Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

A Network-Control Approach to Seizure Intervention via Adjoint Optimal Control

Aadi Deosthaler
Bridgeland High School, Cypress TX
Samarth Muralidhara
samarthmurali19@gmail.com
Bridgeland High School, Cypress TX
Yashas Anil
Bridgeland High School, Cypress TX

Research Article

Keywords: Epilepsy, Network Control, Adjoint Optimal Control, Nonlinear Dynamics, Optimal Control

Posted Date: June 30th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10145295/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

A Network-Control Approach to Seizure Intervention via Adjoint Optimal Control ⋆,⋆⋆

Samarth Muralidhara [a] Aadi Deosthale [a] Yashas Anil [a]

[a] *Bridgeland High School, Cypress, Texas 77433*

Abstract

Most existing open-loop Deep Brain Stimulation (DBS) systems for drug-resistant epilepsy do not take into account the temporal dynamics of seizures or the spread of seizure activity throughout the brain. In this paper, we propose a closed-loop control system that models seizure prevention as a constraint on the nonlinear dynamics of the brain's structural network. The structural network of the brain is modeled using the HCP-1065 connectome, which divides the brain into 360 regions according to the Glasser atlas. While the connectome was built using the full 360-region atlas of the brain, a densely connected sub-network of 60 brain regions was used to develop and test the proposed control system. Each region of the brain is modeled as a nonlinear supercritical Hopf oscillator. A detector continuously monitors certain features of the dynamics of each region of the brain, such as variance and lag-1 autocorrelation, to identify the onset of seizure activity. Two levels of control are applied: a linear-quadratic regulator (Level 1) that maintains the Hopf oscillators' basal state, and a dormant network-wide optimal controller (Level 2) that employs an adjoint-based optimal-control method (Adjoint Optimal Control; AOC) applied directly to the Pontryagin Hamiltonian to provide pre-emptive control should the system enter a region of unsafe operation. The predictive detector displays high accuracy in detecting seizure transitions, achieving an AUC of 0.889 (95% CI $[0.820, 0.950]$) on an independent test fold containing 50 trajectories per class, and demonstrating a mean lead time to the onset of the seizure of 56.9 seconds. Within the modeled regime of neuronal excitability, the model is able to prevent the development of seizure-like bifurcations in 100% of 100 Monte Carlo simulations (95% CI: 96.3–100%).

Key words: Epilepsy; Network Control; Adjoint Optimal Control; Nonlinear Dynamics; Optimal Control.

1 Introduction

Drug-resistant epilepsy is difficult to treat due to the complex dynamics of the epileptic brain and the inability to predict when a seizure will occur [1,5]. Epilepsy is experienced by 50 million individuals across the planet [9,38]. Around 30% of all epileptic individuals experience seizures that are not controlled with anti-epileptic drugs; this percentage has been cited in studies since at least 2000 and continues to be replicated in numerous scientific publications [16]. In cases where drugs are tried and determined to be unable to provide control over the seizures, additional treatments may be utilized. One treatment that is often considered invasive is a resective surgery to the regions of the brain that are exhibiting these seizures [37]. Should those procedures not provide the necessary reduction in seizures for the affected patient, other procedures may be implanted into the patient, such as responsive neurostimulation (RNS) systems [6,19] or deep brain stimulation (DBS) devices [23].

Most patients receiving DBS are managed with open-loop systems that stimulate continuously or according to a fixed schedule, without accounting for the time-varying risk of seizure [23,24]. In this article, we develop and simulate a closed-loop control framework as a computational alternative to these conventional open-loop implementations.

In parallel, epilepsy research is shifting away from investigating the role of individual regions of the brain toward understanding how seizures can propagate throughout the brain [1]. Techniques like diffusion-weighted neuroimaging have allowed for the reconstruction of the structural connections between different regions of the cerebral cortex. Consequently, the question of epilepsy is shifting from where the seizures originate within the

⋆ Research conducted at Bridgeland High School, Cypress, Texas 77433.

⋆⋆These authors contributed equally to this work and share first authorship. Corresponding author S. Muralidhara.

Email addresses: samarthmurali19@gmail.com (Samarth Muralidhara), aadideosthale2021@gmail.com (Aadi Deosthale), yoshiyashasganil@gmail.com (Yashas Anil).

brain toward the structure of the brain that permits the seizures to spread, requiring a model of the brain as a network. Most methods for determining epileptogenic areas or selecting stimulation targets for anti-seizure therapies, however, are based upon linear control theory. The dynamics of seizures are inherently nonlinear and multistable, and linear control theory is a poor framework for understanding or addressing these types of dynamics. For instance, linear control theory often overestimates the control energy required for seizures that are related by a bifurcation and cannot distinguish between the qualitative differences between normal and seizure activity. Furthermore, linear control theory often intends for the system to converge toward a certain goal state, whereas seizure prevention is a problem of keeping the brain away from a certain boundary state.

In light of this, we propose a control approach that prioritizes the safety of brain activity. Control Barrier Functions (CBFs) [25] provide a theoretically rigorous framework for safety-critical control by defining a safe set of states and enforcing its forward invariance. Although the same philosophy of defining a safe set for the system's state and ensuring its invariance under the system's dynamics is employed in this work, the method employed to achieve this invariance is based on an L_2-regularized adjoint-based optimal-control approach applied to the Pontryagin Hamiltonian, subject to a finite actuator saturation bound ($u_{\max} = 0.5$), rather than through an explicit barrier quadratic program. The full CBF formulation is left to future work. This method will only intervene in the system in the event of a need to satisfy the safety constraint, thus minimizing stimulation of the brain while keeping energy costs low.

Accordingly, this paper pursues three main goals: first, to formulate seizure prevention as a constraint upon the dynamics of the connectome; second, to develop a model that incorporates the connectome and optimal control theory to deliver stimulation satisfying that constraint; and third, to evaluate the framework against open-loop DBS, a connectome-naive targeting baseline, and an adaptive PI controller from the literature. Although the connectome model is built on the full 360-region Glasser parcellation [18], the control framework is developed and validated on a densely connected 60-node sub-network as a computational proof-of-concept.

The remainder of the paper is organized as follows: Section 2 reviews the background on connectome models, phenomenological seizure dynamics, and safety-critical control; Section 3 describes the model, control theory, and system architecture; Section 4 presents simulation results; Section 5 discusses limitations and future directions; and Section 6 concludes.

2 Background

2.1 Structural Connectomics and Network Models of Epilepsy

The concept of epilepsy has developed from the notion of focal pathological points to a network perspective in which dynamical patterns of seizure activity are interpreted in terms of large-scale structural connectivity [1]. The brain can be represented as a graph in which each node represents a region of the brain, and the edges represent the white-matter fibers that connect those regions, with inter-regional connectivity quantified by Diffusion Weighted Imaging (DWI) and Diffusion Tensor Imaging (DTI) scans [7]. These imaging techniques measure the anisotropic diffusion of water molecules to estimate white-matter integrity, from which an adjacency matrix A is constructed. Recent studies based on these techniques have revealed abnormalities in the connectivity patterns of patients with Temporal Lobe Epilepsy that match known seizure propagation routes [1]. Moreover, the multimodal Glasser cortical atlas, which divides the brain into 360 different regions, provides sufficient resolution to reveal the complex dynamical patterns associated with seizure propagation throughout the brain [18]. The structural connectivity of the human brain has been shown to possess small-world properties, with a high clustering coefficient and short characteristic path length between regions [32,33,4]. Studies have shown that the small-world nature of the structural connectivity of the brain facilitates rapid transitions from synchronized activity to seizure generalization, and disruptions of these properties have been observed in structural connectivity matrices derived from epileptic brains [4].

2.2 Phenomenological Dynamical Models of Seizure Onset

A trade-off between detail and tractability must be made in creating computational models of seizure dynamics. Phenomenological models use mean-field approximations to describe the dynamics of neuronal populations, avoiding the need to specify every biophysical detail [20]. Among these, the supercritical Hopf bifurcation model is widely used to describe the dynamics of focal seizures [20]. The Hopf model posits that seizure dynamics can be described as a dynamical system with a stable fixed point (representing the interictal state) transitioning into limit-cycle oscillation (representing ictal activity). Here $z \in \mathbb{C}$ represents the reduced neural population state (where $|z|$ is the oscillation amplitude and $\arg(z)$ is the phase), r is the bifurcation parameter representing the excitation–inhibition balance, and ω is the intrinsic oscillation frequency of the population. This transition depends on r: for $r < 0$ the system is stable at the fixed point; for $r > 0$ it bifurcates into limit-cycle dynamics representing seizure activity. The

supercritical Hopf normal form is:

$$\dot{z} = (r + i\omega)z - |z|^2 z. \quad (1)$$

Bifurcation analyses have categorized seizure activity according to the mechanisms underlying onset and termination, including supercritical or subcritical Hopf bifurcations and saddle-node on invariant circle (SNIC) bifurcations. Corticothalamic models have shown how physiological properties interact to determine critical dynamics underlying absence or tonic-clonic seizure activity for different bifurcation types [10]. Studies based on neural mass models have confirmed that complex collective oscillations in heterogeneous populations emerge via supercritical Hopf bifurcations [20,30].

2.3 Safety-Critical Control: From Control Barrier Functions to Adjoint Optimization

Due to the rapid excitation of neurons during a seizure, the focus of this paper is upon constraining the bursts of neuronal activity that occur during an ictal transition. We draw upon the theory of Control Barrier Functions (CBFs) [25] and set invariance in the creation of a control law that ensures the system remains within a safe regime for the neural system. A CBF is a scalar state function $h(x)$ that defines the safe regime for a system: the safe set is $\mathcal{S} = \{x : h(x) \geq 0\}$, and forward invariance of $\mathcal{S}$ is enforced by the condition $\dot{h}(x) \geq 0$ on the boundary of $\mathcal{S}$. In a pure CBF quadratic program, the controller restricts inputs based on this barrier condition, providing a hard mathematical safety guarantee.

In the present work, we adopt this safe-set, set-invariance philosophy as our guiding principle, but approximate the safety constraint through an adjoint-based optimal-control method (AOC) rather than through an explicit quadratic program. Our optimization layer penalizes the system as it approaches the unsafe boundary, driving the state back into the safe region. Rather than strictly preventing the seizure via a hard mathematical barrier, the AOC controller prevents the system from drifting from the safe set toward the unsafe set. This approach does not drive the brain to a single equilibrium, but instead enforces a safety constraint such that the state naturally avoids the unsafe region. The full CBF formulation is left to future work.

2.4 Limitations of Linear Network Control in Epilepsy

Current approaches to network control theory assume linearity in the brain [31,21]. For example, the controllability Gramian is used to calculate the energy required for the brain to move from one state to another state [31]. The assumption of linearity in the brain is limiting in the context of epilepsy, because the dynamics of seizure activity are nonlinear. The brain can reside in either a stable physiological state or a stable seizure state, with transitions between these regimes governed by nonlinear mechanisms that linear models fail to capture. The linearization of these processes fails to account for the separation between basins of attraction, nor does it represent the limit-cycle dynamics exhibited by epileptic brains. Thus, the controllability metrics calculated for these linearized models often misrepresent the true effort required to suppress or redirect seizure activity. Furthermore, the energy calculations made for linear control theory substantially underestimate the control energy that would actually be required relative to the brain's known metabolic budget, limiting their relevance for neural systems [1]. Nonlinear control theory can exploit the energy dynamics of the epileptic brain to formulate control strategies that require less energy to be effective.

