88, P < 0.016$), and persistent modal controllability ($\chi^2(2, 34) = 18.59, P < 1 \times 10^{-4}$) using Friedman's ANOVA across regimes. In post hoc analysis after Bonferroni correction for multiple comparisons of seizure regimes within each metric, we found a significant increase in average controllability from propagation to termination regimes ($t = -3.03, P < 0.01$), with a steady and significant decrease in modal controllability from onset to termination ($t = 3.03, P < 0.01$). Both transient and persistent modal controllability demonstrated a significant increase in value from onset to propagation ($t = -2.78, P < 0.016$, and $t = -3.15, P < 0.01$, respectively), and both decreased again from propagation to termination regimes ($t = 2.3, P < 0.064$, and $t = -4.12, P < 0.001$), although the decrease in transient controllability did not reach statistical significance (Fig. 3). Our results demonstrate that the extent of input energy propagation throughout the network will increase with the seizure evolution and that ease of directing activity toward hard-to-reach brain states is greatest at seizure onset.

Regime-Dependent Optimal Control Energy. The linear dynamical model used to understand state-invariant controllability in

brain networks can also reveal the optimal control energy necessary to drive the brain from an initial brain state x_0 to a final state x_τ over τ time steps (4, 7, 9, 10). Optimal control energy strives to minimize input energy while ensuring a direct trajectory from initial to final states, thus circumventing scenarios in which a controlled seizure might get worse before it gets better. We used the optimal control framework—after verifying that our EC regime model could sufficiently simulate brain state evolution (SI Appendix, Results)—to determine when during the seizure it would be most energetically favorable to drive the brain state from each seizure regime to a seizure-free baseline. Given the success of current RNS treatment practices where current is injected just after seizure onset is detected, we hypothesized that control energy would be smallest in the seizure onset regime, which is a period in which the seizure is often spatially confined (37).

We began by obtaining a representative EC network for each regime in each seizure as the network with the highest average similarity to other networks in the given regime. Per seizure, we quantified an $N \times 1$ initial brain state vector x_0 for the three seizure regimes by averaging the high- γ band power in 1-s time windows at each electrode across all time windows within the regime. We quantified a final, seizure-free brain state x_τ by similarly averaging band power across windows in a preictal time period equal in length to the seizure duration. Then, we measured how much input energy would be required to use each node to drive the brain from each of the three initial states, to the baseline brain activity level recorded in the preictal period. Finally, the nodal optimal control energy values were averaged across all nodes in a particular regime to arrive at three optimal control values for each seizure and regime. These values were then grouped by regime across seizures (Fig. 4A). Again Friedman's ANOVA was conducted to find that seizure regime had a significant effect on optimal control energy ($\chi^2(2, 34) = 18.41, P < 1 \times 10^{-4}$). Specifically, the seizure onset regime showed a significantly lower control energy requirement than the propagation regime at the group level ($N = 34, t = -4.24, P < 6.6 \times 10^{-5}$) and a lower energy than in the termination regime, although not significantly so ($N = 34, t = 2.67, P < 0.02$). The results demonstrate how a perspective of dynamic controllability can identify seizure regimes that may be controlled in an energetically favorable manner.

