

Structure-preserving Koopman model predictive control for closed-loop stabilisation of memristive neural dynamics

XuJiang Tang

txj_262538@163.com

Yangtze University College of Arts and Sciences

QiongLin Li

Yangtze University College of Arts and Sciences

Research Article

Keywords: Koopman operator, Model predictive control, Memristive Hindmarsh–Rose neuron, Closedloop neurostimulation, Lie derivative dictionary, Deep brain stimulation

Posted Date: April 15th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-9200134/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.
[Read Full License](#)

Additional Declarations: No competing interests reported.

Version of Record: A version of this preprint was published at Biological Cybernetics on August 11th, 2026. See the published version at <https://doi.org/10.1007/s00422-026-01054-4>.

Structure-preserving Koopman model predictive control for closed-loop stabilisation of memristive neural dynamics

XuJiang Tang^{†1}

QiongLin Li[†]

Received: date / Accepted: date

Abstract

Closed-loop neurostimulation for pathological neural oscillations requires controllers that are fast, constraint-aware, and robust to the strong nonlinearities inherent in excitable neural membranes. We develop a Koopman operator framework for model predictive control of a four-state memristive Hindmarsh–Rose neuron, a biophysical model that captures electromagnetic induction, coexisting attractors, and chaotic bursting. The key methodological contribution is a structure-preserving dictionary design based on iterated Lie derivatives of the governing vector field: instead of choosing basis functions ad hoc, the dictionary is constructed analytically from the polynomial structure of the neural dynamics, yielding a compact 13-dimensional observable space. A truncation error theorem quantifies the approximation quality of the finite-dimensional lifting, and a stability corollary connects dictionary fidelity to the closed-loop convergence neighbourhood. Extended dynamic mode decomposition with control identifies a lifted linear predictor that is embedded in a constrained receding-horizon controller with nonlinear re-lifting. A dictionary ablation study across four designs confirms that the physics-informed Lie-closure basis matches a generic mixed dictionary of six times its dimension. Monte Carlo benchmarks against five baseline controllers—including PID, LQR, sliding mode control, neural network MPC, and data-enabled predictive control—demonstrate the lowest tracking error (RMSE = 0.091), fastest settling (0.47 s), and 100% stabilisation rate under input saturation, with robustness confirmed under 20% parameter mismatch. These results establish a principled, computationally tractable route from biophysical neural models to real-time closed-loop control, with implications for the design of next-generation adaptive neurostimulation systems.

Keywords: Koopman operator · Model predictive control · Memristive Hindmarsh–Rose neuron · Closed-loop neurostimulation · Lie derivative dictionary · Deep brain stimulation

1 Introduction

The regulation of pathological neural activity is a central problem in computational neuroscience and neuro-engineering. Conditions such as Parkinson’s disease, essential tremor, and drug-resistant epilepsy are characterised by aberrant oscillatory patterns—sustained bursting, hypersynchrony, and bifurcation-mediated regime transitions—that disrupt normal information processing in neural circuits [Hammond et al., 2007, Uhlhaas and Singer, 2006]. Deep brain stimulation (DBS) has emerged as an effective clinical intervention, yet conventional open-loop DBS delivers continuous high-frequency pulses without adapting to the instantaneous neural state, leading to side effects, suboptimal energy use, and limited efficacy during state transitions [Benabid et al., 2009, Lozano et al., 2019]. Closed-loop or adaptive neurostimulation—where the stimulation waveform is computed in real time from neural recordings—has the potential to overcome these limitations, but requires control algorithms that are both fast enough for real-time deployment and rich enough to capture the nonlinear dynamics of excitable neural tissue [Holt et al., 2016, Popovych and Tass, 2012, Little et al., 2013].

From the perspective of control theory, a neural excitability model under stimulation defines a nonlinear, input-constrained dynamical system. Classical linear controllers (PID, LQR) operate on local linearisations that are valid only near a target equilibrium and break down during the large transients typical of bursting and chaotic regimes. Nonlinear approaches—sliding mode control, feedback linearisation, phase-amplitude reduction [Wilson and Moehlis, 2016, Monga et al., 2019]—improve robustness but either require full state access

[†]These authors contributed equally to this work.

that may not be available in neural interfaces or lack systematic constraint handling. Model predictive control (MPC) offers a principled framework for incorporating actuation limits, multi-step prediction, and trajectory optimisation, yet its effectiveness depends critically on the quality of the underlying dynamical model [Schiff, 2012].

Koopman operator theory provides an appealing bridge between nonlinear neural dynamics and linear prediction [Koopman, 1931, Mezić, 2005]. The Koopman operator acts on observable functions of the state and is linear by construction, even when the state dynamics are nonlinear. Applied Koopmanism [Budišić et al., 2012] and dynamic mode decomposition [Schmid, 2010, Tu et al., 2014] provide data-driven finite-dimensional approximations. Extended dynamic mode decomposition with control (EDMDc) [Williams et al., 2015, Proctor et al., 2016] constructs lifted linear predictors suitable for receding-horizon controllers [Korda and Mezić, 2018]. Recent advances include robust tube-based variants [Zhang et al., 2022], dissipativity-constrained learning [Hara et al., 2023], bilinear formulations [Goswami and Paley, 2022], and deep learning approaches [Lusch et al., 2018]. Despite this progress, applications of Koopman-based control to biophysical neuron models remain rare, and the design of Koopman dictionaries for neural systems is still largely ad hoc.

A critical and under-explored question is *how to choose the observable dictionary*. Generic basis functions—radial basis functions (RBFs), random Fourier features, or polynomials up to a fixed degree—guarantee universal approximation in the limit of large dictionaries, but in practice they introduce many irrelevant dimensions, inflate the condition number of the regression, and provide no interpretable link to the underlying biophysics [Li et al., 2017, Kaiser et al., 2021]. The Koopman generator (the infinitesimal operator of the semigroup) is the Lie derivative along the vector field [Mauroy and Mezić, 2016, Mauroy et al., 2020], and for polynomial vector fields the Lie derivative of a monomial yields another polynomial. This algebraic structure suggests a principled dictionary construction: close the seed observables under iterated Lie derivatives, truncate at a specified polynomial degree, and use the resulting basis as the Koopman dictionary.

In this paper, we develop and validate this idea for a biophysically detailed neural model. The memristive Hindmarsh–Rose (MHR) neuron augments the classical three-state Hindmarsh–Rose model [Hindmarsh and Rose, 1984] with a flux-controlled memristive element that emulates slow electromagnetic induction in neural tissue [Lv et al., 2016, Bao et al., 2020, Ma and Tang, 2017]. The resulting four-state system exhibits rich dynamics: coexisting limit cycles and chaotic attractors, mixed-mode oscillations, and bifurcation cascades that are hallmarks of real neural excitability [Njitacke et al., 2020, Parastesh et al., 2021]. As a testbed for closed-loop control, the MHR neuron bridges the gap between abstract control benchmarks and the complexity of real neural circuits, while remaining tractable for rigorous analysis.

The contributions of this work are as follows.

1. **Structure-preserving dictionary design.** We construct the Koopman dictionary from iterated Lie derivatives of the MHR vector field with polynomial degree truncation, yielding a 13-dimensional observable basis. A truncation error theorem (Theorem 1) provides a formal bound on the single-step approximation error as a function of dictionary depth and state magnitude.
2. **Dictionary-dependent stability guarantee.** A corollary (Corollary 1) links the truncation parameters to the closed-loop convergence neighbourhood, showing that enriching the dictionary shrinks the steady-state tracking bound. This result provides a quantitative design guideline: the practitioner can trade dictionary dimension against control precision.
3. **Dictionary ablation and spectral analysis.** Four dictionary designs are compared in both open-loop prediction and closed-loop control across 50 Monte Carlo trials. The compact Lie-closure dictionary matches a mixed design of six times its dimension, while a purely generic RBF dictionary of comparable size performs worst. Eigenvalue analysis of the identified Koopman matrices reveals that the Lie-closure basis captures the biophysical timescale hierarchy (fast membrane oscillation vs. slow recovery and flux relaxation), whereas the generic basis produces scattered spurious modes.
4. **Comprehensive controller benchmarking.** Six controllers—PID, local LQR, sliding mode control, neural network shooting MPC, data-enabled predictive control (DeePC), and the proposed Koopman MPC—are evaluated under identical conditions with statistical significance testing. The Koopman MPC

achieves the lowest tracking error, fastest settling, and 100% stabilisation rate, with robustness confirmed under 20% parameter mismatch. Basin-of-attraction scans show expansion from a fragmented multi-stable partition to a uniformly stabilised region. A synchronisation experiment demonstrates extension to trajectory tracking.

Beyond the methodological contributions, the results carry implications for the design of neurostimulation systems. The Lie derivative dictionary converts dictionary design from an empirical search into a systematic procedure rooted in the biophysics of the target neural circuit. The re-lifting mechanism and the truncation error bound together provide an interpretable quality metric for the Koopman predictor, which is absent in black-box neural network approaches. The per-step computation time (~ 31 ms) fits within realistic sampling budgets for implantable neurostimulators [Stanslaski et al., 2012], suggesting that operator-theoretic MPC is a viable candidate for next-generation adaptive DBS.

The paper is organised as follows. Section 2 provides the neuroscience background and motivates the choice of the MHR model. Section 3 defines the control task. Section 4 presents the Lie-closure dictionary and EDMDc identification. Section 5 describes the constrained Koopman MPC design with stability analysis. Section 6 details baseline controllers and protocol. Section 7 reports simulations and ablation results. Section 8 discusses findings in the context of neurostimulation. Section 9 concludes.

2 Neuroscience Background

This section situates the memristive Hindmarsh–Rose neuron within the landscape of biophysical neural models and discusses the clinical motivation for closed-loop control of excitable neural dynamics.

2.1 From Hodgkin–Huxley to memristive neuron models

The Hodgkin–Huxley formalism established the biophysical foundation of neural excitability by modelling voltage-gated ion channels as conductance variables obeying first-order kinetics [Hodgkin and Huxley, 1952]. Reduced models have since played a central role in computational neuroscience by preserving the essential qualitative features—spiking, bursting, adaptation, bistability—in lower-dimensional systems amenable to phase-plane and bifurcation analysis. The FitzHugh–Nagumo model [FitzHugh, 1961, Nagumo et al., 1962] captures the excitation–recovery dynamics in two variables; the Hindmarsh–Rose model [Hindmarsh and Rose, 1984] adds a slow adaptation variable that generates chaotic bursting; and the Izhikevich model [Izhikevich, 2003] achieves broad dynamical repertoire with minimal computational cost.