2.5 Closed-Loop Neuromodulation and Optimization Frameworks

To address the shortcomings of open-loop systems, closed-loop systems offer real-time optimization of neural states. The Responsive Neurostimulation (RNS) device is an established FDA-approved closed-loop treatment for drug-resistant epilepsy [6,19], demonstrating that closed-loop systems can be effectively used in the management of epilepsy. Although current RNS implementations use fixed detection thresholds and predefined stimulation patterns, next-generation models use real-time optimization based on quadratic programming (QP). Through QP, a control input can be determined that minimizes an objective function while ensuring that safety constraints are satisfied. Research into closed-loop control policies based on Bayesian and QP models has demonstrated their effectiveness in minimizing energy expenditure for neural population models while maintaining control over those populations [8,11].

3 Methods

3.1 Modeling Framework Overview

Our framework is a computational pipeline that extends from the structural connectome to a closed-loop neural controller. The structural connectome is taken from the HCP-1065 dataset [17], and each brain region is represented as a supercritical Hopf oscillator [3,20]. These oscillators are connected according to the structural connectome to form a network model of the brain. We first verify numerically that the resting network is stable and that a perturbation at a focal point propagates throughout the network to produce a seizure. An independent detector monitors the network for signs of an approaching bifurcation and outputs a risk score indicating the likelihood of an impending seizure. Control is performed through a two-level architecture: a continuous local linear-quadratic regulator (Level 1)

that maintains regional homeostasis [21], and a dormant network-wide optimal controller (Level 2) based on the Pontryagin control Hamiltonian [22] and solved by adjoint-based gradient descent (AOC), which is triggered only when the detector flags an unsafe state. The complete system is evaluated against three baselines.

All structural analyses use the full 360-region Glasser atlas [18]; however, all control simulations and Monte Carlo evaluations were conducted exclusively on a 60-node sub-network due to computational constraints (see Section 5 for the implications of this choice). All dynamical variables within the model are dimensionless: because the supercritical Hopf normal form is inherently non-dimensional, all variables—such as the excitability parameter, oscillation amplitude, control variable, and time—are defined in arbitrary units rather than physical units. The implications of this for clinical translation are discussed in Section 5.4. A discussion of how this affects the saturation of stimulation is included in Section 3.7.2.

3.2 Structural Network Construction

The structural connectivity matrix was obtained from the HCP-1065 dataset [17], which is a diffusion-MRI template constructed from 1,065 healthy adults using probabilistic tractography. The edges between regions of interest are determined by the number of streamlines that passed between the regions, and the brain was divided into $N = 360$ regions belonging to the multimodal Glasser atlas [18]. The resulting adjacency matrix $A \in \mathbb{R}^{360\times360}$ contains a value A_{jk} representing the streamline count between region j and region k, used to construct A, which is subsequently normalized. The HCP-1065 dataset already has a symmetric connectivity matrix, so the graph is considered to be undirected. Furthermore, self-connections (diagonal entries) were removed from the graph. No threshold was applied to the connectivity density values because the graph is required to contain all of the connections between regions of interest, and the spectral normalization below bounds the global coupling regardless. For the global coupling to be bounded and for the matrix to be numerically stable, the matrix was normalized by its largest eigenvalue,

$$A_{\text{norm}} = A \,/\, \lambda_{\max}, \tag{2}$$

ensuring that the spectral radius of the coupling matrix was unity (Figure 1). The network used for development and validation was a reduced sub-network containing $N = 60$ nodes within the parcellation, corresponding to the most densely connected module of the full network.

3.3 Node and Network Dynamics

Each region was modeled as a supercritical Hopf oscillator according to the normal form in (1), which we adopt because it is the simplest phenomenological model that

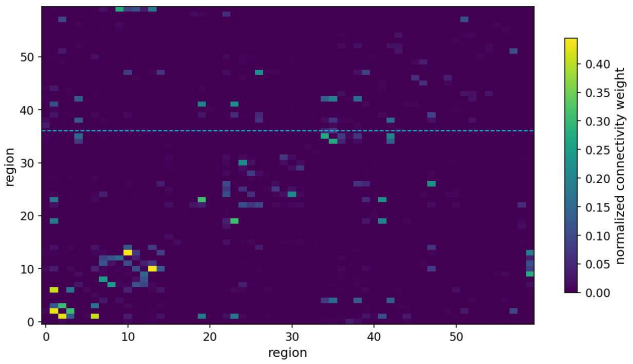


Fig. 1. Normalized structural connectome adjacency matrix A_{norm} on the full 360-node Glasser atlas. Color axis indicates normalized connectivity weight (normalized streamline density).

captures the transition from a stable interictal resting state ($r_j < 0$) to sustained seizure oscillations on a limit cycle ($r_j > 0$) as the excitability parameter crosses the critical threshold $r_j = 0$. The dynamical equation for each region j is:

$$\dot{z}_j = (r_j + i\omega_j)z_j - |z_j|^2 z_j. \tag{3}$$

The excitability value of each region was assigned to the bifurcation parameter r_j. Each region was initialized to its healthy regime, with $r_j = -0.02$ for healthy regions and $r_j = 0.6$ for the single seizure focus. It is assumed that each region within the cortex has the same intrinsic firing frequency; in reality, the frequencies for different regions of the cortex are known to be somewhat heterogeneous. For the model, each region is assumed to have the same intrinsic frequency of $\omega_j = \omega = 2\pi \times 0.10$ rad s^{-1}. This parameter defines a rescaled, dimensionless timescale relative to physiological timescales; the model is currently not sensitive to different frequencies within the cortex, though this is one potential area for future research efforts. The coupling between regions is effected through the normalized connectome matrix A, with a scalar global coupling value of G, and through the stimulation matrix B, leading to the network model:

$$\dot{z}_j = (r_j + i\omega_j)z_j - |z_j|^2 z_j + G\sum_k A_{jk}(z_k - z_j) + B_j u_j. \tag{4}$$

The coupling between identical states is diffusive in nature. The resting state of the network is a fixed point of the coupled system. In each simulation, the global coupling was set to $G = 2.5$.

While each region of the network can be individually actuated, physical DBS systems are able to inject only scalar, real-valued currents. The current model of the control variable u_j perturbs both the real and imaginary components of the complex variable z_j, which represents the theoretical abstraction of the DBS lead sys-

tem. Should this model be applied to a physical lead system with only a single electrode, however, the B_j matrix would have to be constrained to perturb only the real component of z_j. The cost and adjoint gradients are calculated with respect to the complex-valued input variable, however, which may lead to optimization algorithms that prescribe a perturbation that has a non-zero imaginary component; such a component of the control signal does not have a physical counterpart. This limitation of the model to a single physical electrode is discussed in Section 5.4.

The system was integrated using the fourth-order Runge–Kutta (RK4) method with a time step of 0.05 s for the duration of the simulation. The background activity was modeled as an additive Wiener process with amplitude of 0.02 [28]. This noise level was intentionally chosen to be small relative to the radius of the limit cycle of the seizing focus ($\sqrt{r_{\text{focus}}} \approx 0.775$) [20]. The noise was added in the Euler–Maruyama style [28] after the calculation of the RK4 update. As the diffusion coefficient is additive and does not depend on the state variable, the Milstein correction term is zero. Furthermore, the operator-splitting approach—applying RK4 to the deterministic drift and the Euler–Maruyama increment to the diffusion—is valid when the noise amplitude is small relative to the deterministic drift, which was verified throughout the simulations. A formal study of the convergence of the RK4 integration method is presented in Section 5.4.

3.4 Uncontrolled Seizure Propagation

The brain is modeled as a weighted undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$, where $\mathcal{V}$ is the set of $N = 360$ cortical nodes, $\mathcal{E}$ is the set of edges, and $W \in \mathbb{R}^{N\times N}$ is the adjacency matrix of streamline counts [17,7]. The hippocampus, which is located in the mesial temporal lobe, was selected as the focus of the seizures due to its known role in triggering seizure activity in patients with TLE, as well as the high eigenvector centrality of its most central node within the mesial temporal lobe.

Let v be the principal eigenvector associated with the dominant eigenvalue λ_1 of W [26], satisfying:

$$Wv = \lambda_1 v. \tag{5}$$

For a strongly connected graph with positive weights, λ_1 is real and positive, and v is a positive, nonzero vector (Perron–Frobenius theorem). The centrality of node i satisfies:

$$v_i = (1/\lambda_1) \sum_j W_{ij} v_j. \tag{6}$$

By setting $j_{\text{foc}} = \arg\max_{i\in\text{Hippocampus}} v_i$, we choose the initial bifurcation to occur at the node within the hippocampus that has the highest structural influence within the mesial temporal lobe. This choice is based on the assumption that perturbations of high-centrality nodes will spread most rapidly throughout the network's modes of macroscopic activity. Thus, although the most central node in the entire network may have been chosen as a candidate for bifurcation, it is instead this most central node within the hippocampus that is chosen for its anatomical centrality.

To derive the normal form of the dynamics of such a node of interest, let $\dot{x} = F(x, \mu_c)$ describe the dynamics of a cluster of neurons of interest, where x is a vector of the membrane potential values for the neurons in the cluster, and μ_c is the external GABAergic input that acts upon the cluster (the bifurcation parameter, distinct from the control input u defined in Section 3.7). Let the fixed point be $x_0 = 0$. The Jacobian at the critical parameter value is:

$$J = D_x F(0, \mu_c), \quad \lambda_{1,2} = \pm i\omega_c, \tag{7}$$

where all other eigenvalues satisfy $\text{Re}(\lambda_j) < 0$ for $j \geq 3$. The state space is decomposed into a 2-dimensional center eigenspace E_c and an $(n-2)$-dimensional stable eigenspace E_s. Applying a near-identity transformation of the form $x = y + \Phi(y)$ projects the dynamics onto the 2-D center manifold W_c. Variables in the stable manifold E_s decay exponentially to zero and are slaved to the slower varying variables in E_c. Applying a third-order Taylor expansion and transforming to complex coordinates $z = y_1 + iy_2$, the non-resonant nonlinear terms annihilate due to rotational symmetry, leaving the supercritical Hopf normal form:

$$\dot{z} = (\alpha(\mu_c) + i\omega_c)z + \beta_{21}|z|^2 z + \mathcal{O}(|z|^5), \tag{8}$$

where the truncation at $\mathcal{O}(|z|^5)$ is valid when $|z|$ is small; at higher excitability, higher-order terms may become non-negligible. Assuming that the first Lyapunov coefficient $\text{Re}(\beta_{21}) < 0$ (supercritical bifurcation), we can perform a change of variables to yield the isolated nodal equation used in the network. This center manifold reduction applies strictly to the regional population models prior to network embedding. The resulting phenomenological network is 120-dimensional for the $N = 60$ validation sub-network, and global network bifurcations may possess higher-dimensional structure not captured by the isolated-node reduction:

$$\dot{z}_j = (r_j(t) + i\omega_j)z_j - |z_j|^2 z_j. \tag{9}$$

Background noise transforms the deterministic ODE into a high-dimensional Itô SDE [28]:

$$dz = F(z, r)\, dt + \Sigma\, dW_t, \tag{10}$$

where W_t is a $2N$-dimensional Wiener process and $\Sigma = \xi I_{2N}$ is a scalar multiple of the identity (additive, state-independent noise). The system is integrated numeri-

cally using the Euler–Maruyama scheme; the theoretical justification for why critical slowing down (CSD) occurs before the seizure relies on the continuous evolution of the probability density function $P(z,t)$ over $2N$-dimensional state space, governed by the Fokker–Planck equation [29]:

$$\frac{\partial P}{\partial t} = -\sum_i \frac{\partial}{\partial z_i}\big[\mu_i(z)P\big] + \frac{1}{2}\sum_{i,j}\frac{\partial^2}{\partial z_i \partial z_j}\big[D_{ij}(z)P\big], \tag{11}$$

where $\mu_i(z)$ is the drift vector (the i-th component of $F(z,r)$) and $D_{ij}(z)$ is the diffusion tensor. When the focus parameter r_j approaches zero from the negative direction, the restoring gradient of the drift vector μ_i approaches zero along the radial axis. The influence of the diffusion coefficient D then dominates the limiting process for the stationary distribution of $P(z)$. As a result, the variance of the stationary distribution tends to infinity. Thus, the rolling variance of the state variables within the unsafe-state detector can be tracked.