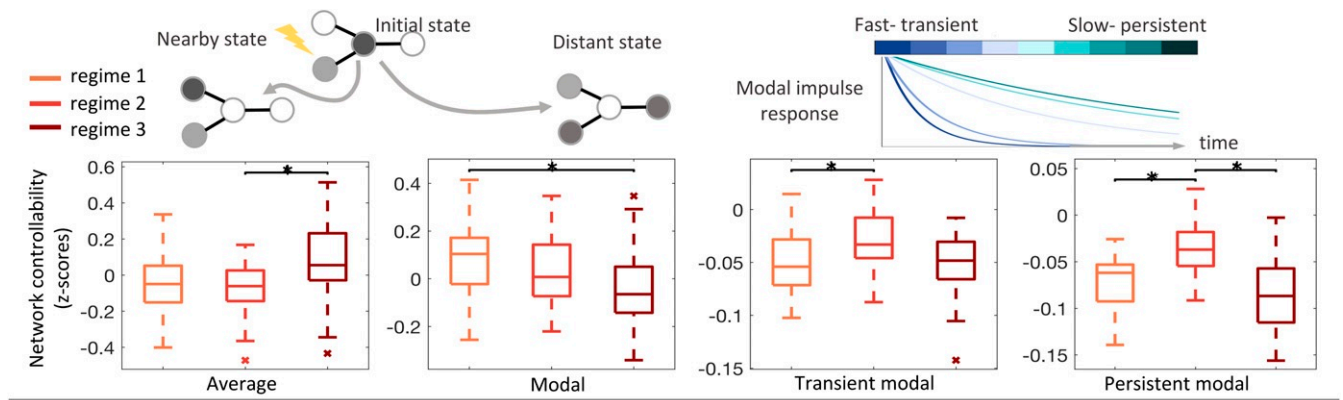


Fig. 3. Group-level controllability dynamics throughout early, middle, and late seizure regimes. A single controllability metric value was obtained for each regime in a seizure, and the distribution of regime metric values across all seizures ($N = 34$) is shown. Average controllability indicates the ease of driving network activity to nearby states and was found to increase throughout seizure regimes ($\chi^2(2, 34) = 9.59, P < 0.01$). Modal controllability indicates the ease of driving network activity to hard-to-reach or distant states and was found to decrease throughout seizure regimes ($\chi^2(2, 34) = 9.59, P < 0.01$). Transient and persistent modal controllability describe the ease of perturbing the slow, sustaining modes of the system or the fast, attenuating modes. A significant effect of seizure regime was found for both transient modal controllability ($\chi^2(2, 34) = 8.88, P < 0.016$) and persistent modal controllability values ($\chi^2(2, 34) = 18.59, P < 1 \times 10^{-4}$). Boxplots indicate the 75% confidence interval (box), median (solid line), 95% confidence interval (whiskers), and outliers (stars). Starred bars indicate significant metric differences between regimes at the $P < 0.017$ level, determined using Friedman's ANOVA.

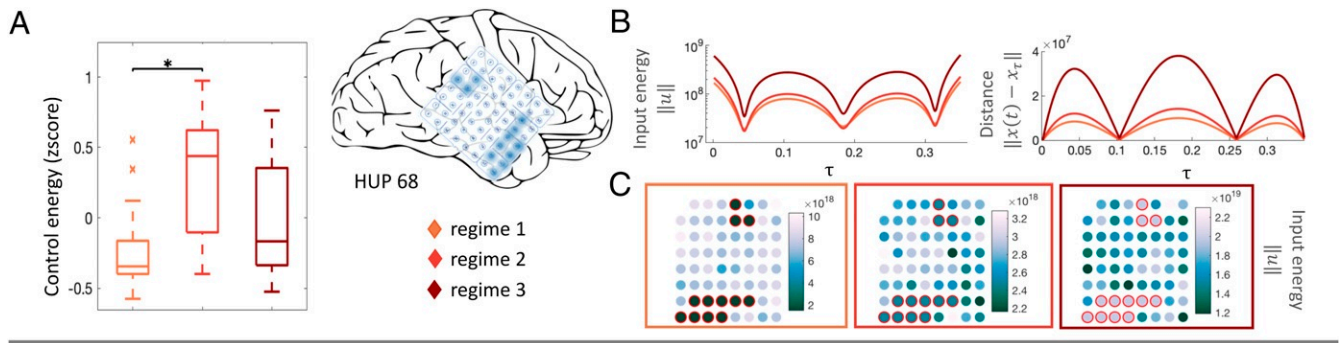


Fig. 4. State-dependent optimal control energy. (A) A significant effect of optimal control energy on seizure regime was found at the group level ($\chi^2(2, 34) = 18.41$, $P < 1 \times 10^{-4}$); energy values were significantly lower in the onset regime compared to propagation. Boxplots indicate the 75% confidence interval (box), median (solid line), 95% confidence interval (whiskers), and outliers (stars). Starred bars indicate significant metric differences between regimes using a t test at the $P < 0.017$ level. (B) The trajectory from seizure state to preictal baseline is controlled through designated SOZ nodes highlighted in blue on the cortical grid for a seizure in subject HUP68 over τ control time steps. (Left) Total input energy into the set of SOZ nodes for subject HUP68 across τ control time steps. (Right) The Euclidean distance of network state at time t to final target network state x_τ along the control trajectory (estimation error, 31.46). (C) Optimal energy values are displayed across the three regimes for a single seizure in subject HUP68, where SOZ energy was significantly lower than the null distribution in the onset regime ($N = 5000$, $P < 0.05$). Nodes are spatially arranged according to physical electrode placement, and the color of each node outside the SOZ reflects the optimal control energy value averaged across resampled control sets in which the node participated. Nodes in the SOZ are outlined in red.