A more recent extension incorporates electromagnetic induction via a flux-controlled memristive element, yielding the memristive Hindmarsh–Rose (MHR) model [Lv et al., 2016, Bao et al., 2020]. The memductance function $W(\varphi) = \alpha + 3\beta\varphi^2$ couples the magnetic flux φ to the membrane current, modelling the effect of ion channel fluctuations under electromagnetic fields. This coupling enriches the bifurcation landscape: the MHR neuron exhibits coexisting periodic and chaotic attractors, hidden bursting firings, and antimonotonicity—features that have been confirmed experimentally in memristive circuit implementations [Ma and Tang, 2017, Njitacke et al., 2020]. For control purposes, the MHR model provides a challenging yet analytically tractable testbed whose polynomial vector field admits the Lie derivative analysis developed in this paper.

2.2 Closed-loop neurostimulation: clinical context

Deep brain stimulation has become the standard surgical therapy for medication-refractory Parkinson’s disease, targeting the subthalamic nucleus or the globus pallidus internus [Benabid et al., 2009, Lozano et al., 2019]. Conventional open-loop DBS delivers continuous high-frequency pulses (~ 130 Hz, 60 – 200 μ s pulse width) that modulate local and network activity through mechanisms that remain incompletely understood [McIntyre et al., 2004]. Although effective, this approach wastes energy during symptom-free intervals, may provoke speech and gait side effects, and cannot adapt to the patient’s changing state [Little et al., 2013].

Adaptive (closed-loop) DBS addresses these limitations by adjusting stimulation parameters based on biomarkers derived from local field potentials, typically the beta-band (13–30 Hz) power in the subthalamic

nucleus [Little et al., 2013, Stanslaski et al., 2012]. Clinical trials have demonstrated that adaptive DBS reduces both total stimulation energy and side effects while maintaining therapeutic efficacy [Little et al., 2016]. However, the control algorithms used in current devices remain simple—typically on–off or proportional controllers [Velisar et al., 2019]—and operate without an explicit dynamical model of the neural target. As the complexity of stimulation patterns increases (e.g., coordinated reset [Popovych and Tass, 2012, Tass, 2003] or phase-locked stimulation [Holt et al., 2016, Wilson and Moehlis, 2016]), there is a growing need for model-based control frameworks that can predict neural state evolution over a finite horizon, respect actuator constraints, and optimise stimulation waveforms in real time.

2.3 Koopman operators as neural system models

Koopman operator theory offers a natural modelling framework for neural control because it (i) transforms nonlinear dynamics into a linear representation in observable space, enabling the use of mature linear control tools; (ii) provides a spectral decomposition that can be related to the timescale hierarchy of neural oscillations [Brunton et al., 2016]; and (iii) accommodates data-driven identification from experimentally recorded neural time series. In the neuroscience context, the observables in a Koopman dictionary correspond to nonlinear combinations of the neural state (membrane potential, recovery variables, conductances) that evolve linearly—a perspective that resonates with the notion that neural computations are low-dimensional despite the high-dimensional state space [Gallego et al., 2017].

In this work, we exploit the fact that the MHR vector field is polynomial to obtain the Koopman dictionary analytically rather than from data. The resulting Lie-closure observables have direct biophysical meaning: they capture the cubic membrane nonlinearity, the quadratic recovery coupling, and the memristive flux interaction that together determine the neuron’s excitability. This physics-informed construction contrasts with black-box approaches (deep Koopman [Lusch et al., 2018], dictionary learning [Li et al., 2017]) in which the observables lack interpretability. For neurostimulation applications, interpretability is valuable because it allows clinicians to understand which dynamical interactions the controller exploits and to assess whether the model remains valid under parameter changes (e.g., disease progression or medication adjustments).

3 Problem Formulation

This section defines the controlled neural dynamics, the actuation constraints, and the quantitative metrics used throughout the paper. All variables are introduced once and reused consistently.

3.1 System model: memristive Hindmarsh–Rose neuron

Let $k \in \mathbb{Z}_{\geq 0}$ denote the discrete-time index with sampling period $\Delta t > 0$. The state is

$$x_k = [x_{1,k} \quad x_{2,k} \quad x_{3,k} \quad x_{4,k}]^\top \in \mathbb{R}^4,$$

where $x_{1,k}$ is the membrane potential, $x_{2,k}$ is the recovery variable (analogous to a slow potassium conductance), $x_{3,k}$ is the adaptation variable (modelling ultra-slow intracellular calcium or metabolic processes), and $x_{4,k}$ is the memristive flux state (representing the magnetic flux across the electromagnetic induction channel). The control input is scalar,

$$u_k \in \mathbb{R},$$

which represents the external stimulation current applied to the neuron—the biophysical equivalent of a charge-balanced current pulse delivered by a neurostimulator electrode.

The continuous-time dynamics are written in control-affine form

$$\dot{x}(t) = f(x(t); \theta) + g(x(t); \theta) u(t),$$

where θ collects the model parameters. For the memristive Hindmarsh–Rose neuron the component equations are [Hindmarsh and Rose, 1984, Lv et al., 2016]

$$\dot{x}_1 = x_2 - ax_1^3 + bx_1^2 - x_3 + I_{\text{ext}} - k_1 W(\varphi)x_1 + u, \quad (1)$$

$$\dot{x}_2 = c - dx_1^2 - x_2, \quad (2)$$

$$\dot{x}_3 = r(s(x_1 - x_r) - x_3), \quad (3)$$

$$\dot{\varphi} = x_1 - k_2 \varphi, \quad (4)$$

where $\varphi \equiv x_4$ is the magnetic flux across the memristive element and the memductance function is $W(\varphi) = \alpha + 3\beta\varphi^2$. Equation (1) governs the fast membrane dynamics: the cubic term $-ax_1^3 + bx_1^2$ produces the characteristic N-shaped current–voltage relation of excitable membranes, x_2 provides fast negative feedback (recovery), x_3 introduces slow adaptation responsible for burst termination, and the memductance coupling $k_1 W(\varphi)x_1$ modulates excitability through electromagnetic induction. The parameter I_{ext} represents a tonic external drive. The control input u enters only the membrane potential equation (1), reflecting the single-channel stimulation constraint typical of DBS electrodes. Nominal parameter values are listed in Appendix A.

The discrete-time model used for identification and control is obtained by fourth-order Runge–Kutta integration with period Δt ,

$$x_{k+1} = F(x_k, u_k; \theta).$$

The map F is treated as unknown in the data-driven modelling stage, while closed-loop simulations use the ground-truth discretisation.

3.2 Control objectives and constraints

The regulation target is a fixed equilibrium $x^* \in \mathbb{R}^4$ associated with a non-oscillatory (quiescent) regime—biophysically representing the suppression of pathological oscillations to a resting state. The control objective is to drive the closed-loop trajectory toward x^* with fast transient suppression of bursting in the membrane potential while respecting actuator limits.

Actuation is subject to hard saturation constraints,

$$u_k \in \mathcal{U}, \quad \mathcal{U} := \{u \in \mathbb{R} : |u| \leq u_{\max}\},$$

with a given bound $u_{\max} > 0$. This constraint reflects the safety limits on charge injection through implanted electrodes: exceeding the charge density threshold causes irreversible tissue damage [Shannon, 1992].

3.3 Performance metrics

Performance is evaluated on a finite post-control horizon $\{k_{\text{on}}, \dots, N\}$ with $N = \lfloor T/\Delta t \rfloor$. The primary metrics are: (i) the post-control root-mean-square tracking error $\text{RMSE}_x := [(N - k_{\text{on}} + 1)^{-1} \sum_{k=k_{\text{on}}}^N \|x_k - x^*\|_2^2]^{1/2}$; (ii) the settling time t_s , defined as the earliest time at which the trajectory enters and remains within an ε -neighbourhood of x^* , a metric motivated by the clinical need for rapid symptom suppression; (iii) the control energy $E_u := \sum_k u_k^2 \Delta t$, which relates to the total charge delivered and hence to tissue safety and battery life; and (iv) the efficiency ratio $\eta := \text{RMSE}_x / E_u$. A trial is counted as stabilised when the final membrane potential satisfies $|x_{1,N} - x_1^*| \leq \varepsilon_f$, and as converged when both the steady-state bias and residual variance over the last W samples fall below prescribed thresholds. All threshold values are recorded in Appendix B.

4 Koopman Operator and EDMDc

This section introduces the lifted linear predictor used for control. The Koopman lifting map is defined on the neuron state, the EDMDc regression is stated in matrix form, and all symbols required for multi-step prediction are fixed for later sections.

4.1 Koopman Operator Theory and Lifting

Consider the autonomous continuous-time system $\dot{x}(t) = f(x(t); \theta)$ on a state space $\mathcal{X} \subset \mathbb{R}^4$. Let $\Phi^t : \mathcal{X} \rightarrow \mathcal{X}$ denote its flow map so that $x(t) = \Phi^t(x(0))$. Let $\mathcal{F}$ be a linear space of scalar observables $\varphi : \mathcal{X} \rightarrow \mathbb{R}$, equipped with pointwise addition and scalar multiplication. The Koopman operator $\mathcal{K}^t : \mathcal{F} \rightarrow \mathcal{F}$ is defined by composition with the flow,

$$(\mathcal{K}^t \varphi)(x) = \varphi(\Phi^t(x)).$$

The operator $\mathcal{K}^t$ is linear even when the state dynamics are nonlinear.

For data-driven prediction and control, a finite collection of observables is selected to construct a lifting map. Let

$$\psi(x) = [\psi_1(x) \quad \psi_2(x) \quad \cdots \quad \psi_p(x)]^\top \in \mathbb{R}^p$$

be a vector of basis functions, and define the lifted state

$$z = \psi(x).$$

The lifted coordinates are designed to preserve the original state through a linear reconstruction map. A matrix $C \in \mathbb{R}^{4 \times p}$ is sought such that

$$x \approx Cz,$$

with C estimated from data. In the controlled setting, the lifting map is applied to measured states at each sampling instant, and the lifted linear model is used for prediction over a finite horizon in the MPC optimisation.