Furthermore, to visualize the propagation heatmap, we order the dynamics based on their structural distance from the target network. The network state is obtained from the nonlinear coupled-Hopf dynamics in Section 3.3; the graph Laplacian $L = D - W$ is only used to compare the nonlinear dynamics with the linear graph diffusion process, not to generate the network state for the coupled-Hopf system [26]. A purely linear diffusion process on the same graph would evolve as:

$$\frac{dx_{\text{lin}}}{dt} = -L\,x_{\text{lin}}, \tag{12}$$

with spatiotemporal kernel e^{-Lt} spreading energy along the minimal paths of L. For the heatmap, the shortest structural path from j_{foc} to all other nodes is obtained via Dijkstra's algorithm, where each edge has a weight of $1/W_{uv}$ (excluding edges with zero weight as they indicate no possible connection between the two nodes), and the y-axis is sorted according to the vector of structural distances (d_s) calculated for each node.

To establish connectome dependence, we used the Maslov–Sneppen rewiring algorithm [27], which uses MCMC methods to sample degree-preserving surrogate networks. Each surrogate was generated by performing $10 \times |\mathcal{E}|$ double-edge swap steps, where $|\mathcal{E}|$ is the number of edges in the network. While node degrees are preserved, the specific connections and higher-order path structure are not, allowing us to isolate the contribution of brain topology. A total of 20 degree-preserving surrogate networks were generated, and recruitment on the true connectome was compared against the surrogate ensemble using a one-sample t-test. We also ablated the onset zone by setting all incoming and outgoing connection weights of the focus node to zero ($W_{j_{\text{foc}},k} = W_{k,j_{\text{foc}}} = 0$), and measured the resulting change in mean non-zone amplitude.

3.5 Predictive Detector (Critical Slowing Down)

The detector operates online by calculating two features of the time-series data within a sliding window of length 200 samples and a step of 1 sample. Within each window, the variance and lag-1 autocorrelation of the signal are calculated directly. These two features are min-max normalized over the length of the record and combined into a single risk score in the range $[0, 1]$ via a weighted sum of the two features with equal weights of 0.5 assigned to each feature [12,13]. These equal weights were assigned to each feature as a modeling choice; optimization of the weights of each term was not performed. Ablation studies of different weighting factors for the two features are outside the scope of this work and represented in future work. The threshold θ for triggering a warning of the approach towards an unsafe state is $\theta = 0.250$, which was determined on a separate fold of the data as the threshold that would yield a classifier with the highest recall (see Section 3.6 for the detailed procedure for determining the threshold). Furthermore, that threshold was not tuned to maximize the detection of bifurcations within the evaluation data. When the risk score calculated within a sliding window crosses the threshold $\theta = 0.250$, an alert is triggered for that region of the signal. However, no controller is implemented; the features of the signal approach the bifurcation in the results (Figure 5) [12,13]. Thus, the critical slowing down of the system is reflected in the vanishing gradient of the restoring force of the supercritical Hopf bifurcation; however, the use of such indicators of an approaching bifurcation as indicators of the occurrence of human seizures is contested in the literature [14,15]. Consequently, such an approach is limited in its reliability; a discussion of this limitation is provided in Section 5.4.

3.6 Independent Detector Validation

The detector was characterized on an independent held-out set of 50 test trajectories per class. Threshold θ was selected as 0.250 on a separate fold of 50 trajectories per class; the threshold was the highest threshold that achieved full recall on the validation fold. The test fold was never examined in the threshold selection process, ensuring that there is no threshold leakage. Furthermore, while 50 trajectories per class represents a modest number of trajectories, the confidence intervals associated with these results are inherently uncertain and not reflected within the presented analysis; collecting additional trajectories for training and testing represents a future direction for this work. Additionally, the detector was trained and tested on the same data upon which the model was trained; it has not been validated outside of

its distribution, or upon real EEG data; this is discussed in Section 5.4. Within each trajectory, the parameter r_j is ramped linearly from its resting value to its seizing value in the presence of noise over 2400 time steps. The true time that r_j crosses from negative to positive is the ground truth seizure onset time t^*. Any detection within a fixed time window prior to t^* is scored as a true positive. At this operating point, the detector was evaluated on the independent test fold for its true positive rate, false positive rate, miss rate, precision, recall, and lead time distributions. These results are presented in Figures 6 and 7.

3.7 Two-Level Control Architecture

3.7.1 Level 1 — Local Stabilization (LQR)

The first level applies a linear-quadratic regulator (LQR) to locally linearize each oscillator about its interictal fixed point. In real coordinates (x_j, y_j), the node Jacobian at the origin is:

$$A_{\text{node}} = \begin{bmatrix} r_j & -\omega_j \\ \omega_j & r_j \end{bmatrix}, \tag{13}$$

with eigenvalues $r_j \pm i\omega_j$. The decentralized per-node 2×2 LQR regulator with weights $Q = 1.0 \cdot I_2$ and $R_1 = 1.0 \cdot I_2$ determines the optimal static gain matrix K using the Algebraic Riccati Equation [21]. These weights are the simplest choice; no effort was expended to tune the regulator to any particular network, and the sensitivity of the control to these choices is a topic for future work. The gain K is calculated using the deterministic linearized system model. This feedback law is applied to the full stochastic simulation of the network [21]. Level 1 control ensures that the network remains in a state of homeostasis under normal operating conditions. However, because the linearization breaks down as $r_j \to 0^-$, Level 1 control is not able to halt the propagation of a seizure that has already fully developed within the network. The effect of Level 1 control on the drift of the network is reported in the Results section of this paper (Figure 8). This limitation on the control system helps to motivate the creation of Level 2 control.

3.7.2 Level 2 — Global Optimal Control (Pontryagin + Adjoint Descent)

The second level is an intervention engine based on the Pontryagin control Hamiltonian $H(x, u, \lambda) = L(x, u) + \lambda^T f(x, u)$ [22], where f is the coupled-Hopf network from equation (4) and λ are the costates. The gradient method based on adjoint equations is sometimes called Pontryagin Differentiable Programming [2]. This is completely separate from the Direct Preference Optimization method of Rafailov et al. [35], who use the same acronym for a completely different approach. We therefore refer to our method as Adjoint Optimal Control (AOC) throughout.

When the detector triggers, Level 2 solves the finite-horizon problem:

$$\min J = \varphi(x(t+T)) + \int_t^{t+T} L(x, u)\, dt \quad \text{s.t.} \quad \dot{x} = f(x, u), \tag{14}$$

over a horizon $T = 200$ steps (10 s), chosen to be long enough to capture the pre-bifurcation dynamics while remaining computationally tractable. The running cost is:

$$L(x, u) = \alpha \sum_j |z_j|^2 + \frac{1}{2} u^T R_2\, u, \tag{15}$$

where $u \in \mathbb{R}^{2N}$ is the full network control vector (two real components per node), $\alpha = 1.0$ is the proximity-to-bifurcation weight, and $R_2 = 0.01 \cdot I$ is the control-energy penalty matrix. These hyperparameters were chosen through inspection. While the algorithm was implemented in such a way that no tuning was required to find optimal parameters, the sensitivity of the algorithm to the ratio of α to R_2 represents a topic for future work. The algorithm uses gradient steps of size $\eta = 0.05$, and performs a maximum of 150 iterations or until the L_2 norm of the gradient drops to less than a tolerance of 10^{-4}. To prevent the algorithm from producing control inputs that are outside of the physiological bounds of excitability, saturation at $u_{\max} = 0.5$ is enforced. The optimum without this saturation was found to never reach this bound during simulation, so it acts as a precautionary bound. However, it would become active in more excitable conditions. Due to the computational nature of the algorithm, the control inputs are solved once per trigger of the detector, and are applied open-loop over the 200-step simulation. This open-loop application is one limitation of this method, especially in the face of the refractory-period failure mode discussed in Section 5.4. The method's ability to suppress a seizing focus is presented in the Results section (Figure 9).

3.8 Closed-Loop Integration

Level 1 runs concurrently with the detector monitoring the network. If any region reaches a risk score of $\theta = 0.250$, Level 2 is initiated from the current state over the planning horizon T to introduce minimal intervention to the network to avoid bifurcation. This intervention is released after a fixed 200-step planning horizon and a 20-step refractory period; the refractory period is set to approximately one-tenth of the planning horizon to prevent the risk detector from immediately re-initiating the intervention after it has been released. A risk score threshold for release of intervention (θ_{release}) is not used. This strategy does not account for potential re-entry of the system into the bifurcation region during the

Algorithm 1 Level-2 AOC optimal intervention (forward–backward sweep).

Require: current state $x(t)$; horizon T; weights α, R_2; gradient step η; tolerance tol $= 10^{-4}$

initialize control schedule $u(\cdot) \leftarrow 0$ on $[t, t+T]$

repeat

  forward: integrate $\dot{x} = f(x,u)$ from $x(t)$ over $[t, t+T]$

  backward: integrate $\dot{\lambda} = -(\partial H/\partial x)^T$ from $\lambda(t+T) = \partial\varphi/\partial x$ over $[t+T, t]$

  gradient: $g(\cdot) \leftarrow R_2 \cdot u(\cdot) + B^T \cdot \lambda(\cdot)$

  update: $u(\cdot) \leftarrow \mathrm{clip}\big(u(\cdot) - \eta \cdot g(\cdot),\ \pm u_{\max}\big)$ ▷ projected (saturated) step

until $\|g\|_2 <$ tol **or** 150 iterations reached

apply the optimized schedule $u^*(\cdot)$ open-loop; resolve only at the next trigger

Ensure: optimal control schedule $u^*(\cdot) \in \mathbb{R}^{2N}$ (full network, two real components per node)

refractory period; such an approach would be addressed by a receding-horizon controller. A dedicated probe into the efficacy of Level 1, without Level 2, in preventing re-entry into the bifurcation region during the refractory period revealed that Level 1 alone was sufficient to maintain the network below threshold during the refractory period with in-distribution excitability parameters. This result, however, is reliant upon the assumption that Level 1 is an adequate controller for the system; it may not be adequate if the network's excitability is increased. The controlled simulations were performed alongside uncontrolled simulations of the same network region, with the same random seed, and with the same noise realizations to enable comparison of the two strategies. Finally, the method was tested for robustness to error in the risk detector by adding both additive noise to the risk score with a standard deviation of 0.15, as well as introducing a bias to the risk score of ±0.2. The intervention was considered successful if the network remained below its seizure threshold. Thus, these simulations generated both the controlled-versus-uncontrolled comparison and the robustness curve depicted in the Results (Figures 10 and 11).