Control Localization in Onset Regime. The optimal control energy results demonstrated that the seizure onset regime is preferred for its low energy requirements. In an effort to find the best control location, we next asked whether controlling nodes at the SOZ required less energy than controlling any other subset of nodes, given that RNS treatments are typically localized to the SOZ (1, 37). Seizures in all but four subjects exhibited focal onsets, for a total of 22 seizures localized within the area of the cortical grid electrodes and used in this analysis. We again measured the control energy required to drive the brain state along an optimal trajectory from the initial states x_0 in each seizure regime to a final state x_τ . This time, however, control input was supplied to each SOZ node simultaneously. We then repeated the simulation in 5,000 random reassignments of the SOZ labels among the non-SOZ electrodes. We found that in only half of the subjects exhibiting localizable SOZs, control using the SOZ nodes required significantly less energy than control from other subsets of nodes in the seizure onset regime ($N = 5,000$, $P < 0.05$) (Fig. 4B and SI Appendix, Fig. S9). We observed no correlation between SOZ significance and the number of SOZ nodes (Mann–Whitney $U = 95.5$, $P > 0.84$), suggesting that our results are not simply due to differences in SOZ size. Factors such as age at seizure onset, years with epilepsy, and seizure count were also not significantly different between groups (SI Appendix, Results). As an example, we show the spatial distribution of control energy throughout the 8×8 electrode grid array in subject HUP68 where the SOZ energy was significantly lower during seizure onset (Fig. 4C). Energy for non-SOZ nodes was calculated as the average energy across all permutation tests in which they participated. These results demonstrate that dynamic controllability analysis can answer spatiotemporal questions regarding stimulation for seizure suppression.

Robustness of Results to Selected Parameters. Our analysis relies on a number of assumptions, as well as parameterized methods for EC network generation and regime delineation, and a parameterized optimal control energy model. We sought to evaluate the reliability of our findings by assessing the sensitivity of our results to our assumptions and parameter values. First, given that spatiotemporal patterns of seizures propagating from the same SOZ may be diverse within a subject (38), we chose to treat each seizure as an independent data point. We completed an alternative analysis after grouping seizures within subjects and found

that the group-level trends of our main findings were preserved, with persistent and transient modal controllability maintaining significant differences across regimes despite the reduction in number of data points (SI Appendix, Fig. S3). We also varied the sparsity tuning hyperparameter, λ , used in EC estimation (SI Appendix, Fig. S5), and adjusted the emphasis on temporal contiguity of community assignments, β , when determining regimes (SI Appendix, Fig. S6). In these analyses we found no significant change in controllability values for any metric after varying each parameter in a neighborhood around the parameter value used in our main results. Finally, we reproduced our group-level optimal control energy results using three additional pairs of time horizon (τ) and distance–energy trade-off (α) parameters. We found that optimal control estimation error increased but that group-level trends were preserved for decreasing values of τ (SI Appendix, Fig. S8). Taken together, we find that our main results are robust to reasonable parameter variation.

Discussion

For patients with medication-resistant epilepsy, RNS as a means of seizure suppression is a relatively new option that provides an alternative to more permanent and invasive choices such as cortical resection or laser ablation. While the majority of patients improve over the course of months to years with treatment, the current trial and error routine for selecting stimulation parameters is far from ideal, and only rarely do patients achieve seizure freedom (1, 2, 37). We completed this study to introduce a control-theory framework for dynamic controllability that could provide support for clinical treatment decisions and highlight the regions and time points for stimulation that require the least input.

Emergence of Quasi-Static Seizure Regimes. In a challenge to a static brain controllability perspective, a number of studies find that multiple patterns of corticocortical coupling may appear on the same structural foundation and that time-evolving EC networks may provide a more accurate view of the connectivity landscape (14, 15, 35, 39). By estimating edges in our EC networks as regularized partial correlations between brain regions, we strive to remove contributions of common latent inputs from subcortical structures and describe only direct corticocortical influences. Using this framework as a linear approximation of the region-to-region dependencies in each seizure regime, we demonstrated

superiority to a static structural model (*SI Appendix, Results*) and uncovered a dynamic profile of controllability with implications for the input energy distribution across brain regions in a given regime. Our results demonstrate increasing average controllability and decreasing modal controllability throughout the majority of seizures. Low average controllability at onset implies that a widespread, diffuse distribution of input energy is hindered at seizure initiation, complementing findings in functional connectivity that show marked decoupling between brain regions at onset due to loss of inhibitory restraint (12, 31). Later throughout propagation and into termination regimes where coupling between brain regions has been shown to increase (12, 29, 31), our results suggest that it will become harder to direct the brain into energetically unfavorable states (low modal controllability).