4.2 Structure-Preserving Dictionary via Lie Derivative Closure

A principal difficulty in Koopman-based modelling is the choice of the dictionary ψ . Generic basis functions such as radial basis functions guarantee universal approximation in the limit $p \rightarrow \infty$, yet they introduce no structural preference aligned with the dynamics and typically require large p for adequate accuracy [Li et al., 2017, Kaiser et al., 2021]. This subsection develops a systematic dictionary construction for polynomial control-affine systems by exploiting the algebraic structure of the Koopman generator through iterated Lie derivatives [Mauroy and Mezić, 2016, Mauroy et al., 2020]. Rigorous spectral convergence guarantees for data-driven Koopman approximations have been established in [Colbrook and Townsend, 2024]. The resulting dictionaries are compact, physics-informed, and equipped with a quantifiable approximation error bound.

4.2.1 Koopman generator and Lie derivatives

The infinitesimal generator of the Koopman semigroup for the autonomous flow $\dot{x} = f(x)$ is the Lie derivative operator [Mezić, 2005]

$$(\mathcal{L}_f \varphi)(x) = \nabla \varphi(x) \cdot f(x), \quad (5)$$

which is the directional derivative of φ along the vector field f . For a polynomial vector field f and a monomial observable φ , the Lie derivative $\mathcal{L}_f \varphi$ is itself a polynomial whose degree depends on the degrees of f and φ . If $\mathcal{L}_f \varphi$ can be expressed exactly as a linear combination of the observables $\psi_1, \dots, \psi_p$, then the Koopman matrix representation is exact on $\text{span}(\psi)$. Conversely, any monomial produced by $\mathcal{L}_f$ that lies outside $\text{span}(\psi)$ contributes to the truncation error of the finite-dimensional approximation.

4.2.2 Seed set and iterative closure for MHR dynamics

The MHR vector field (1)–(4) contains three nonlinear terms that enter the membrane potential equation: x_1^3, x_1^2 , and the memductance coupling $W(\varphi)x_1 = (\alpha + 3\beta\varphi^2)x_1$. In the recovery and adaptation equations, the only nonlinearity is x_1^2 . The flux equation is bilinear $(x_1 - k_2\varphi)$. Collecting these terms together with the state itself yields the seed set

$$\mathcal{M}_0 = \{x_1, x_2, x_3, \varphi\} \cup \{x_1^2, x_1^3, \varphi^2 x_1\}. \quad (6)$$

Definition 1 (m -th order Lie closure). Given a polynomial vector field $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and a finite seed set $\mathcal{M}_0$ of polynomial observables, define the iterated Lie closure of order m with degree truncation D as

$$\mathcal{D}_m^D = \Pi_{\leq D} \left(\text{span} \bigcup_{j=0}^m \mathcal{L}_f^j(\mathcal{M}_0) \right),$$

where $\mathcal{L}_f^0 \phi = \phi$, $\mathcal{L}_f^j \phi = \mathcal{L}_f(\mathcal{L}_f^{j-1} \phi)$, and $\Pi_{\leq D}$ retains only monomials of total degree at most D .

Applying one Lie derivative iteration to $\mathcal{M}_0$ with the MHR vector field and truncating at total degree $D = 3$ produces the extended set

$$\mathcal{D}_1^3 = \mathcal{M}_0 \cup \{x_1 x_2, x_1 x_3, \phi x_1, \phi^2, \phi x_1^2, x_1^2 x_2\}. \quad (7)$$

The derivation proceeds by computing the gradient of each element of $\mathcal{M}_0$ and evaluating the inner product with f . For example, $\mathcal{L}_f(x_1^2) = 2x_1 \dot{x}_1$ generates terms proportional to $x_1 x_2, x_1^4, x_1^3, x_1 x_3, \phi^2 x_1^2$, and x_1 . All terms of degree four and above are discarded by $\Pi_{\leq 3}$, while the remaining monomials are added to the dictionary. The memductance observable $\phi^2 x_1$ produces ϕx_1^2 and ϕ^2 through the flux dynamics $\dot{\phi} = x_1 - k_2 \phi$, capturing the electromagnetic coupling structure. The complete dictionary $\psi(x) \in \mathbb{R}^{13}$ is

$$\psi(x) = [x_1, x_2, x_3, \phi, x_1^2, x_1^3, \phi^2 x_1, x_1 x_2, x_1 x_3, \phi x_1, \phi^2, \phi x_1^2, x_1^2 x_2]^\top. \quad (8)$$

4.2.3 Truncation error bound

The closure procedure terminates because the degree truncation $\Pi_{\leq D}$ prevents unbounded growth. The monomials discarded at each iteration constitute the truncation residual. The following theorem quantifies this residual.

Theorem 1 (Truncation error for Lie-closure dictionaries). Let $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a polynomial vector field of degree δ and let $\psi = [\psi_1, \dots, \psi_p]^\top$ be the dictionary constructed from $\mathcal{D}_m^D$ (Definition 1). Denote the Galerkin projection of the Koopman generator onto $\text{span}(\psi)$ by the matrix $L \in \mathbb{R}^{p \times p}$ and let $K = e^{L\Delta t}$ be the corresponding discrete-time Koopman matrix. Then the single-step approximation error satisfies

$$\|\psi(F(x)) - K\psi(x)\|_2 \leq C_{D,m,\delta} \|x\|_2^{D+1} \Delta t \quad (9)$$

for all x in a compact set $\mathcal{X} \subset \mathbb{R}^n$ containing the origin, where $F(x) = x + f(x)\Delta t + \mathcal{O}(\Delta t^2)$ is the Euler-discretised flow and $C_{D,m,\delta} > 0$ depends on the coefficients of f , the truncation parameters (D, m) , and the degree δ of f , and is independent of x and Δt .

Proof. Write the continuous-time Koopman evolution of each observable as $\frac{d}{dt} \psi_i(x(t)) = (\mathcal{L}_f \psi_i)(x(t))$. Decompose the Lie derivative into its projection onto $\text{span}(\psi)$ and its orthogonal residual:

$$\mathcal{L}_f \psi_i(x) = \sum_{j=1}^p L_{ji} \psi_j(x) + r_i(x),$$

where $r_i(x)$ collects all monomials of degree exceeding D that arise from $\mathcal{L}_f \psi_i$ and are discarded by $\Pi_{\leq D}$. Since ψ_i has degree at most D and f has degree δ , the residual r_i consists of monomials of degree between $D+1$ and $D+\delta-1$. On the compact set $\mathcal{X}$, each such monomial x^α with $|\alpha| \geq D+1$ satisfies $|x^\alpha| \leq \|x\|_\infty^{|\alpha|} \leq \|x\|_2^{D+1} \cdot \text{diam}(\mathcal{X})^{|\alpha|-D-1}$, so there exists a constant $\bar{r} > 0$ (depending on the coefficients of f and the number of discarded monomials) such that $|r_i(x)| \leq \bar{r} \|x\|_2^{D+1}$.

At the Euler level,

$$\psi_i(F(x)) = \psi_i(x) + \Delta t \mathcal{L}_f \psi_i(x) + \mathcal{O}(\Delta t^2),$$

while

$$[K\psi(x)]_i = [e^{L\Delta t} \psi(x)]_i = \psi_i(x) + \Delta t [L\psi(x)]_i + \mathcal{O}(\Delta t^2).$$

The difference is

$$\psi_i(F(x)) - [K\psi(x)]_i = \Delta t r_i(x) + \mathcal{O}(\Delta t^2),$$

from which (9) follows by taking vector norms with $C_{D,m,\delta} = \sqrt{p} \bar{r} + \mathcal{O}(\Delta t)$. $\square$

Remark 1. For the MHR system, $\delta = 3$ (cubic nonlinearities in (1)). With $D = 3$ and $m = 1$, the dictionary $\mathcal{D}_1^3$ spans $p = 13$ observables and the residual consists of degree-4 and degree-5 monomials. On the operating set $\mathcal{X} = \{x : \|x - x^*\|_\infty \leq r\}$, the bound (9) scales as r^4 , so the approximation improves rapidly as trajectories approach the equilibrium. Increasing D to 4 or 5 would capture additional monomials and reduce the constant $C_{D,m,\delta}$, at the cost of enlarging p . This trade-off is quantified empirically in Section 7.3.

4.2.4 Hybrid dictionaries

The Lie-closure dictionary captures the algebraic structure of the dynamics and provides a quantifiable approximation error, yet its polynomial basis has limited capacity for representing spatially localised features far from the equilibrium. A hybrid design augments the Lie-closure monomials with a moderate number of radial basis functions to improve representation capacity in transient regions where the state deviates significantly from x^* . The hybrid dictionary takes the form

$$\psi_{\text{hyb}}(x) = \underbrace{\psi_{\text{Lie}}(x)}_{\mathcal{D}_1^3}, \underbrace{\phi_1(x), \dots, \phi_{n_r}(x)}_{\text{RBF}} \top \in \mathbb{R}^{p_{\text{Lie}} + n_r}, \quad (10)$$

where $\phi_j(x) = \exp(-\|x - c_j\|_2^2/w^2)$ with centres c_j selected by k -means clustering on the training data [Williams et al., 2015]. The Lie-closure monomials provide guaranteed local accuracy near the equilibrium (Theorem 1), while the RBFs extend coverage over the full operating envelope.

4.3 Extended Dynamic Mode Decomposition with Control (EDMDc)

Given the dictionary ψ and a dataset $\{(x_k, u_k)\}_{k=0}^N$, form lifted snapshot matrices Z_-, Z_+, U_- in the standard way. EDMDc fits $A \in \mathbb{R}^{p \times p}$ and $B \in \mathbb{R}^{p \times 1}$ by regularised least squares,

$$(A, B) = \arg \min_{A, B} \|Z_+ - AZ_- - BU_-\|_F^2 + \lambda \|A\|_F^2,$$

yielding a closed-form solution. By construction, the first four entries of $z_k = \psi(x_k)$ are the physical state, so the reconstruction map is $C = [\mathbf{I}_4 \mathbf{0}_{4 \times (p-4)}]$.

4.4 Multi-step Prediction and Validation

The open-loop i -step prediction follows $\hat{z}_{k+i+1|k} = A\hat{z}_{k+i|k} + Bu_{k+i|k}$, with $\hat{x}_{k+i|k} = C\hat{z}_{k+i|k}$. Closed-loop operation re-lifts from measured states at each sample, $\hat{z}_{k+i|k+i} = \psi(\tilde{x}_{k+i})$, which aligns with receding-horizon deployment where fresh measurements are always available.