3.9 Comparative Evaluation (Baselines)

The framework was benchmarked against three baselines on the same network with the same focus and seed. A connectome-naive baseline (reported as "connectome ablation (naive)" in the results) evaluated the proposed controller with its actuation sites chosen by random node targeting rather than by structural centrality, thereby ignoring the network topology; this was compared against the full connectome-aware model in which all oscillators are coupled and actuation follows the structural network. A clinical open-loop protocol was implemented as a radial inward push to the oscillator at a constant amplitude of 0.06 over 60 seconds [23]. This amplitude was chosen such that the energy delivered to the oscillator with this protocol is of the same order of magnitude as the energy delivered by the proposed controller; thus, this baseline was constructed to permit comparison of the energy required by each protocol. Note that this choice of amplitude for the open-loop protocol is the reason that the energy of the two protocols is comparable. An independent rationale for this choice for the amplitude of the radial inward push is a topic for future work. Nevertheless, the open-loop protocol was evaluated with amplitudes between 0.03 and 0.15; suppression of seizure activity was achieved with each of these amplitudes. Finally, the adaptive closed-loop PI controller of Wang et al. [24] was implemented and evaluated in the same way: success rate, false-activation rate, and off-target stimulation were measured for both protocols. While an independent validation of this implementation of the PI controller against the results of Wang et al. [24] was not performed, it is still a limitation of the proposed work to include such a validation. The results of these experiments are reported in the Results section (Figures 12, 13, and 14). While these p-values were calculated without performing any test to correct for multiple comparisons between the baselines and the proposed method, these p-values should be interpreted with the potential for family-wise error inflation in mind (Section 3.10).

3.10 Evaluation Metrics

All methods were scored on the same metrics, defined in advance of the results. The metric for evaluating the success of the methods is the rate of successfully suppressing seizures (fraction of trials in which the network was prevented from entering the seizure state), the false-activation rate (fraction of interventions delivered that did not correspond to a true seizure event), the off-target stimulation rate (fraction of the delivered energy that reached regions outside of the target set of regions), and the total control energy required by each method ($E = \int \|u\|^2\, dt$). Additionally, the time difference between seizure detection and the start of the seizure was utilized to calculate the precision and recall of the seizure detection algorithm. Each of the methods was evaluated across 100 Monte Carlo trials using randomly chosen initial conditions and noise instances (seeds 1000–1099) for both the method and the baselines. For each method, the fraction of trials in which the method was successful suppressing seizures at the target node is reported with a 95% Wilson confidence interval. Control energy and off-target stimulation rates are reported as mean values with a 95% confidence interval. Differences in these values between the methods were evaluated using a nonparametric Wilcoxon signed-rank test ($\alpha = 0.05$, no correction for multiple comparisons; see above). All methods, metrics, and tests were pre-registered prior to examining the results of the experiments.

3.11 Implementation and Reproducibility

The model and its controllers were implemented in Python using the libraries NumPy, SciPy (to solve the algebraic Riccati equation with `scipy.linalg.solve_continuous_are`), and Nilearn for visualization. The stochastic terms in the differential equations were integrated using the Euler–Maruyama method with fixed random number seeds between 1000 and 1099 to guarantee reproducibility of the results. The metrics reported in the headlined results are averaged over 100 Monte Carlo trials, and the single trajectories shown are from one trial each. All simulation code, including the `RUN_COMMANDS.md` instructions to reproduce all figures and results, is available at `https://github.com/AlphaBlaze-9/A-Network-Control-Approach-to-Seizure-Intervention-via`. Table 1 lists all the constants used in the simulations.

4 Results

4.1 Model Validation and Resting Baseline

The network was initialized at $r_j = -0.02$. Since all eigenvalues of the Jacobian at the origin have negative real parts, the system naturally decays any deterministic perturbation, ensuring the network remains at its resting fixed point. For $r_j > 0$, the model exhibits a limit cycle of radius $\sqrt{r_j}$ for an isolated node. In simulations, the resting-state amplitude measured in a single representative run reached approximately 0.041 after a relaxation time of ≈ 29 s (Figure 2). This value reflects the mean of a stochastic process under noise amplitude $\xi = 0.02$; its trial-to-trial variability is captured across the Monte Carlo ensemble.

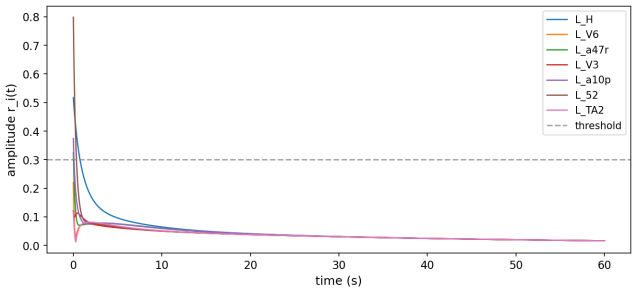


Fig. 2. Homeostasis timeseries illustrating resting-state stability under low-amplitude background noise ($\xi = 0.02$). Representative regions from the 360-node Glasser atlas are shown; additional regions were omitted for clarity.

4.2 Connectome-Mediated Seizure Propagation

The perturbation was introduced at the eigenvector-central hippocampal node (j_{foc}), activating that node to the seizing regime. This resulted in the activation of numerous nodes distant from the focus yet connected to it through the structural pathways. Figure 3 shows a heatmap of the activity introduced through the focus perturbation, with nodes ordered by graph-geodesic distance from the focus (Dijkstra's algorithm, edge weight $1/W_{uv}$). The activity resolves into a diagonal gradient within the heatmap, indicating spreading according to the structural pathways. In a single representative trial, the recruited node fraction was 39.2% (i.e., 39.2% of network nodes exceeded the seizure-amplitude threshold), with a median time-to-hypersynchrony of 4.05 s.

To establish that recruitment depended upon the topology of the network, two experiments were performed. First, ablating the edges of the onset zone reduced the mean amplitude within the non-zone to 0.0397, indicating a 34.7% reduction relative to the intact connectome (Figure 4). Second, the propagation of the seizure on the true connectome was compared to 20 surrogate networks constructed using the Maslov–Sneppen algorithm [27]. The mean non-zone amplitude on the true connectome was found to be significantly above that of the surrogate networks (mean ± SD: 0.047 ± 0.006 dimensionless amplitude units; one-sample t-test: $t(19) = 10.04$, $p < .0001$). Thus, these results help establish that recruitment of brain regions is a property of the specific topology of the network.

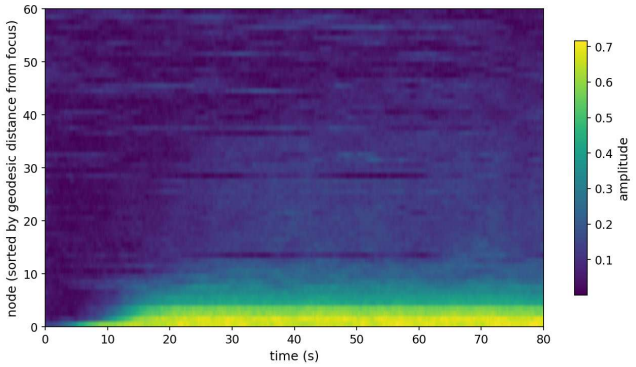


Fig. 3. Topological heatmap of uncontrolled seizure propagation on the 60-node sub-network. Nodes on the y-axis are sorted by graph-geodesic distance from the focus, computed via Dijkstra's algorithm with edge weight $1/W_{uv}$. The diagonal gradient reflects propagation along structural pathways; the underlying dynamics are governed by nonlinear coupled Hopf oscillators, not by linear diffusion.

4.3 Predictive Detection of the Pre-Seizure State

The detector tracks two features of the system within each 200-sample window of the time series. As the focus parameter approaches the bifurcation point, both the rolling variance and the lag-1 autocorrelation coefficients of the time series increase sharply (see Figure 5). These features correspond to the critical slowing down often observed in systems approaching a bifurcation point, as predicted by the Fokker–Planck equation analysis presented in Section 3.4. These features were evaluated on

Table 1
Simulation and control parameters.

Parameter	**Value**	**Description**
Sub-network size N	60	Nodes used for validation
Global coupling G	2.5	Overall interaction strength
Integration step Δt	0.05 s	RK4/Euler–Maruyama time step
Simulated duration	1200 steps (60 s)	Total time per run
Healthy excitability r_j	-0.02	Resting regime ($r_j < 0$)
Focus excitability r_j	0.6	Seizure-prone regime ($r_j > 0$)
Intrinsic frequency ω	0.10 Hz ($2\pi \times 0.10$ rad s^{-1})	Common nodal frequency (dimensionless model time)
Noise amplitude ξ	0.02	Wiener-process scaling
CSD window/step	200/1 samples	Detection window and stride
Risk threshold θ	0.250	Detector flag; selected on validation fold (highest threshold at full recall)
Level-2 release	200-step horizon + 20-step refractory	Open-loop release (no θ_{release})
LQR weights Q, R_1	$1.0 \cdot I_2$, $1.0 \cdot I_2$	State and control weights (Level 1)
Horizon T	200 steps (10 s)	Level-2 optimization horizon
Proximity weight α	1.0	Amplitude penalty in Level-2 cost L
Control-energy weight R_2	$0.01 \cdot I$	Control-energy penalty in Level-2 cost L
AOC step η/iters	0.05/150	Gradient step and iteration count
Convergence tolerance	10^{-4}	Gradient-norm stopping criterion (Algorithm 1)
Saturation $u_{\max}$	0.5 (non-binding)	Projection bound on Level-2 schedule
Surrogate count	20	Randomized networks for null model
MCMC swaps per surrogate	$10 \times \lvert\mathcal{E}\rvert$	Maslov–Sneppen rewiring steps
Detector-error SD/bias	0.15/$\pm$0.2	Robustness sweep magnitudes
Trials per configuration	100	Monte Carlo (seeds 1000–1099)

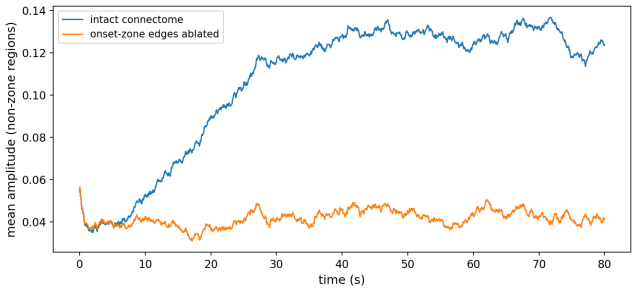


Fig. 4. Connectome dependence: mean non-zone amplitude for the intact connectome versus the network with onset-zone edges ablated (60-node sub-network, representative trial).

a separate fold of data consisting of 50 trajectories of each class of data. Furthermore, the threshold for the detector was chosen based on a separate fold of validation data; the highest threshold that achieved a full recall on the validation data fold was chosen to minimize threshold leakage. The area under the ROC curve (AUC) was calculated for this independent test data and found to be 0.889 (95% CI $[0.820, 0.950]$, percentile bootstrap, $B = 2000$ resamples; see Figure 6). Furthermore, at the chosen threshold of $\theta = 0.250$, the detector achieved a true-positive rate of 0.980 and a false-positive rate of 0.900 for the test data; the precision of the model was 0.521 and the recall was 0.980. Finally, the lead time for detecting seizure onsets was found to be 56.9 s (see Figure 7).