Interestingly, the temporal trend of persistent and transient modal controllability differs from overall modal control in the onset regime. The relatively low metric values at onset reflect the combination of two factors—first, that input attenuation governed by the extreme eigenmodes could be relatively slower in general at seizure onset than during propagation, and second, that external perturbations may be less aligned with the extreme eigenmodes at onset. Our findings are consistent with a prior investigation of ictal eigenmode stability in a dynamical linear framework showing that relatively faster decay in the slowest eigenmodes occurred during midseizure, perhaps reflecting greater synchrony with compensatory inhibition during the propagation period (32). Our results indicate that decoupling between brain regions at onset could make it particularly difficult to control the very fast or very slow dynamics associated with the extreme eigenmodes of the brain's activity in that regime.

Spatiotemporal Control as a Tool for RNS Intervention. From the perspective of clinical RNS treatment, our model supports the current practice of intervening at seizure onset (37), with group-level results indicating that the amount of control energy required for a transition to a seizure-free state is smallest at onset. We note that increased energy requirements for propagation and termination regimes may be driven by our effort to counteract the natural ictal progression, by fully controlling brain state along a direct trajectory to seizure freedom, at an energetic cost. Our analysis on locally differentiated optimal control energy between SOZ and non-SOZ nodes found that the SOZ required significantly less energy for seizure suppression in seizures in only half of the subjects. This finding does not support the SOZ as a universally energetically favorable stimulation target and begs the question of whether energetically favorable points outside the SOZ could still be clinically relevant stimulation targets. There are a number of possible explanations and factors that drive this result. Epilepsy is theorized to be a network disorder (30, 40, 41); thus, it is likely that multiple cortical locations may be targeted to modulate the same pathologic network. Additionally, in resective surgery, removal of the SOZ alone is often insufficient, indicating that additional brain regions within the broader epileptogenic zone contribute to support seizures (28). These facts combined with evidence suggesting that seizure precursors may appear outside of the SOZ (40), that tissue damage at the SOZ may have remote effects on the functionality of other brain regions (42), and that the vicinity of stimulus location to the SOZ is not always significant for patient outcomes (43) indicate that energetically favorable control points selected by our model may still be effective even if outside the SOZ. Among other factors, our results may also be constrained by the linear control framework, which may capture shorter time horizons for control but could oversimplify the true trajectory back to a seizure free state (*SI Appendix, Results*).

Although current RNS technology is limited by a restricted number of seizure detection criteria and stimulation response patterns (2), our methodology could enrich a number of model-

based control approaches that could be adapted to improve seizure suppression in future implantable devices. For example, model predictive control (MPC) is a multivariable control algorithm that periodically reassesses the control function every few time steps along the planned trajectory (44). MPC relies on an internal representation of the dynamical environment, and our partial correlation EC model may improve upon prior estimates by reflecting direct corticocortical influences. Other recent work has theoretically demonstrated that seizures may be controlled via static output feedback of a linear dynamical system, wherein the underlying model is derived using multivariate autoregression (MVAR) (45). This work could be extended by removing spurious links in an MVAR model using a partial correlation constraint that assesses the correlation between a pair of nodes after removing the linear effects of other variables in the same and previous time points (46). Exciting future directions include testing the hypotheses generated from these models using intracranial data from patients with RNS implants and prospectively in new stimulation paradigms.

Methodological Considerations. Our model of dynamic controllability has a number of limitations. First, we chose to operationalize brain state here as high- γ band power. While this measure is commonly used in the epilepsy literature, it does not perfectly capture other relevant features of brain activity including the slope of the power spectral density, nonsinusoidal features of the LFP, etc. (47, 48). Additionally, our model uses only a single seizure-free brain state, when in reality it could be represented by many state vectors. Second, our model assumes noise-free, linear network dynamics within each time window and assumes that the corticocortical influences constraining brain state evolution during a given seizure regime will remain stationary with respect to the time horizon for optimal control. Although the brain is nonlinear and it is possible that the dynamics of the EC network will move away from its regime configuration after stimulation, linear approximations on a short timescale can capture broad dynamics in epilepsy while allowing for the application of linear control theory and ease of interpretability (11, 12). Extending this work to models of nonlinear control theory could similarly investigate ways to control the brain to a seizure free manifold, rather than a single state, and would complement the hypotheses generated here (18, 27). While beyond the scope of this work, clinical deployment of our method would necessitate real-time seizure regime classification, such as described in the algorithm validated by Baldassano et al. (49), and implementation details will depend heavily on the capabilities of the stimulation platform. Finally, while stimulation is occasionally used to induce, and then inhibit, a seizure while in the epilepsy monitoring unit, such occurrences are rare, and our dataset did not contain instances of stimulation. Immediate next steps will include testing the validity of our proposed method with data containing stimulation events at time points spanning all ictal regimes.