5 Koopman-based Model Predictive Control

This section formulates the constrained MPC problem in the lifted space using the EDMDc predictor. The lifted state dynamics and reconstruction map are combined with quadratic objectives and hard input bounds, yielding a nonlinear shooting problem solved at each sampling instant via bounded quasi-Newton optimisation.

5.1 Lifted-space MPC Formulation

Let $z_k = \psi(x_k) \in \mathbb{R}^p$ be the lifted state and let the identified predictor be

$$z_{k+1} = Az_k + Bu_k, \quad x_k = Cz_k,$$

where $A \in \mathbb{R}^{p \times p}$ and $B \in \mathbb{R}^{p \times 1}$ are obtained by EDMDc and $C \in \mathbb{R}^{4 \times p}$ is the reconstruction map. By construction of the dictionary (Section 4.3), the first four entries of z_k are the physical state x_k itself, so $C = [\mathbf{I}_4 \mathbf{0}_{4 \times (p-4)}]$. The regulation target is $x^* \in \mathbb{R}^4$, and the cost function penalises deviation directly in the physical state space rather than requiring an explicit lifted target z^* .

Fix a prediction horizon $N_p \in \mathbb{Z}_{\geq 1}$. At time k , the decision variables are the input sequence

$$\mathbf{u}_k := [u_{k|k} \quad u_{k+1|k} \quad \cdots \quad u_{k+N_p-1|k}]^\top \in \mathbb{R}^{N_p},$$

and the corresponding predicted lifted states $\hat{z}_{k+i|k}$ for $i = 1, \dots, N_p$. The predicted dynamics follow a shooting scheme with intermediate re-lifting: at each predicted step, the lifted state is propagated by the linear predictor and the resulting physical-state prediction is re-lifted through the full dictionary. Concretely,

$$\begin{aligned} \hat{z}_{k+i+1|k} &= \mathbf{A} \hat{z}_{k+i|k} + \mathbf{B} u_{k+i|k}, \\ \hat{x}_{k+i+1|k} &= \mathbf{C} \hat{z}_{k+i+1|k}, \\ \hat{z}_{k+i+1|k} &\leftarrow \psi(\hat{x}_{k+i+1|k}), \end{aligned}$$

for $i = 0, \dots, N_p - 2$, with $\hat{z}_{k|k} = \psi(x_k)$. This re-lifting corrects the nonlinear basis functions at each horizon step and prevents error accumulation in the higher-order dictionary entries, yielding more accurate multi-step predictions than a purely linear rollout.

The MPC objective penalises state deviation and actuation effort. With symmetric positive semidefinite weights $\mathbf{Q} \in \mathbb{R}^{4 \times 4}$ and $\mathbf{Q}_f \in \mathbb{R}^{4 \times 4}$ and a scalar $R > 0$, the finite-horizon cost is

$$J_k(\mathbf{u}_k) = \sum_{i=0}^{N_p-1} (\|\hat{x}_{k+i|k} - x^*\|_{\mathbf{Q}}^2 + R u_{k+i|k}^2) + \|\hat{x}_{k+N_p|k} - x^*\|_{\mathbf{Q}_f}^2, \quad (11)$$

where $\|v\|_{\mathbf{Q}}^2 = v^\top \mathbf{Q} v$ and $\|u\|_R^2 = R u^2$.

Input saturation is enforced by the hard constraint

$$u_{k+i|k} \in \mathcal{U}, \quad \mathcal{U} := \{u \in \mathbb{R} : |u| \leq u_{\max}\},$$

for $i = 0, \dots, N_p - 1$. If state constraints are included, they are expressed as $\hat{x}_{k+i|k} \in \mathcal{X}$ for a prescribed admissible set $\mathcal{X}$.

The optimisation problem solved at time k is

$$\mathbf{u}_k^* = \arg \min_{\mathbf{u}_k} J_k(\mathbf{u}_k) \quad \text{s.t. dynamics \& constraints.} \quad (12)$$

This problem is a quadratic program when the predicted dynamics use a purely linear rollout. In the present implementation, the shooting scheme re-lifts the predicted state through the nonlinear dictionary, which modifies the cost landscape as described below.

Remark 2. Because the shooting scheme re-lifts the predicted state through the nonlinear dictionary at each horizon step, the mapping from the decision vector $\mathbf{u}_k$ to the total cost is no longer quadratic. The resulting nonlinear program is solved with the L-BFGS-B algorithm [Byrd et al., 1995], a bounded quasi-Newton method that exploits analytically supplied gradients. Each gradient evaluation propagates cost sensitivities backward through the shooting rollout, producing an exact gradient of the single-step objective at a cost linear in N_p . The bound constraints on u are handled natively by L-BFGS-B, so no auxiliary slack variables or active-set iterations are needed. Warm-starting from the shifted solution of the previous sampling instant reduces the typical iteration count to fewer than 15 per step.

5.2 Integral Correction

Finite-dimensional lifting can introduce steady offsets. An integral state η_k accumulates the membrane-potential error $e_k = x_{1,k} - x_1^*$ and enters the cost as an augmented penalty

$$J_k^{\text{aug}}(\mathbf{u}_k) = J_k(\mathbf{u}_k) + \sum_{i=0}^{N_p-1} \rho \eta_{k+i|k}^2, \quad (13)$$

with $\rho \geq 0$. The same ρ is used across controllers that support integral action.

Table 1: Koopman EDMDc MPC at sampling instant k .

Input: Measured state x_k ; EDMDc matrices $\mathbf{A}, \mathbf{B}$; dictionary ψ ; target x^* ; horizon N_p ; weights $\mathbf{Q}, \mathbf{R}, \mathbf{Q}_f$; bounds $u_{\max}$; integral state η_k	
1.	$z_k \leftarrow \psi(x_k)$ (lift from measurement)
2.	$e_k \leftarrow x_{1,k} - x_1^*$
3.	$\eta_k \leftarrow \eta_{k-1} + e_k \Delta t$
4.	Solve (12) via L-BFGS-B with re-lifting and initial condition $\hat{z}_{k k} = z_k$
5.	$u_k \leftarrow u_{k k}^* + k_i \eta_k$ (apply with integral correction)
Output: Control input u_k ; updated integral state η_k	

5.3 Algorithm Summary

At each sampling instant, the controller lifts the measured state $z_k = \psi(x_k)$, solves (12) via L-BFGS-B with re-lifting, and applies $u_k = u_{k|k}^* + k_i \eta_k$. The procedure is summarised in Table 1.

5.4 Stability and Feasibility Analysis

The closed-loop dynamics couple the nonlinear neuron model with a receding-horizon controller built on an approximate lifted predictor. The true lifted evolution includes an additive disturbance $d_k \in \mathbb{R}^p$ that captures regression error and lifting mismatch,

$$z_{k+1} = \mathbf{A}z_k + \mathbf{B}u_k + d_k. \quad (14)$$

Assumption 1. *There exists a constant $\bar{d} \geq 0$ such that $\|d_k\|_2 \leq \bar{d}$ for all $k \geq 0$ along the closed-loop trajectory within the operating region $\mathcal{X}$.*

Assumption 2. *There exists a local stabilising gain $K_{\text{loc}} \in \mathbb{R}^{1 \times p}$ for the nominal lifted system $(\mathbf{A}, \mathbf{B})$ under input saturation such that the terminal weight $\mathbf{Q}_f$ satisfies*

$$\mathbf{C}^\top \mathbf{Q}_f \mathbf{C} \succeq \mathbf{C}^\top \mathbf{Q} \mathbf{C} + K_{\text{loc}}^\top \mathbf{R} K_{\text{loc}}. \quad (15)$$

In practice, $\mathbf{Q}_f = \alpha \mathbf{Q}$ with $\alpha > 1$ is tuned to dominate the right-hand side over the operating region.

Proposition 1 (Practical stability under bounded re-lifting error). *Let Assumptions 1 and 2 hold and let the feasible set be nonempty for all $k \geq 0$. Then the closed-loop trajectory satisfies*

$$\limsup_{k \rightarrow \infty} \|x_k - x^*\|_2 \leq \gamma(\bar{d}),$$

where γ is a class- $\mathcal{K}$ function depending on $\mathbf{Q}, \mathbf{Q}_f, \mathbf{R}$ and the spectral properties of $\mathbf{A}$.

Corollary 1 (Dictionary quality determines the stability neighbourhood). *If the dictionary ψ is constructed from $\mathcal{D}_m^D$ (Definition 1), then by Theorem 1 the disturbance satisfies $\bar{d} \leq C_{D,m,\delta} \sup_{x \in \mathcal{X}} \|x\|_2^{D+1} \Delta t =: \bar{d}(D, m)$, and the stability neighbourhood scales as $\gamma(\bar{d}(D, m)) \rightarrow 0$ as $D \rightarrow \infty$ or as the operating radius shrinks.*

Proofs of Proposition 1 and Corollary 1 are given in Appendix E. In the experimental configuration, $\mathbf{Q}_f = 8\mathbf{Q}$ was verified to satisfy (15) over the operating region.

6 Baseline Controllers and Experimental Protocol

This section defines the baseline controllers and the evaluation protocol used for all comparisons. The same sampling period, actuation bounds, simulation horizon, and initial-condition sets are used across methods. Controller-specific parameters are tuned using a common validation set and then fixed for all reported tests.

6.1 Baseline Controllers

Five baselines share the same measured state x_k , saturation bound $|u_k| \leq u_{\max}$, and sampling period Δt . **PID** operates on the membrane potential error with gains (K_p, K_i, K_d) tuned on a validation set. **Local LQR** linearises the discretised dynamics at (x^*, u^*) and solves the discrete algebraic Riccati equation; the resulting state-feedback gain is applied with saturation. **Sliding mode control (SMC)** uses a scalar sliding surface on (x_1, x_2) with a smooth tanh approximation to the sign function. **Neural network shooting MPC** learns a one-step predictor $\hat{x}_{k+1} = \mathcal{N}(x_k, u_k)$ on the same training data used for Koopman identification, then evaluates 48 random input sequences per step and selects the lowest-cost candidate—a sampling-based strategy that trades exploration efficiency for implementation simplicity. **DeePC** [Coulson et al., 2019] constructs Hankel matrices from a persistently excited offline dataset and solves a regularised data-enabled predictive control problem on the scalar output $y_k = x_{1,k}$; the scalar formulation and cold-start transient partially explain the low stabilisation rate observed in Section 7. All tuning parameters are recorded in Appendix C.