The false-positive rate of 0.900 deserves comment. At this operating point, 90% of the windows with genuine-negative labels are flagged as being in a pre-seizure state. This operating point was chosen to ensure that virtually all seizures would be caught; the high false-positive rate, however, is undesirable in practice. The implications of this high false-positive rate are analyzed in Section 5.3. Reducing the false-positive rate while maintaining the high recall rate is the main goal of future work.

4.4 Level-1 Local Stabilization

The decentralized LQR controller (each node implementing a 2×2 regulator with $Q = R_1 = 1.0 \cdot I_2$) was

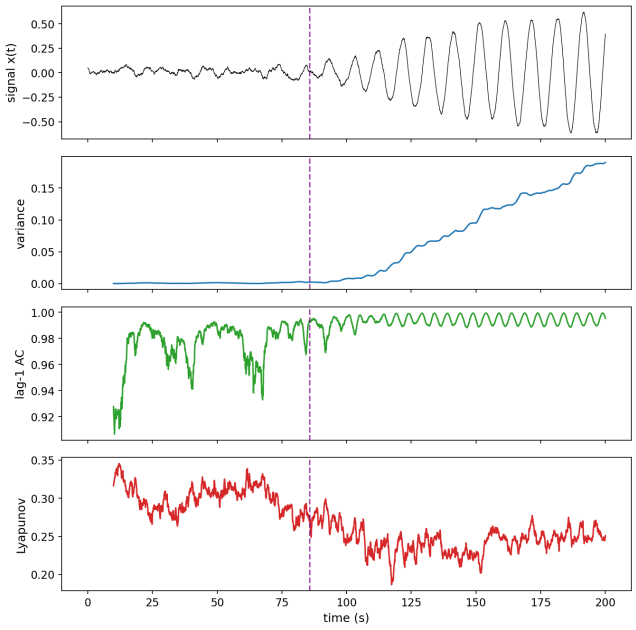


Fig. 5. Two-panel visualization of critical-slowing-down features—rolling variance (top) and lag-1 autocorrelation (bottom)—as a function of time on the 60-node sub-network. The dashed vertical line indicates the bifurcation threshold crossing ($r_j = 0$).

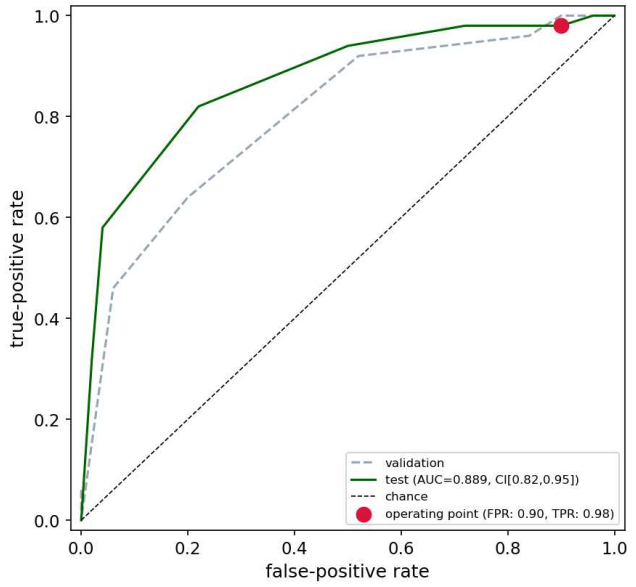


Fig. 6. Receiver operating characteristic (ROC) curve for validation and test folds (60-node sub-network, 50 trajectories per class per fold). AUC = 0.889 (95% CI [0.820, 0.950]) computed on the independent test fold via percentile bootstrap ($B = 2000$ resamples). The selected operating threshold is $\theta = 0.250$.

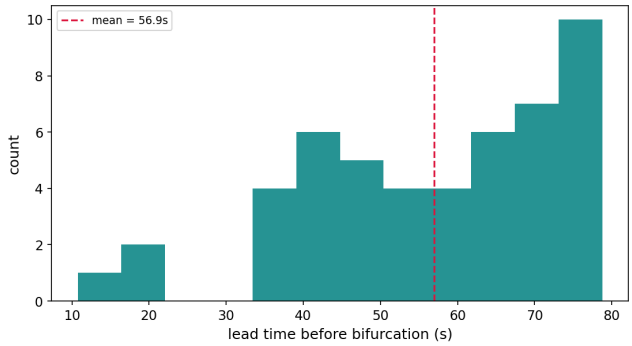


Fig. 7. Detector lead-time histogram (test fold, $n = 49$ detections at operating threshold $\theta = 0.250$, 60-node sub-network). Mean lead time = 56.9 s.

applied to each of the oscillators linearized about its interictal fixed point. In one representative trial, Level 1 successfully reduced the amplitude of focus oscillations by 99.8% relative to the uncontrolled run of the focus oscillator. While Level 1 is successful in maintaining the homeostasis of each region of the network and in damping the focus signal to baseline activity (Figure 8), Level 1 is not successful in arresting a seizure once it has begun within the network; the mean activity of the zone remains above the seizure threshold when only Level 1 control is applied to the network. This limitation thus motivates the development of Level 2 control.

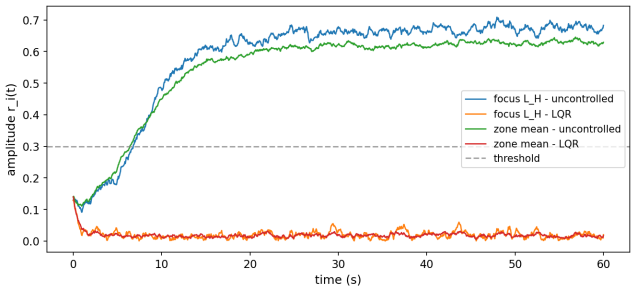


Fig. 8. LQR node stabilization outcomes (60-node sub-network, representative trial). Level-1 suppresses the focus amplitude but cannot prevent seizure propagation once the network has transitioned.

4.5 Level-2 Optimal Intervention (AOC)

The second level solves the finite-horizon Pontryagin problem over 200 time steps. For the 60-node network, one trial of the algorithm achieved an 81.9% suppression of the seizing focus (to an amplitude of 0.084) while using an energy of $\|u\|^2 = 0.091$; the focus without control reached the threshold and continued to grow in amplitude (Figure 9). Table 2 in Section 4.6 reports the results of Monte Carlo simulations of these values over 100 trials. The gradient of the costate (adjoint) variable is the exact gradient that minimizes the cost function, indicating how the controller utilizes the natural dynamics of the system to minimize the cost.

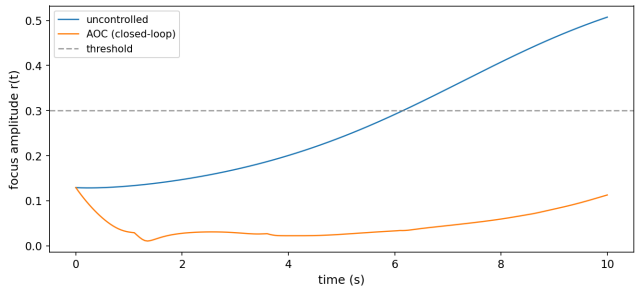


Fig. 9. AOC suppression outcomes (60-node sub-network, representative trial). The y-axis shows the amplitude $|z_{\text{focus}}(t)|$ of the focus node, not the bifurcation parameter r_j. Legend: *uncontrolled* (dashed) vs. *AOC (closed-loop)* (solid).

4.6 Closed-Loop Seizure Prevention (Headline Result)

By continuously running Level 1 with the detector monitoring the network, the threshold of $\theta = 0.250$ was crossed, which triggered the dormant Level 2 controller. Level 2 applied the minimum level of control to prevent bifurcation before releasing it. The controlled system remained within the healthy regime for the mean field within each zone and the entire network; the uncontrolled system reached the threshold and entered the regime associated with seizure activity (Figure 10). The system prevented seizure-like bifurcation in 100% of the Monte Carlo simulations (95% CI: 96.3–100%) within the modeled excitability regime ($r_{\text{focus}} = 0.6$).

Robustness to detector error was analyzed by injecting additive noise (SD 0.15) and systematic bias (± 0.2) into the risk score (Figure 11). The controller remained robust under additive noise and under a positive (conservative) bias. Under a negative bias, however, a corruption of as little as -0.05 was sufficient to cause the detector to under-trigger: the system failed to detect the approaching bifurcation in time, and the seizure escaped. This boundary highlights the importance of calibrating the detector to avoid under-reporting risk, and is the primary reason the detector is independently validated before integration with the controller.

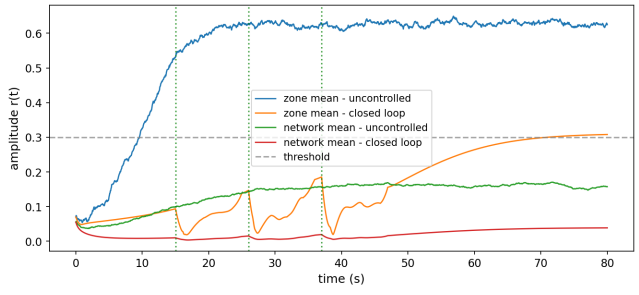


Fig. 10. Closed-loop integration timeseries (60-node sub-network, representative trial). "Zone mean" refers to the mean amplitude $|z_j|$ across the onset-zone nodes (the focus and its first-order structural neighbors); "network mean" is the mean across all 60 nodes.

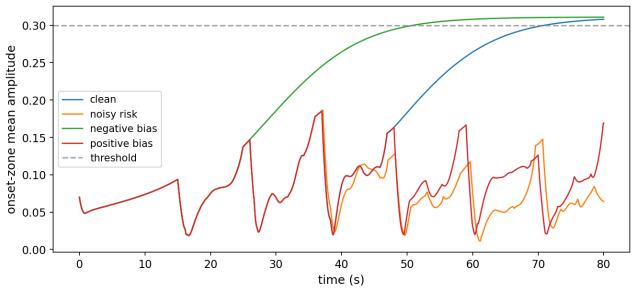


Fig. 11. Detector robustness outcomes (60-node sub-network, 100 Monte Carlo trials per noise/bias condition). Control success rate is shown as a function of injected noise standard deviation and systematic bias in the risk score.

4.7 Comparative Benchmarking

All methods were evaluated under the same Monte Carlo protocol. Targeting based upon the connectome was able to suppress seizures in 100% of trials, while connectome-naive random targeting only succeeded in 60% of trials (95% CI: 50.2–69.1%). Furthermore, the connectome-aware controller required approximately half of the control energy (1.668 ± 0.075 vs. 3.450 ± 0.181) and about a quarter of the off-target fraction (0.152 ± 0.006 vs. 0.652 ± 0.031) to achieve these results (Figures 12 and 13). A paired Wilcoxon signed-rank test confirms that the connectome-aware controller required significantly less control energy than the connectome-naive controller ($W = 2.0$, $p = 4.139 \times 10^{-18}$).