In conclusion, we employed a method of measuring controllability of effective networks and optimal seizure control energy across three temporal seizure stages. Our results present preliminary work in building theoretically informed hypotheses for future empirical validation.

Materials and Methods

Subject Data. The ECoG data used for this study were recorded from subdural electrodes implanted in 14 epilepsy patients undergoing EEG monitoring prior to surgical treatment for medically refractory epilepsy at the Hospital of the University of Pennsylvania (HUP) (512 Hz sample rate) or at the Mayo Clinic in Rochester, MN (500 Hz sampling rate). All subjects included in this study gave written informed consent. This study was approved by the University of Pennsylvania Institutional Review Board and Mayo Clinic Institutional Review Board. The implanted surface electrodes (ad Tech Medical Instruments) were composed of both grid and strip arrays with 2.3-mm-diameter contacts at 10 mm intercontact spacing. Electrode placement was

planned according to clinical protocol by a team of epileptologists with a separate reference electrode placed distant to the site of seizure onset, as in prior studies (30, 39). The following characterizations were determined by an epileptologist and reviewed in a clinical conference at HUP: seizure categorization as complex partial and/or secondarily generalized, the earliest electrographic onset, unequivocal electrographic onset, and focal SOZ or lack of seizure focus. The deidentified patient data were retrieved from the online International Epilepsy Electrophysiology Portal (IEEG Portal) (50).

EC Networks. For each epileptic event, EC networks were constructed from 1-s time windows of both ictal and preictal data. All time series data from ECoG channel recordings were referenced to the common average reference to avoid broad field effects (30, 35, 38). Each channel was digitally filtered with 60 Hz notch, 120 Hz low-pass, and 1 Hz high-pass filters using a fourth-order Butterworth design to remove line noise, drift, and high-frequency noise, respectively (38). The preprocessed data were partitioned into consecutive 1-s time windows spanning the length of the seizure and preictal period. For each window, the GLASSO method was used to find a regularized partial correlation EC matrix (33). See *SI Appendix, Methods*, for details.

Community Detection to Identify Seizure Regimes. A similarity matrix was computed using the Pearson correlation coefficient between all pairwise combinations of the T distinct EC networks for a single seizure; each element was linearly weighted by a contiguity parameter $\beta = 0.01$, such that networks at two distant time points were weighted less than those at adjacent time points. Using MATLAB, the Louvain algorithm from ref. 51 was used to assign networks to communities from the network similarity matrix. The number of communities discovered can be tuned by the scaling parameter γ . For each EC network, the Louvain algorithm was run 100 times using a value of γ that produced three communities, and a consensus partition was selected from the 100 trial partitions. See *SI Appendix, Methods*, for details.

Controllability Metrics and Optimal Control Calculations. Average, modal, persistent modal, and transient modal controllability metrics were calculated in MATLAB with custom scripts and code from ref. 7. Before calculating a given controllability metric, each function first ensures that the network is stable by enforcing Schur stability and then subtracting the identity matrix to normalize the eigenvalues (25). Optimal control energy was also calculated in MATLAB using functions referenced in ref. 7. Optimal control energy requires the selection of the value for τ , the number of time steps over which to perform the control trajectory optimization, and α , an energy-distance trade-off parameter. Parameter selection details can be found in *SI Appendix, Methods*.

Statistical Methods. Friedman's ANOVA, t tests, and Bonferroni corrections for multiple comparisons were performed using functions from MATLAB's statistics toolbox. We wrote custom MATLAB scripts to create null distributions for our energy analysis by randomly sampling with replacement

5,000 sets of nodes containing the same number of nodes as our set of interest. Our random resampling method is an extension of similar methods used to compute subject-specific confidence intervals for nodal metric values (52).