6.2 Dataset Generation and Protocol

Training data are generated by simulating the MHR neuron with a persistently exciting random input ($|u_k| \leq u_{\max}$) over a single trajectory of 5000 steps. The same dataset is used for both Koopman EDMDC and the neural network predictor. Initial conditions for the 50-trial benchmark are drawn from $x_1 \in [\underline{x}_1, \bar{x}_1]$, $x_4 \in [\underline{x}_4, \bar{x}_4]$ with (x_2, x_3) at nominal values (bounds in Appendix B). For robustness tests, a 20% parameter mismatch is introduced while keeping controller parameters fixed. All metrics (Section 3.3) are computed with the same activation index k_{on} and horizon. Pairwise statistical significance is assessed via Wilcoxon rank-sum tests with Holm correction at $\alpha = 0.05$.

7 Results

This section reports model validation, closed-loop tremor suppression, controller benchmarking, basin scans, robustness under parameter mismatch, and synchronisation tests. All figures use the same notation defined in Sections 3–5. The image files are included from the `figures/` directory.

7.1 Koopman EDMDC Multi-step Prediction Validation

Figures 1 and 2 validate the EDMDC predictor. The 200-step open-loop test with re-lifting shows high fidelity for recovery and adaptation states; phase and bias errors accumulate at longer horizons in the strongly nonlinear membrane mode. The shorter horizon aligned with the MPC window confirms that prediction accuracy is sufficient for control purposes.

7.2 Closed-loop Tremor Suppression with Koopman MPC

Figures 3–5 show a representative trial. Koopman MPC drives the membrane potential into the ε -neighbourhood of x_1^* within ~ 0.5 s, saturating during the initial transient and relaxing below 0.8 V within three seconds. The phase portrait confirms monotonic contraction toward the target, in contrast to the sustained bursting of the uncontrolled dynamics.

7.3 Dictionary Ablation Study

This subsection evaluates the influence of dictionary composition on Koopman model quality and closed-loop MPC performance. Four dictionaries are compared: (i) GenericRBF, consisting of the state plus 80 radial basis functions ($p = 84$); (ii) LieD1, consisting of the state plus the nine Lie-closure monomials from $\mathcal{D}_1^3$ in (7) ($p = 13$); (iii) Hybrid, combining the Lie-closure monomials with 50 RBFs ($p = 63$); and (iv) CurrentMixed, the original design with five ad hoc polynomials plus 80 RBFs ($p = 89$). All four designs share the same training data, MPC tuning, and 50-trial evaluation protocol.




Figure 1: Koopman EDMDc 200-step prediction with re-lifting for three initial conditions (solid: ground truth; dashed: prediction).




Figure 2: Lifted-model prediction check over the MPC-aligned horizon.




Figure 3: Membrane potential under Koopman MPC (vertical marker: control activation).




Figure 4: Control input under saturation constraints.




Figure 5: Phase-space trajectories: uncontrolled vs. Koopman MPC.

Table 2 reports lifted dimension, spectral radius $\rho(\mathbf{A})$, condition number $\kappa(\mathbf{A})$, and per-sample EDMDc fitting residual for each dictionary. The LieD1 dictionary achieves a condition number of 4.91, four orders of magnitude lower than the 8.64×10^4 of CurrentMixed, confirming that the Lie-closure monomials form a well-conditioned basis for the MHR dynamics.

Table 2: Koopman model diagnostics for each dictionary design.

Dictionary	p	$\rho(\mathbf{A})$	$\kappa(\mathbf{A})$	Residual/ N
GenericRBF	84	1.0000	1.17×10^4	6.41×10^{-4}
LieD1	13	1.0001	4.91	7.62×10^{-3}
Hybrid	63	1.0001	7.40×10^3	6.85×10^{-3}
CurrentMixed	89	1.0000	8.64×10^4	2.81×10^{-3}

Figure 6 shows the mean prediction error over a 200-step horizon with re-lifting for each dictionary. The Lie-closure monomials provide accurate short-horizon predictions that align with the MPC operating window, while the hybrid dictionary maintains low error across the full horizon.




Figure 6: Multi-step prediction error (mean over 10 initial conditions) by dictionary type. The re-lifting interval matches the MPC sampling period.


Table 3 summarises closed-loop MPC performance. The most striking finding is that the compact LieD1 dictionary ($p = 13$) achieves an RMSE of 0.104 and a settling time of 0.67 s, rivalling the CurrentMixed design ($p = 89$, RMSE 0.090, $t_s = 0.52$ s) at one-seventh the lifted dimension. The GenericRBF dictionary, despite having $p = 84$, exhibits the worst tracking error (0.252) and is the only design with stabilisation failures (96%), demonstrating that large lifted dimension without structural guidance degrades rather than improves control performance. The Hybrid dictionary ($p = 63$) offers a balanced trade-off between accuracy and conditioning, closely matching CurrentMixed in all metrics while reducing the condition number by an order of magnitude.

Figure 7 visualises the trial-level distributions: GenericRBF shows heavy upper tails in both metrics, whereas the three physics-informed dictionaries cluster tightly.

Wilcoxon rank-sum tests confirm that the Hybrid dictionary significantly outperforms GenericRBF in RMSE ($p = 0.006$), while the differences between Hybrid, LieD1, and CurrentMixed are not statistically significant ($p > 0.12$). The lack of significant difference between LieD1 ($p = 13$) and CurrentMixed ($p = 89$) is itself a notable result: it demonstrates that the nine Lie-closure monomials capture the dominant nonlinear interactions required for MPC-relevant prediction, validating the truncation error analysis in Theorem 1.

Table 3: Dictionary ablation: closed-loop MPC performance (mean $\pm$ std, 50 trials). Best values in bold.

Dictionary	RMSE _x	t_s (s)	E_u	η	Stab. %
GenericRBF	0.252 $\pm$ 0.242	5.49 $\pm$ 7.67	64.7 $\pm$ 61.9	0.00665	96
LieD1	0.104 $\pm$ 0.075	0.67 $\pm$ 0.69	67.0 $\pm$ 53.5	0.00197	100
Hybrid	0.093 $\pm$ 0.075	0.68 $\pm$ 0.86	66.8 $\pm$ 52.8	0.00178	100
CurrentMixed	0.090$\pm$0.074	0.52$\pm$0.57	66.3$\pm$51.9	0.00171	100

Figure 7: Box plots of RMSE_x and control efficiency η across 50 trials for each dictionary design.

The GenericRBF dictionary, which lacks any physics-informed observables, stands as a clear outlier in both accuracy and reliability.

7.4 Koopman Spectral Analysis

To connect the empirical performance differences to the operator-theoretic foundations, we examine the eigenvalue spectra of the EDMDC matrices $\mathbf{A}$ obtained from each dictionary. Figure 8 plots the eigenvalues $\{\lambda_j\}$ of $\mathbf{A}$ in the complex plane for LieD1 ($p = 13$) and GenericRBF ($p = 84$).

The LieD1 spectrum exhibits a clear bimodal structure: (i) a pair of complex-conjugate eigenvalues with argument $\theta \approx \pm 0.30$ rad and $|\lambda| \approx 0.979$, corresponding to the fast membrane–recovery oscillation at a period of ~ 1.0 s ($\approx 2\pi\Delta t/\theta$); and (ii) a cluster of real eigenvalues in $[0.93, 1.00]$, corresponding to the slow adaptation (x_3) and flux (φ) relaxation dynamics. This separation mirrors the biophysical timescale hierarchy of the MHR neuron: the membrane potential oscillates on a fast time-scale set by the cubic x_1 – x_2 coupling, while recovery and memristive variables evolve an order of magnitude more slowly. Because the Lie-closure monomials are algebraically aligned with these couplings, the 13 eigenvalues faithfully represent the relevant dynamical modes, with no spurious eigenvalues outside or far inside the unit circle.

By contrast, the GenericRBF spectrum contains 84 eigenvalues scattered broadly inside the unit disk. While a few eigenvalues coincide approximately with the LieD1 modes, the majority have $|\lambda_j| \in [0.3, 0.9]$ and carry no identifiable physical meaning—they arise from the high-dimensional RBF basis fitting noise and local data structure rather than the Koopman generator. The presence of near-zero eigenvalues inflates the condition number $\kappa(\mathbf{A}) = 1.17 \times 10^4$ and amplifies regression sensitivity, explaining both the poor prediction accuracy and the stabilisation failures observed in the ablation study.

Remark 3. The spectral radius $\rho(\mathbf{A}) \approx 1.000$ for all dictionaries (Table 2) is consistent with the marginally stable oscillatory dynamics of the uncontrolled MHR neuron. The oscillatory eigenvalue pair of LieD1 can be identified with the Koopman eigenfunction associated with the limit-cycle frequency of the autonomous flow, a direct manifestation of the Koopman spectral decomposition [Mezić, 2005]. That this mode emerges from only 13 observables—constructed without any spectral computation—validates the Lie-closure strategy as a




Figure 8: Eigenvalue spectra of the identified Koopman operator $\mathbf{A}$ for LieD1 (left) and GenericRBF (right). The unit circle is shown in grey. Marker colour encodes the magnitude $|\lambda_j|$.

physics-based alternative to data-driven spectral methods such as Hankel-DMD for capturing the dominant Koopman modes of excitable neural dynamics.

7.5 Benchmarking Against Baseline Controllers

Figure 9 summarises quantitative comparisons. Koopman MPC achieves the lowest mean RMSE (0.091 ± 0.074), followed by LQR (0.141 ± 0.121) and SMC (0.190 ± 0.218); PID (0.519), NN shooting (0.456), and DeePC (0.888) are substantially worse. All five pairwise improvements are significant (Wilcoxon, $p < 0.014$). Settling time is 0.47 ± 0.50 s for Koopman MPC vs. 1.20 s (LQR) and > 12 s (PID, NN). Control energy (66.3) is comparable to the classical baselines and an order of magnitude below NN (193.2) and DeePC (456.4). The efficiency ratio $\eta = \text{RMSE}_x/E_u$ ranks Koopman MPC first (0.00137).

Figure 10 reports stabilisation and convergence rates. Koopman MPC and LQR both achieve 100%/100%; SMC stabilises 86%; PID 64%; NN 50%; DeePC 28% stabilisation/50% convergence.

7.6 Representative Trajectory Comparison

Figure 11 compares membrane potential and control input for a fixed initial condition. Koopman MPC and LQR suppress the burst within ~ 1 s; SMC follows with higher switching activity; PID settles after ~ 10 s; NN shooting oscillates irregularly; and DeePC fails to converge.