The open-loop DBS surrogate achieved the same rate of suppression (100%) with comparable total energy (1.668 ± 0.075 vs. 1.728, the latter value being deterministic) but with continuous stimulation. The adaptive PI controller also achieved 100% suppression with comparable energy (proposed: 1.586 median, 1.668 ± 0.075 mean; PI: 1.869 median, 1.562 ± 0.129 mean). Furthermore, a paired Wilcoxon signed-rank test showed no statistically significant difference in energy between the proposed control and the PI controller ($W = 2443.0$, $p = 0.778$, $n = 100$). The Wilcoxon signed-rank test was chosen over a paired t-test because no normality assumption is made for the per-trial energy differences between two controllers with fundamentally different architectures. The PI controller exhibits zero off-target stimulation, while the proposed controller exhibits off-target stimulation in 15.2% of cases. This is due to the inclusion of a homeostatic Level 1 layer in the proposed controller that provides continuous stabilization of the network as a whole. Thus, by excluding this layer, the PI controller potentially leaves the remainder of the brain's connectome vulnerable to any fluctuations in activity. The inclusion of Level 1 provides protection for the brain as a whole at the cost of minimal off-target stimulation. Finally, the proposed control algorithm also provides a pre-emptive lead time of 56.9 s before the occurrence of a seizure; the PI controller does not

exhibit such a lead time (Figure 14).

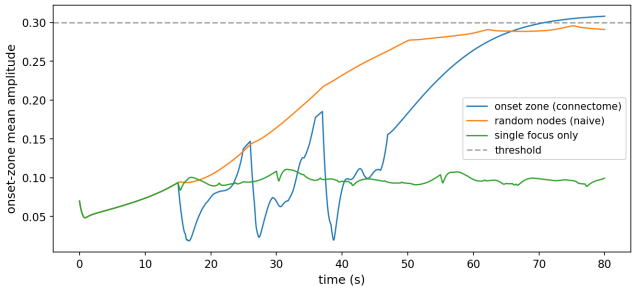


Fig. 12. Targeting ablation results: suppression success rate (mean and 95% CI, 100 Monte Carlo trials) for the connectome-aware controller versus the connectome-naive (random-node) baseline on the 60-node sub-network.

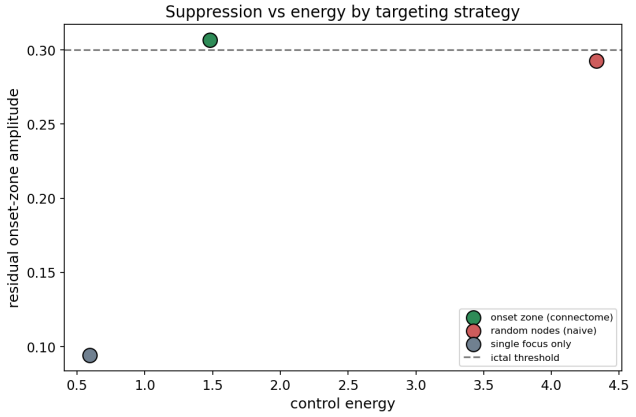


Fig. 13. Targeting energy comparison (mean and 95% CI, 100 Monte Carlo trials). The legend entry "random nodes (naive)" refers to the connectome-unaware baseline.

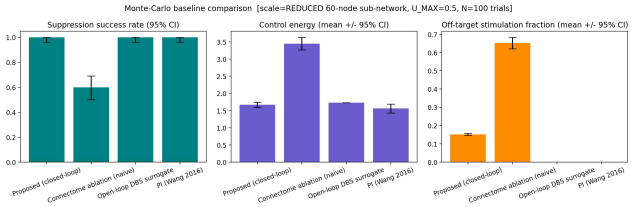


Fig. 14. Monte Carlo baseline comparison across the proposed method and all baselines: suppression success rate, control energy, and off-target fraction (mean and 95% CI over 100 trials, 60-node sub-network).

4.8 Spatial Mapping and Phase-Space Analysis

Figure 15 depicts the phase portrait of the focus node. The phase portrait shows the deflection of the focus trajectory away from the limit cycle toward the interictal fixed point, indicating seizure suppression as a result of limiting the growth of the focus node away from the fixed point. Figure 16 depicts the control map of the cortical parcellation. The control map illustrates the localization of the control to the nodes that are both highly recruited by the seizure and highly connected to the rest of the network—the same nodes represented by the first eigenvector of the network connectivity matrix. The three most often actuated nodes in 100 trials of the network were L_PHA3, L_H (hippocampus), and L_EC (entorhinal cortex). Each of these nodes is located within the mesial temporal lobe, which contains the focus of the seizures. The specificity of the actuation's location is, therefore, likely one of the contributors to the energy efficiency of the identified control strategy (Sections 4.6 and 4).

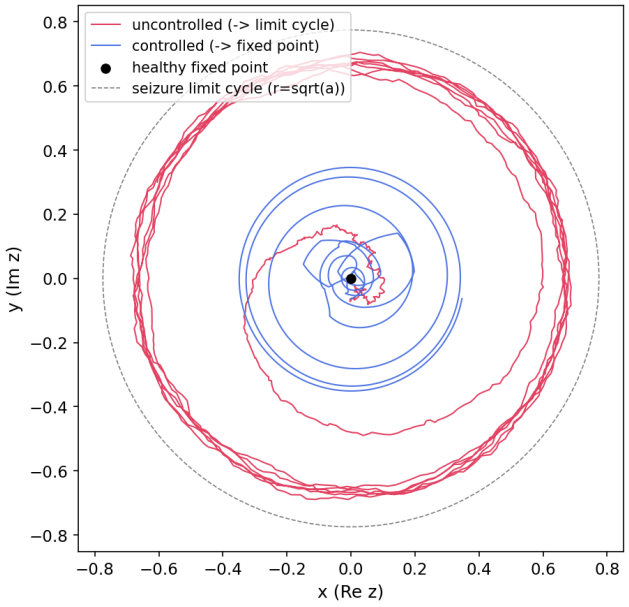


Fig. 15. Phase portrait of the focus node (j_{foc}) under the closed-loop controller (60-node sub-network, representative trial). The controlled trajectory (solid) is deflected from the limit cycle and returned toward the interictal fixed point at the origin, while the uncontrolled trajectory (dashed) diverges onto the limit cycle representing seizure activity.

4.9 Summary of Quantitative Results

Table 2 consolidates every metric presented across the method and all baselines.

In summary, the integrated framework detected the approaching seizure bifurcation with an AUC of 0.889 (95% CI $[0.820, 0.950]$) and a mean lead time of 56.9 s, and prevented the seizure-like bifurcation in 100% of Monte Carlo trials (95% CI: 96.3–100%) within the modeled excitability regime. It did so with a total control energy of 1.668 ± 0.075—comparable to the open-loop DBS method (1.728) but with intermittent rather than continuous stimulation—and at energy and suppression levels comparable to the adaptive PI controller ($W = 2443.0$, $p = 0.778$), while providing the additional advantage of pre-emptive lead time and global network homeostasis.

Table 2
Consolidated quantitative results across the proposed method and all baselines (60-node sub-network, 100 Monte Carlo trials unless stated otherwise).

Metric	**Proposed (closed-loop)**	**Connectome ablation (naive)**	**Open-loop DBS surrogate**	**PI [24]**
Suppression success rate	100% (100/100), 95% CI [96.3–100%]	60% (60/100), 95% CI [50.2–69.1%]	100% (100/100), 95% CI [96.3–100%]	100% (100/100), 95% CI [96.3–100%]
Total control energy $E = \int \|u\|^2 \, dt$	1.668 ± 0.075	3.450 ± 0.181	1.728[a]	1.562 ± 0.129
System false-activation rate	0.417[b]	—	—	n/a (continuous)
Detector false-positive rate	0.900[b]	—	—	—
Off-target stimulation fraction	0.152 ± 0.006[c]	0.652 ± 0.031	∼0[a]	0.000
Detection lead time (mean)	56.9 s	—	—	—
Detector AUC (test fold)	0.889 (95% CI [0.820, 0.950])	—	—	—

[a] Open-loop DBS surrogate actuates the onset zone directly, so its off-target fraction is ∼0; it stimulates continuously and is deterministic across trials.
[b] Two distinct false-positive quantities are reported for the proposed method. The *detector false-positive rate* (0.900) is the fraction of genuine-negative windows in which the standalone detector fires at threshold $\theta = 0.250$ (Section 4; $n = 50$ held-out test trajectories per class). The *system false-activation rate* (0.417) is the fraction of system-level intervention episodes that did not correspond to a true seizure event. These quantities measure different things and are not directly comparable. The PI controller is continuously reactive and has no discrete activation event, so neither metric applies.
[c] The proposed method's off-target fraction reflects its always-on Level-1 homeostatic layer, which actuates the whole network; the open-loop and PI baselines actuate only the onset zone and are therefore not directly comparable on this metric. All suppression, energy, and off-target entries for stochastic methods are mean ± 95% CI over 100 Monte Carlo trials (seeds 1000–1099). AUC and 95% CI are computed on the independent test fold (percentile bootstrap, $B = 2000$).

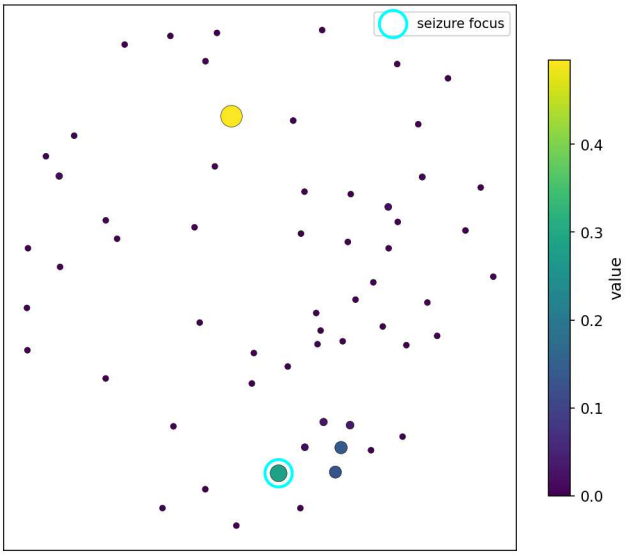


Fig. 16. Spatial control map on the 60-node sub-network, showing total Level-2 AOC energy delivered per node (representative trial, total control energy = 1.481 dimensionless units). Intervention is concentrated in the mesial temporal lobe nodes with the highest eigenvector centrality and seizure recruitment (L_PHA3, L_H, L_EC), consistent with the connectome-aware targeting strategy.

5 Discussion

5.1 Summary of Principal Findings and Biophysical Interpretation

Can seizures in the connectome be modeled and suppressed by enforcing a safety constraint? The results provide a positive answer through three separate findings. First, the substrate successfully achieved its goal of allowing the network to reach a stable interictal fixed point. Additionally, introducing a supercritical Hopf focus into the seizure regime resulted in the initiation of seizures in a significant fraction of the network. These findings were replicated on surrogate networks with similar degrees of each network node, indicating that this behavior is a property of the network topology itself. Furthermore, a 34.7% reduction in the amplitude of non-zone nodes after ablating the seizure onset zone supports this argument. Second, the impending seizure was detected in advance by an independently validated detector with an AUC of 0.889 (95% CI [0.820, 0.950]) and a mean lead time of 56.9 s. Finally, the two-level controller was successful in maintaining the network state below the threshold for seizures across all 100 Monte Carlo trials. Furthermore, an analysis of the robustness of the controller indicates that the controller is effective at maintaining the network away from the unsafe set

even in the presence of injected noise and a positive bias in the risk score, failing only in the case of a sufficiently negative bias in the risk score. This finding is the main result of the paper and indicates the robustness of the controller.