Diversity Statement. Recent work in several fields of science has identified a bias in citation practices such that papers from women and other minorities are undercited relative to the number of such papers in the field (53–57). Here we sought to proactively consider choosing references that reflect the diversity of the field in thought, form of contribution, gender, and other factors. We obtained predicted gender of the first and last author of each reference by using databases that store the probability of a name being carried by a woman (57, 58). By this measure (and excluding self-citations to the first and last authors of our current paper), our references contain 22.92% woman(first)/woman(last), 26.05% man/woman, 6.67% woman/man, and 44.36% man/man. This method is limited in that 1) names, pronouns, and social media profiles used to construct the databases may not, in every case, be indicative of gender identity and 2) it cannot account for intersex, non-binary, or transgender people. Second, we obtained predicted racial/ethnic category of the first and last author of each reference by databases that store the probability of a first and last name being carried by an author of color (59, 60). By this measure (and excluding self-citations), our references contain 9.29% author of color (first)/author of color(last), 12.25% white author/author of color, 23.54% author of color/white author, and 54.92% white author/white author. This method is limited in that 1) names, Census entries, and Wikipedia profiles used to make the predictions may not be indicative of racial/ethnic identity and 2) it cannot account for Indigenous and mixed-race authors or those who may face differential biases due to the ambiguous racialization or ethnicization of their names. We look forward to future work that could help us to better understand how to support equitable practices in science.

Data Availability. Anonymized iEEG were retrieved from the publicly accessible database hosted on the International Epilepsy Electrophysiology Portal (www.ieeg.org) (50).

ACKNOWLEDGMENTS. We thank both Jason Z. Kim for helpful discussions regarding control theory and Dr. Xiaosong He for helpful discussions regarding structural and EC model comparison. This work was primarily funded by NSF award BCS-1631550 (to D.S.B.) and National Institute of Neurological Disorders and Stroke award R01 NS099348 (to B.L. and D.S.B.). D.S.B. acknowledges additional support from the John D. and Catherine T. MacArthur Foundation, the Alfred P. Sloan Foundation, the Institute for Scientific Interchange Foundation, the Paul Allen Foundation, the Army Research Laboratory (W911NF-10-2-0022), and the Army Research Office (Bassett-W911NF-14-1-0679 and Grafton-W911NF-16-1-0474). B.L. acknowledges additional support from the Mirowski Family Foundation, the Pennsylvania Health Research Formula Fund, Johnathan and Bonnie Rothberg, and Neil and Barbara Smit. The content is solely the responsibility of the authors and does not necessarily represent the official views of any of the funding agencies.