7.7 Basin Scans and Robustness

Figure 12 contrasts open-loop and closed-loop basin scans over a grid in the (x_1, x_4) plane. The uncontrolled scan reveals a fragmented multistable partition; under Koopman MPC, the entire grid converges, confirming expansion from a fragmented basin to a uniformly stabilised region.

Figure 13 reports robustness under 20% parameter mismatch. Koopman MPC retains suppression ability with a small residual offset consistent with the practical stability bound (Proposition 1), whereas local LQR exhibits larger overshoot and slower convergence due to the shifted linearisation point.




Figure 9: Benchmark metrics across controllers (error bars: ± 1 std).




Figure 10: Stabilisation and convergence rates over $\mathcal{X}_0$.




Figure 11: Representative trajectory comparison (top: x_1 ; bottom: u ; vertical marker: activation).




Figure 12: Basin scans in (x_1, x_4) : uncontrolled (left) vs. Koopman MPC (right).




Figure 13: Robustness: Koopman MPC vs. LQR under 20% mismatch.

7.8 Synchronisation Experiments

Figure 14 reports master–slave synchronisation. The error norm drops below 10^{-2} within ~ 1 s of activation, demonstrating that the same Koopman MPC framework extends to trajectory tracking by replacing the fixed target x^* with the time-varying master state.




Figure 14: Synchronisation: error norm (left) and phase-space trajectories (right).

8 Discussion

8.1 Transient suppression and the role of lifting

The MHR neuron combines fast oscillatory membrane dynamics, slow recovery, and memristive coupling—features that make a Jacobian-based model valid only in a small neighbourhood of x^* . The Koopman EDMDC predictor propagates a linear recursion in the 13-dimensional lifted space, which captures the dominant nonlinear interactions over the MPC horizon while re-lifting from measured states at each sample corrects accumulated drift. The result is the shortest settling time (0.47 ± 0.50 s) among all tested controllers, with peak input approaching the saturation bound during the initial transient and relaxing thereafter—behaviour consistent with finite-horizon constrained optimisation.

8.2 Dictionary design: theory meets biophysics

The ablation study confirms the predictions of the truncation error bound (Theorem 1). The compact LieD1 dictionary ($p = 13$) achieves a condition number of 4.91, four orders of magnitude below CurrentMixed ($p = 89$, $\kappa = 8.64 \times 10^4$), yet its RMSE and settling time are statistically indistinguishable ($p > 0.12$, Wilcoxon). GenericRBF ($p = 84$), which contains no physics-informed observables, is the only design with stabilisation failures (96%), demonstrating that large lifted dimension without structural guidance degrades performance. The Lie derivative framework therefore converts dictionary design from an empirical search into a systematic procedure guided by the vector field.

The biophysical significance of this result merits emphasis. The nine Lie-closure monomials added to the state in equation (7)— x_1x_2 , x_1x_3 , ϕx_1 , ϕ^2 , ϕx_1^2 , $x_1^2x_2$ —correspond to specific nonlinear interaction channels in the neural dynamics: x_1x_2 captures the fast excitation–recovery coupling, ϕx_1 and ϕ^2 encode the electromagnetic induction pathway, and x_1x_3 represents the slow adaptation modulation of the membrane. That these nine observables suffice for accurate MPC suggests that the neuron’s controllability structure is dominated by a small set of biophysically meaningful mode interactions, not by the full combinatorial space of nonlinear terms.

8.3 Short-horizon prediction anomaly of LieD1

An apparent anomaly (Fig. 6) is that LieD1 shows the highest short-horizon prediction error despite strong closed-loop performance. Because $\mathcal{D}_1^3$ retains only degree- ≤ 3 monomials, the truncation residual consists of quartic and quintic terms that dominate when $\|x\|$ is large. As the horizon extends and re-lifting resets the lifted state, the effective $\|x\|$ decreases and the error decays. In the MPC loop, re-lifting at every sample and tolerance of moderate prediction bias—provided the cost gradient points toward the target—prevent the short-horizon deficit from degrading regulation quality.

8.4 Spectral structure and timescale hierarchy

The eigenvalue spectrum of the identified Koopman operator A (Fig. 8) provides a direct connection between the operator-theoretic framework and the biophysics of neural excitability. The LieD1 spectrum exhibits a clear bimodal structure: (i) a pair of complex-conjugate eigenvalues with argument $\theta \approx \pm 0.30$ rad and $|\lambda| \approx 0.979$, corresponding to the fast membrane–recovery oscillation at a period of ~ 1.0 s; and (ii) a cluster of real eigenvalues in $[0.93, 1.00]$, corresponding to the slow adaptation (x_3) and flux (ϕ) relaxation. This separation mirrors the biophysical timescale hierarchy of the MHR neuron: the membrane potential oscillates on a fast time-scale set by the cubic x_1 – x_2 coupling, while recovery and memristive variables evolve an order of magnitude more slowly. That these modes emerge from only 13 observables—constructed without any spectral computation, purely from the algebraic structure of the vector field—validates the Lie-closure strategy as a physics-based alternative to data-driven spectral methods.

By contrast, the GenericRBF spectrum contains 84 eigenvalues scattered broadly inside the unit disk. The majority carry no identifiable physical meaning and arise from the high-dimensional RBF basis fitting noise and local data structure. This spectral clutter explains both the poor prediction accuracy and the stabilisation failures: the controller must contend with spurious dynamical modes that do not correspond to any real neural process.

8.5 Implications for closed-loop neurostimulation

The results of this study, while obtained on a computational model, carry several implications for the design of adaptive neurostimulation systems.

From model neurons to neural circuits. The MHR neuron is a single-cell model, whereas clinical DBS targets populations of neurons in subcortical nuclei. Nevertheless, the mean-field dynamics of coupled neural populations often reduce to low-dimensional systems with polynomial nonlinearities and timescale separation [Wilson and Moehlis, 2016]—precisely the structural features exploited by the Lie-closure dictionary. Extending the present framework to mean-field or neural mass models [Deco et al., 2008] is a natural and tractable next step: the Lie derivative construction applies to any polynomial vector field regardless of state dimension, and the truncation error bound generalises directly.

Interpretability and clinical trust. A practical obstacle for deploying advanced controllers in neurostimulation devices is the lack of interpretability: clinicians must understand what the controller is doing and why. The Koopman dictionary provides a transparent intermediate representation—each observable is a named polynomial function of the neural state, and the Koopman matrix A encodes how these observables evolve. Changes in the eigenvalue spectrum under disease progression or medication adjustment can be monitored as biomarkers of model validity. This level of interpretability is absent in deep learning-based controllers and represents a practical advantage for regulatory approval and clinical adoption.

Computational feasibility for implantable devices. The mean per-step computation time of ~ 31 ms fits within the 50 ms sampling budget of current investigational neurostimulators [Stanslaski et al., 2012]. Modern system-on-chip designs for implantable devices provide sufficient floating-point capability for the matrix–vector operations required by the lifted linear predictor. The compact 13-dimensional LieD1 dictionary is

particularly attractive in this context: it reduces the memory footprint and computational load by a factor of ~ 7 compared to the 89-dimensional mixed design, with no statistically significant loss in control quality.

Charge-balanced stimulation and safety. The input saturation constraint $|u| \leq u_{\max}$ enforced by the MPC formulation provides a first layer of safety. In practice, neurostimulation waveforms must be charge-balanced to prevent electrolysis and tissue damage. The MPC framework can accommodate charge-balance constraints by augmenting the optimisation with equality constraints on the time-integrated control input [Shannon, 1992]. This extension is straightforward within the L-BFGS-B formulation and does not alter the dictionary design.

Robustness as a clinical requirement. Neural dynamics are inherently non-stationary: electrode impedance drifts over months, disease progression alters network properties, and medication cycles introduce parameter variation on the timescale of hours. The 20% parameter mismatch test in Section 7.7 confirms that the Koopman MPC retains stable convergence under substantial modelling error—a prerequisite for any clinically deployed controller. Periodic re-identification of the EDMDC model from logged stimulation–response data could extend this robustness to long-term drift, with the Lie-closure dictionary providing a fixed structural template that only requires re-fitting the matrix entries (A, B) .

8.6 Limitations and future directions

Several limitations should be acknowledged. First, the MHR neuron is a phenomenological model; the memristive element is an abstraction of electromagnetic induction whose quantitative parameters have not been calibrated to specific experimental preparations. Second, the single-neuron setting does not capture the network-level dynamics (synchronisation, propagation delays, heterogeneity) that are central to DBS efficacy. Third, the Lie-closure framework requires explicit knowledge of the vector field; for data-driven settings where the governing equations are unknown, hybrid approaches combining Lie-inspired polynomial dictionaries with data-driven refinement [Kaiser et al., 2021] may be necessary. Fourth, the current study uses a fixed-parameter Koopman model; online adaptation of the EDMDC matrices to changing neural dynamics is an important direction for long-term deployment.

Future work should address these limitations by (i) validating the framework on biophysically calibrated multi-compartment neuron models and neural population models; (ii) incorporating output feedback (partial observability) to reflect the limited measurements available from DBS electrodes; (iii) extending the MPC formulation to include charge-balance and pulse-shape constraints; and (iv) investigating online re-identification strategies that exploit the fixed dictionary structure for efficient model updates.

9 Conclusion

This paper presented a structure-preserving Koopman operator framework for real-time model predictive control of a biophysical neuron model exhibiting chaotic excitability, with the goal of informing the design of closed-loop neurostimulation systems. Four contributions were established.

First, a Lie derivative closure procedure was developed that systematically generates physics-informed Koopman observables from the governing vector field of the memristive Hindmarsh–Rose neuron. The resulting 13-dimensional dictionary captures the dominant nonlinear interactions—cubic membrane dynamics, excitation–recovery coupling, and electromagnetic induction—and is equipped with a formal truncation error bound (Theorem 1) that links dictionary depth to approximation quality.

Second, a dictionary-dependent stability result (Corollary 1) connects the Koopman predictor truncation error to the closed-loop convergence neighbourhood, providing a quantitative design guideline: enriching the dictionary shrinks the steady-state tracking bound.

Third, a dictionary ablation study confirmed that the compact Lie-closure basis ($p = 13$) achieves control performance statistically indistinguishable from a mixed design of six times its dimension ($p = 89$), while a purely generic RBF dictionary of comparable size performed worst with the only stabilisation failures. Spectral

analysis revealed that the Lie-closure basis captures the biophysical timescale hierarchy of the neuron, whereas the generic basis produces scattered spurious modes.