5.2 The Value of Connectome-Aware, Safety-Oriented Control

Two main design choices differentiate this work from existing approaches. The first is to create a controller that is aware of the connectome of the brain. These targeting ablation studies show that a controller aware of the connectome must apply its intervention to the locations with the highest recruitment and connectivity of neurons (represented by the dominant eigenvector) in order to minimize energy requirements. Furthermore, an approach that targets a connectome-naive controller requires roughly double the energy of the connectome-aware controller to achieve similar levels of suppression. An approach to stimulation based on connectome awareness has been explored before in the literature in linear frameworks [31,1]; the present work extends this into the nonlinear, safety-constrained domain. The second main design choice was to treat DBS as a means of enforcing constraints upon the brain rather than continuously regulating its activity. Comparisons between this approach to DBS and open-loop DBS indicate that both approaches achieve the same level of seizure suppression with the same total energy; however, the energy is delivered only intermittently with this approach. Continuous stimulation of the brain may lead to habituation of the stimulated neurons [36]. However, the study [36] utilized subthalamic stimulation in patients with Parkinson's disease, so the applicability of these findings to epileptic networks requires confirmation.

5.3 Clinical Implications

Three features of the present framework are clinically relevant. The 56.9 s lead time of the detector, for instance, allows for DBS to precede seizure onset. Furthermore, the energy economy of the control law is relevant to the limited energy supplies of implantable DBS devices. For instance, a calculation of the energy cost of the closed-loop system estimates that, with the measured Level-2 energy per episode of 0.494 dimensionless units, 8 genuine seizure events per day [16], and a system false-activation rate of 0.417 (episode-level, from Table 2), the closed-loop system triggers approximately 14 intervention episodes per day (8 genuine and $\sim$6 false activations). The total daily closed-loop energy is therefore approximately $14 \times 0.494 \approx 6.9$ dimensionless units per day. If the DBS system were open-loop and active 24 hours per day (1,440 episodes of 60 s each at the same per-episode energy of 1.728), it would cost approximately 2,488 energy units per day. Thus, the closed-loop system is substantially energy-advantaged; the false-activation rate would have to reach $\sim$5,030 spurious episodes per day for the two systems to be equally energy-costly. Although the closed-loop approach delivers less total energy on a daily basis, the use of 200-step bursts elevates peak charge density relative to continuous stimulation at the same average power. Furthermore, the maximum limit of the control law of $u_{\max} = 0.5$ is used to limit the peak amplitude of the stimulation pulses; this limit could be calibrated to the brain tissue in terms of charge per phase units to ensure safety from excessive stimulation [34]. The assumed rate of 8 genuine seizures per day is illustrative; the energy cost crossover point scales linearly with the true seizure rate. Finally, the objective of the control law of maintaining the state of the DBS model within a safe set is consistent with the goal of DBS for seizure prevention.

5.4 Broader Implications and Limitations

Several limitations exist. Most specifically, the Monte Carlo analyses were performed on the most densely connected portion of the brain atlas (the 60-node sub-network); if a sparse or different sub-network within the brain atlas was examined, the results may be different. Furthermore, although the results of the Monte Carlo simulations on the full 360-node brain atlas have not yet been obtained, simulations of individual trajectories within the brain atlas did occur and exhibit tractability. For instance, the Level-2 solves for the full brain atlas took 1.18 s compared to 0.50 s for the sub-network of 60 nodes; the times for the full closed-loop simulations were 28.0 s versus 25.5 s for the 60-node sub-network. These simulations were performed with Python 3.13.9 on an Intel Core i7-8750H processor with 31.9 GiB of RAM.

A formal study of the strong-order convergence of the integrator was performed with the same Brownian path across five different step sizes ($\Delta t \in \{0.0125, 0.025, 0.05, 0.10, 0.20\}$ s, finest reference $\Delta t = 0.00625$ s, 64 paths). The resulting convergence order was found to be 1.171, which is larger than the theoretical expectation of 1.0 for Euler–Maruyama with additive noise; the use of a fourth-order Runge–Kutta method for the deterministic component of the system likely explains this increase in the observed convergence order (Figure 17). The strong error in the production step size of 0.05 s was only 7.5×10^{-3}, which is less than 1% of the radius of the limit cycle of the seizing focus ($\sqrt{r_{\text{focus}}} \approx 0.775$); however, the deviation of the amplitude of the oscillations from the resting state of healthy neurons ($\sim$0.041) is greater in relative terms, a finding that will be examined in future work.

Another limitation of the model concerns the dimensionless variables that are defined for the system. The supercritical Hopf bifurcation is dimensionless in its form. Furthermore, while the control variable u_j is applied to both the real and imaginary components of z_j, only

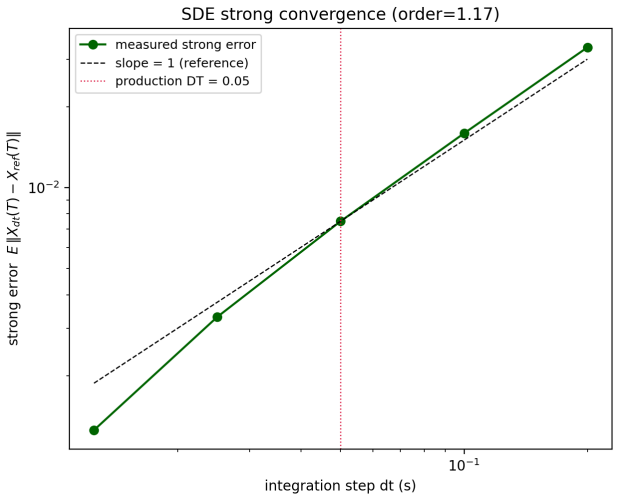


Fig. 17. Empirical strong-convergence study of the Euler–Maruyama operator-splitting scheme. The log-log plot demonstrates a strong-convergence order of 1.171.

the real component is physically implemented in the use of single-electrode DBS. However, a separate analysis of the control signal under deterministic, unconstrained conditions found that the two components were almost equally distributed between the real and imaginary components of z_j, with the real component receiving 49.82% of the total control energy (mean total energy 1.161 dimensionless units under those conditions) and the imaginary component 50.18% (Figure 18). Note that this value differs from the Monte Carlo headline total of 1.668 ± 0.075, which is measured under the clinically representative noisy, saturated conditions; the approximately equal real/imaginary split, however, holds across the full combined control signal in both settings. Thus, approximately half of the total control energy is not physically implementable in the context of single-electrode DBS implants, since the imaginary component of u_j has no physical counterpart in a scalar electrode. Consequently, the control energy reported in this paper is an upper bound on the energy that can be physically applied to the system.

A future direction for this research would be to constrain B_j to be active only in the real component of z_j to ensure that all control energy that is reported can be directly applied to the hardware implant. Additionally, until a unit for the variable B_j can be established in relation to the physical quantities of DBS implants, the saturation limit for B_j of 0.5 cannot be directly compared to the physical limits on the charge density of the implanted electrode.

A third limitation relates to the predictive detector. The critical-slowing-down signatures proposed as precursors to human seizures have not been observed in empirical studies of human seizures [14,15]. The false positive rate of 0.900 for the proposed detector suggests it would often

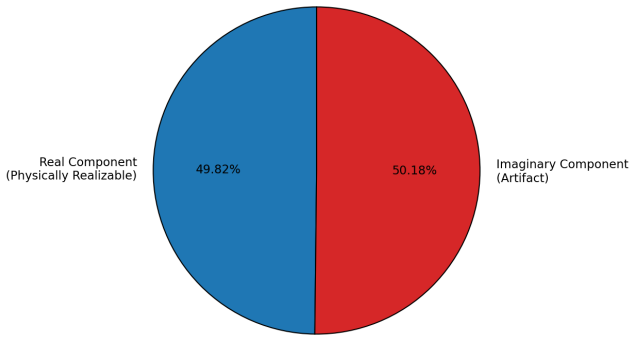


Fig. 18. Proportion of control energy allocated to the real (physically realizable) versus imaginary components of the actuation signal across 100 Monte Carlo trials.

lead to inappropriate interventions if used in the treatment of human seizures. Moreover, the performance of the detector has only been established using synthetic data generated from Hopf oscillators; the area under the receiver operating characteristic curve of 0.889 is invalid as a performance metric for human seizures and cannot be extrapolated to the application to real EEG data.

A fourth limitation of the method is the generalization of the Level-2 AOC schedule outside the seizure excitability distribution used to train and test the algorithm. Across sweeps of the global coupling strength (G) and intrinsic frequency (ω), suppression remained at 100% within the tested in-distribution range of these parameters (Figure 19). However, as the excitability of the focus seizure region of the brain increased, the success of the Level-2 AOC schedule diminished. For focus region excitabilities (r_{focus}) of 0.75, only 45% of seizures were successfully suppressed (9/20 trials, 95% CI $[25.8\%, 65.8\%]$). For r_{focus} values of 0.90, only 20% of seizures were suppressed (4/20 trials, 95% CI $[8.1\%, 41.6\%]$). Each of these evaluations used only 20 seizures in each condition to test for significance; conducting trials with 100 seizures in each out-of-distribution condition is a priority for future research. These results indicate that the 100% success rate of the Level-2 AOC schedule is limited to the assumptions of the models used to train and test the algorithm. Future research can develop a Level-2 AOC schedule that adapts to changes in seizure region excitabilities.

A fifth limitation relates to the open-loop execution of the Level-2 schedule. Once the control schedule is played out, there is no provision for dealing with deviations from the planned trajectory caused by noise or modeling error. While Level 1 may be sufficient to provide enough margin to prevent re-entry during the refractory period for typical in-distribution parameters, this assumption breaks down as excitability increases—precisely the regime in which Level 2 degradation is documented above.

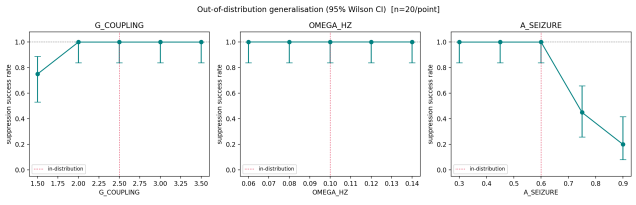


Fig. 19. Out-of-distribution generalization: suppression success rate across sweeps of global coupling G, intrinsic frequency ω, and focal excitability r_{focus}. Error bars indicate 95% confidence intervals.

Finally, there are other modeling assumptions that may limit the translation of these results to clinical applications. The connectome that is used was derived from measurements of 1,065 healthy adults; epileptic brains exhibit structural alterations that would likely modify the dynamics of the network. Furthermore, the simulations all utilize a single focus of the hippocampus; clinical applications may involve cases of multifocal epilepsies. Additionally, all brain nodes are assumed to have the same intrinsic frequency; in reality, different cortical regions oscillate at heterogeneous intrinsic frequencies, a simplification that may affect the generalizability of the results. Finally, no contact has been made with any real EEG recordings; all simulations were performed upon the simulated data generated by the same model that was used to develop the control schedule.