- G. K. Bergey *et al.*, Long-term treatment with responsive brain stimulation in adults with refractory partial seizures. *Neurology* **84**, 810–817 (2015).
- A. Hartshorn, B. Jobst, Responsive brain stimulation in epilepsy. *Ther. Adv. Chronic Dis.* **9**, 135–142 (2018).
- F. Pasqualetti, S. Zampieri, F. Bullo, Controllability metrics, limitations and algorithms for complex networks. *IEEE Trans. Control Netw. Syst.* **1**, 40–52 (2014).
- E. Tang *et al.*, Control of brain network dynamics across diverse scales of space and time. *Phys. Rev.* **101**, 062301 (2020).
- S. F. Muldoon *et al.*, Stimulation based control of dynamic brain networks. *PLoS Comput. Biol.* **12**, e1005076 (2016).
- A. N. Khambhati *et al.*, Functional control of electrophysiological network architecture using direct neurostimulation in humans. *Network Neurosci.* **3**, 848–877 (2019).
- J. Stiso *et al.*, White matter network architecture guides direct electrical stimulation through optimal state transitions. *Cell Rep.* **28**, 2554–2566.e7 (2019).
- J. Z. Kim *et al.*, Role of graph architecture in controlling dynamical networks with applications to neural systems. *Nat. Phys.* **14**, 91–98 (2018).
- S. Gu *et al.*, Controllability of structural brain networks. *Nat. Commun.* **6**, 1–10 (2015).
- R. F. Betzel, S. Gu, J. D. Medaglia, F. Pasqualetti, D. S. Bassett, Optimally controlling the human connectome: The role of network topology. *Sci. Rep.* **6**, 1–14 (2016).
- S. M. Smith *et al.*, Network modelling methods for FMRI. *Neuroimage* **54**, 875–891 (2011).
- M. A. Kramer, S. S. Cash, Epilepsy as a disorder of cortical network organization. *Neuroscientist* **18**, 360–372 (2012).
- C. Kirst, M. Timme, D. Battaglia, Dynamic information routing in complex networks. *Nat. Commun.* **7**, 11061 (2016).
- A. Palmigiano, T. Geisel, F. Wolf, D. Battaglia, Flexible information routing by transient synchrony. *Nat. Neurosci.* **20**, 1014–1022 (2017).
- D. Battaglia, A. Witt, F. Wolf, T. Geisel, Dynamic effective connectivity of inter-areal brain circuits. *PLoS Comput. Biol.* **8**, e1002438 (2012).
- V. K. Jirsa, W. C. Stacey, P. P. Quilichini, A. I. Ivanov, C. Bernard, On the nature of seizure dynamics. *Brain* **137**, 2210–2230 (2014).
- A. Spiegler, E. C. A. Hansen, C. Bernard, A. R. McIntosh, V. K. Jirsa, Selective activation of resting-state networks following focal stimulation in a connectome-based network model of the Human brain. *eneuro* **3**, ENEURO.0068–16.2016 (2016).
- J. G. T. Zanudo, G. Yang, R. Albert, Structure-based control of complex networks with nonlinear dynamics. *Proc. Natl. Acad. Sci. U.S.A.* **114**, 7234–7239 (2017).
- B. Horwitz, The elusive concept of brain connectivity. *Neuroimage* **19**, 466–470 (2003).
- V. Sakkalis, Review of advanced techniques for the estimation of brain connectivity measured with EEG/MEG. *Comput. Biol. Med.* **41**, 1110–1117 (2011).
- Y. Ren *et al.*, Transient seizure onset network for localization of epileptogenic zone: Effective connectivity and graph theory-based analyses of ECoG data in temporal lobe epilepsy. *J. Neurol.* **266**, 844–859 (2019).
- P. J. Franaszczuk, G. K. Bergey, Application of the directed transfer function method to mesial and lateral onset temporal lobe seizures. *Brain Topogr.* **11**, 13–21 (1998).
- S. A. Weiss *et al.*, Ictal high frequency oscillations distinguish two types of seizure territories in humans. *Brain* **136**, 3796–3808 (2013).
- S. A. Weiss *et al.*, Seizure localization using ictal phase-locked high gamma: A retrospective surgical outcome study. *Neurology* **84**, 2320–2328 (2015).
- T. M. Karrer *et al.*, A practical guide to methodological considerations in the controllability of structural brain networks. *J. Neural. Eng.* **17**, 026031 (2020).