Fourth, a six-controller benchmark across 50 Monte Carlo trials demonstrated the lowest tracking error (RMSE = 0.091), fastest settling time (0.47 s), and 100% stabilisation rate under input saturation, with robustness confirmed under 20% parameter mismatch.

These results demonstrate that the Koopman operator approach provides a principled, interpretable, and computationally tractable route from biophysical neural models to constrained closed-loop control. The Lie-closure dictionary converts dictionary design from an empirical search into a systematic procedure rooted in the structure of the neural dynamics, and the per-step computation (~ 31 ms) is compatible with the timing requirements of current implantable neurostimulation devices. Future work will extend this framework to neural population models, incorporate output feedback under partial observability, and investigate online model adaptation for long-term clinical deployment.

Acknowledgements

This work was supported by the Jingzhou Municipal Science and Technology Bureau (Grant No. 2023LHX08).

Author contributions

QL conceived the neural model analysis and performed the MATLAB simulations. XT designed the Koopman dictionary framework, developed the Python benchmark suite, proved the theoretical results, and wrote the manuscript. Both authors reviewed and approved the final version.

Data availability

The MATLAB and Python source code used to generate all results are available at <https://github.com/TXJ2006/Koopman-MHR-MPC>.

Conflict of interest

The authors declare no competing interests.

A MHR Model Parameters

This appendix records the parameters used for the memristive Hindmarsh–Rose (MHR) neuron and the numerical settings used to generate data and evaluate closed-loop performance. The continuous-time model is written in control-affine form

$$\dot{x}(t) = f(x(t); \theta) + g(x(t); \theta) u(t), \quad x = [x_1, x_2, x_3, x_4]^\top,$$

where x_1 is membrane potential, x_2 is recovery, x_3 is adaptation, and x_4 is the memristive flux state.

The simulations use a fixed sampling period Δt and numerical integration inside each sampling interval. The discrete-time map is

$$x_{k+1} = F(x_k, u_k; \theta),$$

obtained by integrating the continuous-time dynamics on $[k\Delta t, (k+1)\Delta t]$.

Table 4 reports the parameter vector θ and the operating point (x^*, u^*) used in the regulation experiments. The equilibrium is computed by solving $f(x^*; \theta) + g(x^*; \theta)u^* = 0$ under $u^* = 0$ unless stated otherwise.

For robustness tests, the mismatched parameter vector is denoted by $\tilde{\theta}$ and satisfies

$$\frac{\|\tilde{\theta} - \theta\|_2}{\|\theta\|_2} = 0.2.$$

Table 4: MHR model parameters and numerical settings used in experiments.

Symbol	Value
Δt	0.05 s
Integrator	Fourth-order Runge–Kutta (RK4), fixed step Δt
$u_{\max}$	5.0
a, b, c, d	1, 3, 1, 5
r, s, x_r	0.006, 4, -1.6
I	3.2 (external stimulus)
k	0.5 (memristive coupling strength)
α, β	0.1, 0.02 (memductance parameters)
k_2	0.5 (flux decay)
EM forcing	$0.35 \sin(0.35t)$ (sinusoidal)
x^*	$[-0.686, -1.356, 3.654, -1.373]^\top$
u^*	0

All controller parameters are fixed at their nominal values when evaluating robustness.

B Shared Protocol Parameters

Table 5 records the simulation and evaluation parameters that are common to all controllers throughout the benchmark. These values ensure that comparisons are conducted under identical operating conditions.

Table 5: Shared protocol parameters used across all controllers and trials.

Parameter	Value
Sampling period Δt	0.05 s
Control activation index k_{on}	100 (activation at $t = 5$ s)
Simulation horizon T_{sim}	500 steps (25 s total)
Number of Monte Carlo trials	50
Input saturation bound $u_{\max}$	5.0
IC range: x_1	$[-2.5, 1.5]$
IC range: x_4	$[-3.0, 1.0]$
IC fixed: (x_2, x_3)	$(-1.356, 3.654)$
Settling tolerance ε	0.1
Settling window	Last 100 steps (5 s)
Convergence bias threshold	0.05
Convergence variance threshold	0.01

Table 6 reports the mean per-step computation time for each controller, measured on a single core using the Python benchmark implementation. The timing excludes data collection and initialisation overhead, isolating the online control computation.

Remark 4. *The Koopman MPC step time is dominated by the L-BFGS-B optimisation, which involves repeated dictionary evaluations and gradient back-propagation through the shooting rollout. The large standard deviation reflects the variation in iteration counts across time steps: early transient steps require more iterations due to aggressive correction, while steady-state steps converge in fewer than five iterations with warm-starting. Classical controllers (PID, LQR, SMC) evaluate closed-form expressions and are orders of magnitude faster. DeePC and neural network shooting occupy an intermediate range. For neurostimulation applications with sampling periods of 50 ms, the mean Koopman MPC computation of approximately 31 ms remains within the real-time budget, though the tail of the distribution may require iteration-count limits in deployment.*

Table 6: Per-step computation time for each controller (mean $\pm$ std over 200 steps).

Controller	Time per step (ms)
PID	0.004 ± 0.001
Local LQR	0.005 ± 0.001
SMC	0.006 ± 0.002
DeePC	0.373 ± 0.061
NN shooting	0.947 ± 0.109
Koopman MPC	30.667 ± 48.717

C Controller Tuning Parameters

This appendix records the tuning choices used across controllers and clarifies which parameters are shared to support fair comparison. The same sampling period Δt , control activation time index k_{on} , evaluation horizon N , and saturation bound u_{max} are used throughout.

C.1 Koopman EDMDc identification settings

Let $\psi : \mathbb{R}^4 \rightarrow \mathbb{R}^p$ be the lifting map used in EDMDc. The dictionary is chosen to balance prediction fidelity and numerical conditioning. The main benchmark results in Sections 7.5–7.7 use the CurrentMixed dictionary ($p = 89$), which provides the lowest tracking error. The dictionary ablation study in Section 7.3 compares this design against three alternatives derived from the Lie-closure theory in Section 4.2.

Table 7 records the lifting and regression settings for the CurrentMixed dictionary used in the main benchmark.

Table 7: EDMDc identification settings for the Koopman-lifted predictor (CurrentMixed dictionary).

Item	Setting
Lift dimension p	89 (= 4 + 5 polynomial + 80 RBF)
Dictionary ψ	$[x; x_1^2, x_1^3, \varphi^2 x_1, x_1 x_2, \varphi^2; 80 \text{ RBF}]$
Regularisation λ	pseudo-inverse (pinv, implicit Tikhonov)
Training trajectories	1 trajectory, $T = 5000$ steps (250 s)
Input excitation	uniform random $u \in [-1.5, 1.5]$ i.i.d. per step
RBF centres	80 centres via k -means on training states
RBF width σ	2.0
Reconstruction map	First 4 rows of ψ (identity embedding)

C.2 Koopman MPC tuning

The MPC design is defined by the horizon N_p , weights (Q, Q_f, R) , and the optional integral penalty ρ in Section 5.2. The horizon length is selected to match the prediction fidelity of the Koopman model under re-lifting and to respect real-time computation constraints.

Table 8 records the MPC tuning values.

C.3 Baseline controller tuning

PID, LQR, SMC, neural network shooting MPC, and DeePC tuning values are recorded in Table 9. Each method is tuned on the same validation set with the same saturation bound.

Table 8: Koopman MPC tuning parameters.

Item	Setting
Prediction horizon N_p	12
State weight Q	$\text{diag}(25, 5, 5, 3)$ on physical states
Terminal weight Q_f	$8Q = \text{diag}(200, 40, 40, 24)$
Input weight R	0.01
Integral gain k_i (on x_1)	0.008
Solver	L-BFGS-B (shooting with re-lifting)
Solver tolerance	$\text{ftol} = 10^{-10}$, $\text{maxiter} = 60$
Warm start	Enabled (shifted previous solution)

Table 9: Baseline controller tuning parameters used in experiments.

Controller	Parameter	Value
PID	K_p	1.0
	K_i	0.015
	K_d	0.2
Local LQR	Q_ℓ	$\text{diag}(10, 2, 2, 1)$
	R_ℓ	0.1
SMC	k_s (switching gain)	2.5
	k_{lin} (linear gain)	0.5
	ε (boundary layer)	0.08
NN shooting MPC	N_p	6
	Q_{nn}	$\text{diag}(10, 2, 2, 1)$
	R_{nn}	0.1
	Samples per step	48 (random shooting)
DeePC	T_{ini}	20
	N_p	8
	λ_g (regularisation)	1.0
	λ_u (input regularisation)	0.1

D Extended Statistical Tables and Koopman-specific Analyses

Tables 10–14 provide extended statistics for the benchmark metrics.

Table 10: Extended statistics for post-control tracking error RMSE_x .

Controller	Mean	Std	Median	Trials
PID	0.5188	0.1526	0.5576	50
Local LQR	0.1410	0.1210	0.1072	50
SMC	0.1899	0.2175	0.1241	50
NN shooting	0.4558	0.2830	0.4084	50
DeePC	0.8880	0.3362	0.8990	50
Koopman MPC	0.0907	0.0740	0.0659	50

Table 11: Extended statistics for settling time t_s computed on converged trials.

Controller	Mean	Std	Median	Converged trials	Total trials
PID	12.69	3.43	12.00	30	50
Local LQR	1.20	0.83	1.00	50	50
SMC	1.43	1.88	0.35	42	50
NN shooting	19.06	1.32	19.70	18	50
DeePC	19.85	0.14	19.93	6	50
Koopman MPC	0.47	0.50	0.35	50	50

Table 12: Extended statistics for control energy E_u and peak input u_{peak} .

Controller	E_u mean	E_u std	u_{peak} mean	u_{peak} std
PID	59.03	47.07	2.048	0.736
Local LQR	69.18	52.99	4.804	0.602
SMC	58.34	48.30	2.867	0.465
NN shooting	193.19	104.72	4.968	0.047
DeePC	456.44	17.29	5.000	0.000
Koopman MPC	66.31	51.72	4.566	0.769

Table 13: Extended statistics for control efficiency ratio $\eta = \text{RMSE}_x/E_u$ (lower is better).

Controller	η (mean)	Rank
PID	0.00879	6
Local LQR	0.00204	3
SMC	0.00326	5
NN shooting	0.00236	4
DeePC	0.00195	2
Koopman MPC	0.00137	1

D.1 Dictionary ablation: extended statistical results

Table 15 reports Wilcoxon rank-sum test results comparing the Hybrid dictionary against each alternative design. The test statistic and p -value are computed on the RMSE_x metric across 50 trials.