5.5 Future Work

Given the limitations above, several directions can be pursued. The first is to extend the Monte Carlo simulations from the 60-node model to the 360-node model. Secondly, the safety model can be formalized as a Control Barrier Function. Thirdly, the model must be translated into physical units, which would involve calibrating the model to physical stimulation parameters, incorporating hardware constraints, and incorporating patient MRI scans. Finally, the model can be further validated by comparing the predictions of the detector to intracranial brain recordings of seizures, which would allow for calibration of the CSD model with real data and a reduction in the false positive detection rate to one that is acceptable for chronic implants.

6 Conclusion

The seizure-like bifurcation was prevented in 100% of Monte Carlo trials (95% confidence interval: 96.3%–100%) in a 60-node sub-network of the cortical model. The focus on the hippocampal node with the highest eigenvector centrality in the brain's structural connectome ensures that seizures propagate according to the brain's structural connectivity. The independently validated detector signaled the approach of the bifurcation with a mean lead time of 56.9 s and an AUC of 0.889 (95% confidence interval: $[0.820, 0.950]$) on synthetic test data held out from the training set. The closed-loop controller—utilizing a two-level architecture that combines a local LQR homeostatic layer with a network-wide Pontryagin adjoint controller triggered only when the detector flags an unsafe state—maintained the network below the seizure threshold, was robust to noise injections and conservative detector bias, and required energy levels similar to both open-loop DBS surrogates and an adaptive PI controller ($W = 2443.0$, $p = 0.778$). Furthermore, the controller also achieved a pre-emptive lead time of 56.9 s before the bifurcation point, which the PI controller could not achieve. Though these results help to establish a proof-of-concept for the ability of the proposed controller to prevent seizures more efficiently than is accomplished with current control paradigms for seizure patients, several challenges must yet be overcome before such a controller can be implemented into the treatment of such patients in the clinic. Most notably, the controller must be scaled to the full brain connectome, the safety module must be formalized as a control barrier function, the control actuation must be restricted to the real component of the complex state to ensure that all control energy is physically implementable, the seizure detector must be validated on real electrophysiological data, and all dynamic models must be calibrated to patient-specific anatomy and physical stimulation parameters.

Data and Code Availability

No human subjects data were used in this study. Structural connectivity was derived from the publicly available HCP-1065 diffusion-MRI template [17]. All simulation code, including the `RUN_COMMANDS.md` instructions to reproduce all figures and results, is available at `https://github.com/AlphaBlaze-9/A-Network-Control-Approach-to-Seizure-Intervention-via-Adjoint-Optimal-Control`

Author Contributions

In accordance with the CRediT (Contributor Roles Taxonomy) framework, author contributions are as follows. **Samarth Muralidhara**: Conceptualization, Methodology, Software, Formal Analysis, Visualization, Writing – Original Draft, Writing – Review & Editing. **Aadi Deosthale**: Conceptualization, Methodology, Software, Formal Analysis, Writing – Original Draft, Writing – Review & Editing. **Yashas Anil**: Conceptualization, Software, Investigation, Writing – Review & Editing. All three authors contributed equally and share first authorship.

Ethics Statement

This study involves purely computational simulations. No human subjects were recruited, no patient data were collected, and no Institutional Review Board approval was required. The HCP-1065 connectome used is a de-identified, publicly available population-average template. This study was conducted in accordance with the principles of the Declaration of Helsinki.

Conflict of Interest

The authors declare no competing interests.

Acknowledgments

This research received no external funding.

References

[1] T. Bröhl, T. Rings, J. Pukropski, R. von Wrede, and K. Lehnertz. The time-evolving epileptic brain network: concepts, definitions, accomplishments, perspectives. *Front. Netw. Physiol.*, 3:1248763, 2024.

[2] W. Jin, Z. Wang, Z. Yang, and S. Mou. Pontryagin Differentiable Programming: An End-to-End Learning and Control Framework. *Adv. Neural Inf. Process. Syst. (NeurIPS)*, 33:7048–7060, 2020.

[3] P. Donnelly-Kehoe, V. M. Saenger, N. Lisofsky, S. Kühn, M. L. Kringelbach, J. Schwarzbach, and G. Deco. Reliable local dynamics in the brain across sessions are revealed by whole-brain modeling of resting state activity. *Hum. Brain Mapp.*, 40(10):2967–2980, 2019.

[4] A. I. Luppi and E. A. Stamatakis. Combining network topology and information theory to construct representative brain networks. *Netw. Neurosci.*, 5(1):96–114, 2021.

[5] M. Mula, H. Coleman, and S. J. Wilson. Neuropsychiatric and cognitive comorbidities in epilepsy. *Continuum*, 28(2):457–482, 2022.

[6] J. A. Roa, K. A. Dhamsania, V. Olortegui Anaya, et al. Responsive neurostimulation of the thalamus for the treatment of refractory epilepsy. *Front. Hum. Neurosci.*, 16:863152, 2022.

[7] B. Q. Rosen and E. Halgren. A whole-cortex probabilistic diffusion tractography connectome. *eNeuro*, 8(1):ENEURO.0416-20.2020, 2021.

[8] M. M. Shanechi, J. J. Chemali, M. Liberman, K. Solt, and E. N. Brown. A brain-machine interface for control of medically induced coma. *PLoS Comput. Biol.*, 9(10):e1003284, 2013.

[9] A. K. Sharma, E. Rani, A. Waheed, and S. K. Rajput. Pharmacoresistant epilepsy: a current update on non-conventional pharmacological and non-pharmacological interventions. *J. Epilepsy Res.*, 5(1):1–8, 2015.

[10] M. Steriade. Impact of network activities on neuronal properties in corticothalamic systems. *J. Neurophysiol.*, 86(1):1–39, 2001.

[11] W. Xiong, E. S. Nurse, E. Lambert, M. J. Cook, and T. Kameneva. Seizure forecasting using long-term electroencephalography and electrocardiogram data. *Int. J. Neural Syst.*, 31(8):2150039, 2021.

[12] M. Scheffer, J. Bascompte, W. A. Brock, V. Brovkin, S. R. Carpenter, V. Dakos, H. Held, E. H. van Nes, M. Rietkerk, and G. Sugihara. Early-warning signals for critical transitions. *Nature*, 461(7260):53–59, 2009.

[13] V. Dakos, M. Scheffer, E. H. van Nes, V. Brovkin, V. Petoukhov, and H. Held. Methods for detecting early warnings of critical transitions in time series illustrated using simulated ecological data. *PLoS One*, 7(7):e41010, 2012.

[14] T. Wilkat, T. Rings, and K. Lehnertz. No evidence for critical slowing down prior to human epileptic seizures. *Chaos*, 29(9):093104, 2019.

[15] T. Rings, P. Mazrooei, T. Wilkat, and K. Lehnertz. Critical speeding up prior to bifurcations in complex systems. *Chaos*, 29(12):121105, 2019.

[16] P. Kwan and M. J. Brodie. Early identification of refractory epilepsy. *N. Engl. J. Med.*, 342(5):314–319, 2000.

[17] D. C. Van Essen, S. M. Smith, D. M. Barch, T. E. J. Behrens, E. Yacoub, K. Ugurbil, and the WU-Minn HCP Consortium. The WU-Minn Human Connectome Project: an overview. *NeuroImage*, 80:62–79, 2013.

[18] M. F. Glasser, T. S. Coalson, E. C. Robinson, C. D. Hacker, J. Harwell, E. Yacoub, K. Ugurbil, J. Andersson, C. F. Beckmann, M. Jenkinson, S. M. Smith, and D. C. Van Essen. A multi-modal parcellation of human cerebral cortex. *Nature*, 536(7615):171–178, 2016.

[19] G. K. Bergey, M. J. Morrell, E. M. Mizrahi, A. Goldman, D. King-Stephens, D. Nair, S. Srinivasan, B. Jobst, R. E. Gross, D. C. Shields, G. Barkley, V. Salanova, P. Morrell, and C. Faught. Long-term treatment with responsive brain stimulation in adults with refractory partial seizures. *Neurology*, 84(16):1667–1675, 2015.

[20] G. Deco, M. L. Kringelbach, V. K. Jirsa, and P. Ritter. The dynamics of resting fluctuations in the brain: metastability and its dynamical cortical core. *Sci. Rep.*, 7(1):3095, 2017.

[21] H. Kwakernaak and R. Sivan. *Linear Optimal Control Systems*. Wiley-Interscience, New York, 1972.

[22] L. S. Pontryagin, V. G. Boltyanskii, R. V. Gamkrelidze, and E. F. Mishchenko. *The Mathematical Theory of Optimal Processes*. Interscience, New York, 1962.

[23] T. Hu, H. Xie, M. Shan, Z. Yao, Q. Zhang, S. Guo, J. Jia, Y. Li, and C. Zhang. Long-term efficacy of anterior nucleus of the thalamus deep brain stimulation in drug-resistant epilepsy: a 5-year multicenter study. *BMC Med.*, 23:615, 2025.

[24] J. Wang, E. Niebur, J. Hu, and X. Li. Suppressing epileptic activity in a neural mass model using a closed-loop proportional-integral controller. *Sci. Rep.*, 6:27344, 2016.

[25] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada. Control barrier function based quadratic programs for safety critical systems. *IEEE Trans. Autom. Control*, 62(8):3861–3876, 2017.

[26] F. R. K. Chung. *Spectral Graph Theory*. CBMS Regional Conference Series in Mathematics, Vol. 92. American Mathematical Society, Providence, RI, 1997.

[27] S. Maslov and K. Sneppen. Specificity and stability in topology of protein networks. *Science*, 296(5569):910–913, 2002.

[28] P. E. Kloeden and E. Platen. *Numerical Solution of Stochastic Differential Equations*. Springer, Berlin, 1992.

[29] H. Risken. *The Fokker–Planck Equation: Methods of Solution and Applications*. 2nd ed. Springer, Berlin, 1989.

[30] S. H. Strogatz. *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering*. Westview Press, Cambridge, MA, 1994.

[31] S. Gu, F. Pasqualetti, M. Cieslak, Q. K. Telesford, A. B. Yu, A. E. Kahn, J. D. Medaglia, J. M. Vettel, M. B. Miller, S. T. Grafton, and D. S. Bassett. Controllability of structural brain networks. *Nat. Commun.*, 6:8414, 2015.

[32] D. J. Watts and S. H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442, 1998.

[33] E. Bullmore and O. Sporns. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat. Rev. Neurosci.*, 10(3):186–198, 2009.

[34] R. V. Shannon. A model of safe levels for electrical stimulation. *IEEE Trans. Biomed. Eng.*, 39(4):424–426, 1992.

[35] R. Rafailov, A. Sharma, E. Mitchell, S. Ermon, C. D. Manning, and C. Finn. Direct Preference Optimization: your language model is secretly a reward model. *Adv. Neural Inf. Process. Syst. (NeurIPS)*, 36, 2023.

[36] A. L. Benabid, S. Chabardes, J. Mitrofanis, and P. Pollak. Deep brain stimulation of the subthalamic nucleus for the treatment of Parkinson's disease. *Lancet Neurol.*, 8(1):67–81, 2009.

[37] S. Wiebe, W. T. Blume, J. P. Girvin, and M. Eliasziw. A randomized, controlled trial of surgery for temporal-lobe epilepsy. *N. Engl. J. Med.*, 345(5):311–318, 2001.

[38] World Health Organization. *Epilepsy: A Public Health Imperative*. WHO, Geneva, 2019.