26. J. M. Shine *et al.*, Human cognition involves the dynamic integration of neural activity and neuromodulatory systems. *Nat. Neurosci.* **22**, 289–296 (2019).
27. E. Wu-Yan *et al.*, Benchmarking measures of network controllability on canonical graph models. *J. Nonlinear Sci.* **18**, 1–39 (2018).
28. L. Jehi, The epileptogenic zone: Concept and definition. *Epilepsy Current* **18**, 12–16 (2018).
29. A. N. Khambhati, K. A. Davis, T. H. Lucas, B. Litt, D. S. Bassett, Virtual cortical resection reveals push-pull network control preceding seizure evolution. *Neuron* **91**, 1170–1182 (2016).
30. A. N. Khambhati *et al.*, Dynamic network drivers of seizure generation, propagation and termination in human neocortical epilepsy. *PLoS Comput. Biol.* **11**, e1004608 (2015).
31. P. Jiruska *et al.*, Synchronization and desynchronization in epilepsy: Controversies and hypotheses. *J. Physiol.* **591**, 787–797 (2013).
32. T. Mullen, G. Worrell, S. Makeig, “Multivariate principal oscillation pattern analysis of ICA sources during seizure” in *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, I. Korhonen, Ed. (Institute of Electrical and Electronics Engineers, Piscataway, NJ, 2012), pp. 2921–2924.
33. J. Friedman, T. Hastie, R. Tibshirani, Sparse inverse covariance estimation with the graphical lasso. *Biostatistics* **9**, 432–441 (2008).
34. M. J. Rosa *et al.*, Sparse network-based models for patient classification using fMRI. *Neuroimage* **105**, 493–506 (2015).
35. A. Das, D. Sexton, C. Lainscek, S. S. Cash, T. J. Sejnowski, Characterizing brain connectivity from human electrocorticography recordings with unobserved inputs during epileptic seizures. *Neural Comput.* **31**, 1271–1326 (2019).
36. E. C. Hansen, D. Battaglia, A. Spiegler, G. Deco, V. K. Jirsa, Functional connectivity dynamics: Modeling the switching behavior of the resting state. *Neuroimage* **105**, 525–535 (2015).
37. M. J. Morrell, Responsive cortical stimulation for the treatment of medically intractable partial epilepsy. *Neurology* **77**, 1295–1304 (2011).
38. K. Ma *et al.*, Coalescence and fragmentation of cortical networks during focal seizures. *J. Neurosci.* **30**, 10076–10085 (2010).
39. P. Shah *et al.*, Characterizing the role of the structural connectome in seizure dynamics. *Brain* **142**, 1955–1972 (2019).
40. K. Lehnertz, H. Dickten, S. Porz, C. Helmstaedter, C. E. Elger, Predictability of uncontrollable multifocal seizures—Towards new treatment options. *Sci. Rep.* **6**, 24584 (2016).
41. S. S. Spencer, Neural networks in human epilepsy: Evidence of and implications for treatment. *Epilepsia* **43**, 219–227 (2002).
42. H. Ung *et al.*, Interictal epileptiform activity outside the seizure onset zone impacts cognition. *Brain* **140**, 2157–2168 (2017).
43. M. Seeck *et al.*, Electrode location and clinical outcome in hippocampal electrical stimulation for mesial temporal lobe epilepsy. *Seizure* **22**, 390–395 (2013).
44. M. Morari, C. E. Garcia, D. M. Pretz, Model predictive control: Theory and practice. *IFAC Proc. Vol.* **21**, 1–12 (1988).
45. A. Ashourvan *et al.*, Model-based design for seizure control by stimulation. *J. Neural Eng.* **17**, 026009 (2020).
46. B. Wild *et al.*, A graphical vector autoregressive modelling approach to the analysis of electronic diary data. *BMC Med. Res. Methodol.* **10**, 28 (2010).
47. M. A. Colombo *et al.*, The spectral exponent of the resting EEG indexes the presence of consciousness during unresponsiveness induced by propofol, xenon, and ketamine. *Neuroimage* **189**, 631–644 (2019).
48. N. Schaworonkow, V. V. Nikulin, Spatial neuronal synchronization and the waveform of oscillations: Implications for EEG and MEG. *PLoS Comput. Biol.* **15**, e1007055 (2019).
49. S. Baldassano *et al.*, A novel seizure detection algorithm informed by hidden Markov model event states. *J. Neural Eng.* **13**, 036011 (2016).
50. J. B. Wagenaar, B. H. Brinkmann, Z. Ives, G. A. Worrell, B. Litt, “A multimodal platform for cloud-based collaborative research” in *2013 6th International IEEE/EMBS Conference on Neural Engineering (NER)*, B. He, M. Akay, Eds. (Institute of Electrical and Electronics Engineers, Piscataway, NJ, 2013), pp. 1386–1389.
51. V. D. Blondel, J. L. Guillaume, R. Lambiotte, E. Lefebvre, Fast unfolding of communities in large networks. *J. Stat. Mech. Theor. Exp.* **2008**, P10008 (2008).
52. E. C. Conrad *et al.*, The sensitivity of network statistics to incomplete electrode sampling on intracranial EEG. *Network Neurosci.* **4**, 484–506 (2020).
53. S. M. Mitchell, S. Lange, H. Brus, Gendered citation patterns in international relations journals. *Int. Stud. Perspect.* **14**, 485–492 (2013).
54. D. Maliniak, R. Powers, B. F. Walter, The gender citation gap in international relations. *Int. Organ.* **67**, 889–922 (2013).
55. N. Caplar, S. Tacchella, S. Birrer, Quantitative evaluation of gender bias in astronomical publications from citation counts. *Nat. Astron.* **1**, 0141 (2017).
56. M. L. Dion, J. L. Sumner, S. M. Mitchell, Gendered citation patterns across political science and social science methodology fields. *Polit. Anal.* **26**, 312–327 (2018).
57. J. D. Dworkin *et al.*, The extent and drivers of gender imbalance in neuroscience reference lists. *Nat. Neurosci.* **23**, 918–926 (2020).
58. D. Zhou *et al.*, “Gender diversity statement and code notebook v1.0. Zenodo. 10.5281/zenodo.4104748. Deposited 18 October 2020.
59. A. Ambekar, C. Ward, J. Mohammed, S. Male, S. Skiena, “Name-ethnicity classification from open sources” in *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge discovery and Data Mining - KDD '09* (Association for Computing Machinery Press, New York, 2009), p. 49.
60. G. Sood, S. Laohaprapanon, Predicting race and ethnicity from the sequence of characters in a name. arXiv:1805.02109 (5 May 2018).