Table 14: Extended statistics for stabilisation and convergence rates.

Controller	Stabilisation rate (%)	Convergence rate (%)
PID	64	100
Local LQR	100	100
SMC	86	100
NN shooting	50	100
DeePC	28	50
Koopman MPC	100	100

Table 15: Wilcoxon rank-sum tests: Hybrid dictionary vs. alternatives (RMSE_x, 50 trials). Significant at $\alpha = 0.05$ marked with †.

Comparison	p -value	Significant
Hybrid vs. GenericRBF	0.006	†
Hybrid vs. LieD1	0.130	
Hybrid vs. CurrentMixed	0.607	

Table 16 provides per-dictionary trial statistics for all five closed-loop metrics.

Table 16: Dictionary ablation: detailed per-metric statistics (mean $\pm$ std, median in parentheses, 50 trials).

Metric	GenericRBF	LieD1	Hybrid	CurrentMixed
RMSE _x	0.252 $\pm$ 0.242 (0.133)	0.104 $\pm$ 0.075 (0.085)	0.093 $\pm$ 0.075 (0.069)	0.090 $\pm$ 0.074 (0.065)
t_s (s)	5.49 $\pm$ 7.67 (0.55)	0.67 $\pm$ 0.69 (0.45)	0.68 $\pm$ 0.86 (0.35)	0.52 $\pm$ 0.57 (0.35)
E_u	64.7 $\pm$ 61.9 (43.6)	67.0 $\pm$ 53.5 (51.3)	66.8 $\pm$ 52.8 (50.2)	66.3 $\pm$ 51.9 (49.7)
η	0.00665	0.00197	0.00178	0.00171
Stab. %	96	100	100	100

E Stability Proofs

Proof of Proposition 1. Define $V_k := \min_{\mathbf{u}_k \in \mathcal{F}(z_k)} J_k(\mathbf{u}_k)$ and let $\mathbf{u}_k^*$ be the minimiser. A shifted candidate sequence $\tilde{\mathbf{u}}_{k+1} = (\mathbf{u}_{k+1|k}^*, \dots, \mathbf{u}_{k+N_p-1|k}^*, K_{\text{loc}} \hat{z}_{k+N_p|k})$ is feasible at time $k+1$. Writing the stage and terminal costs as $\ell(z, u) = \|\mathbf{C}z - x^*\|_Q^2 + Ru^2$ and $\ell_f(z) = \|\mathbf{C}z - x^*\|_{Q_f}^2$, the cost of the candidate satisfies

$$\begin{aligned}
V_{k+1} &\leq V_k - \ell(z_k, \mathbf{u}_{k|k}^*) \\
&\quad + \ell(\hat{z}_{k+N_p|k}, K_{\text{loc}} \hat{z}_{k+N_p|k}) \\
&\quad + \ell_f(\hat{z}_{k+N_p+1|k}) - \ell_f(\hat{z}_{k+N_p|k}) + \Delta_d.
\end{aligned}$$

Assumption 2 ensures the three middle terms sum to ≤ 0 in the nominal case. The disturbance deviation $z_{k+1} = \hat{z}_{k+1|k} + d_k$ propagates to give $|\Delta_d| \leq c_1 \bar{d}^2 + c_2 \bar{d} \|z_k - z^*\|_2$, with constants depending on $\|\mathbf{A}\|_2, N_p, \|\mathbf{Q}_f\|_2, \|\mathbf{C}\|_2$. Thus

$$V_{k+1} - V_k \leq -\|\mathbf{C}z_k - x^*\|_Q^2 + c_1 \bar{d}^2 + c_2 \bar{d} \|z_k - z^*\|_2.$$

Outside a ball of radius $r^* = (c_2 \bar{d} + \sqrt{c_2^2 \bar{d}^2 + 4\lambda_{\min}(\mathbf{Q})c_1 \bar{d}^2}) / (2\lambda_{\min}(\mathbf{Q}))$, the quadratic penalty dominates and V_k decreases. Since $\|x_k - x^*\|_2 \leq \|\mathbf{C}\|_2 \|z_k - z^*\|_2$, the bound $\limsup_{k \rightarrow \infty} \|x_k - x^*\|_2 \leq \|\mathbf{C}\|_2 r^* =: \gamma(\bar{d})$ follows, with $\gamma(0) = 0$. $\square$

Proof of Corollary 1. The re-lifting disturbance satisfies $\|d_k\|_2 = \|\psi(F(x_k, u_k)) - A\psi(x_k) - Bu_k\|_2$. The truncation component is bounded by Theorem 1, giving $\bar{d} \leq C_{D,m,\delta} \sup_{x \in \mathcal{X}} \|x\|_2^{D+1} \Delta t$. Substituting into Proposition 1 yields the result. $\square$

References

- Han Bao, Aihuang Hu, Wenbo Liu, and Bocheng Bao. Hidden bursting firings and bifurcation mechanisms in memristive neuron model with threshold electromagnetic induction. *IEEE Transactions on Neural Networks and Learning Systems*, 31(2):502–511, 2020. doi: 10.1109/TNNLS.2019.2905137.
- Alim Louis Benabid, Stephan Chabardes, John Mitrofanis, and Pierre Pollak. Deep brain stimulation of the subthalamic nucleus for the treatment of Parkinson’s disease. *The Lancet Neurology*, 8(1):67–81, 2009. doi: 10.1016/S1474-4422(08)70291-6.
- Steven L. Brunton, Bingni W. Brunton, Joshua L. Proctor, and J. Nathan Kutz. Koopman invariant subspaces and finite linear representations of nonlinear dynamical systems for control. *PLOS ONE*, 11(2):e0150171, 2016. doi: 10.1371/journal.pone.0150171.
- Marko Budišić, Ryan Mohr, and Igor Mezić. Applied Koopmanism. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 22(4):047510, 2012. doi: 10.1063/1.4772195.
- Richard H. Byrd, Peihuang Lu, Jorge Nocedal, and Ciyu Zhu. A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing*, 16(5):1190–1208, 1995. doi: 10.1137/0916069.
- Matthew J. Colbrook and Alex Townsend. Rigorous data-driven computation of spectral properties of Koopman operators for dynamical systems. *Communications on Pure and Applied Mathematics*, 77(1):221–283, 2024. doi: 10.1002/cpa.22125.
- Jeremy Coulson, John Lygeros, and Florian Dörfler. Data-enabled predictive control: In the shallows of the DeePC. *Proceedings of the European Control Conference (ECC)*, pages 307–312, 2019. doi: 10.23919/ECC.2019.8795639. Extended preprint: arXiv:1811.05890.
- Gustavo Deco, Viktor K. Jirsa, Peter A. Robinson, Michael Breakspear, and Karl Friston. The dynamic brain: from spiking neurons to neural masses and cortical fields. *PLOS Computational Biology*, 4(8):e1000092, 2008. doi: 10.1371/journal.pcbi.1000092.
- Richard FitzHugh. Impulses and physiological states in theoretical models of nerve membrane. *Biophysical Journal*, 1(6):445–466, 1961. doi: 10.1016/S0006-3495(61)86902-6.
- Juan A. Gallego, Matthew G. Perich, Lee E. Miller, and Sara A. Solla. Neural manifolds for the control of movement. *Neuron*, 94(5):978–984, 2017. doi: 10.1016/j.neuron.2017.05.025.
- Debdipta Goswami and Derek A. Paley. Bilinearization, reachability, and optimal control of control-affine nonlinear systems: A Koopman spectral approach. *IEEE Transactions on Automatic Control*, 67(6):2715–2728, 2022. doi: 10.1109/TAC.2021.3088802.
- Constance Hammond, Hagai Bergman, and Peter Brown. Pathological synchronization in Parkinson’s disease: networks, models and treatments. *Trends in Neurosciences*, 30(7):357–364, 2007. doi: 10.1016/j.tins.2007.05.004.
- Keita Hara, Masaki Inoue, and Noboru Sebe. Dissipativity-constrained learning of MPC with Koopman operators. *Automatica*, 157:111271, 2023. doi: 10.1016/j.automatica.2023.111271.
- J. L. Hindmarsh and R. M. Rose. A model of neuronal bursting using three coupled first order differential equations. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 221(1222):87–102, 1984. doi: 10.1098/rspb.1984.0024.

- Alan L. Hodgkin and Andrew F. Huxley. A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology*, 117(4):500–544, 1952. doi: 10.1113/jphysiol.1952.sp004764.
- Alexandra B. Holt, Dan Wilson, Maxwell Shinn, Jeff Moehlis, and Theoden I. Netoff. Phasic burst stimulation: A closed-loop approach to tuning deep brain stimulation parameters for Parkinson’s disease. *PLoS Computational Biology*, 12(7):e1005011, 2016. doi: 10.1371/journal.pcbi.1005011.
- Eugene M. Izhikevich. Simple model of spiking neurons. *IEEE Transactions on Neural Networks*, 14(6):1569–1572, 2003. doi: 10.1109/TNN.2003.820440.
- Eurika Kaiser, J. Nathan Kutz, and Steven L. Brunton. Data-driven discovery of Koopman eigenfunctions for control. *Machine Learning: Science and Technology*, 2(3):035023, 2021. doi: 10.1088/2632-2153/abf0f5.
- B. O. Koopman. Hamiltonian systems and transformation in Hilbert space. *Proceedings of the National Academy of Sciences*, 17(5):315–318, 1931. doi: 10.1073/pnas.17.5.315.
- Milan Korda and Igor Mezić. Linear predictors for nonlinear dynamical systems: Koopman operator meets model predictive control. *Automatica*, 93:149–160, 2018. doi: 10.1016/j.automatica.2018.03.046.
- Qianxiao Li, Felix Dietrich, Erik M. Bollt, and Ioannis G. Kevrekidis. Extended dynamic mode decomposition with dictionary learning: A data-driven adaptive spectral decomposition of the Koopman operator. *Chaos*, 27(10):103111, 2017. doi: 10.1063/1.4993854.
- Simon Little, Alek Pogosyan, Stephane Neal, Baltazar Zavala, Ludvic Zrinzo, Marwan Hariz, Tom Foltynie, Patricia Limousin, Keyoumars Ashkan, James FitzGerald, Alexander L. Green, Tipu Z. Aziz, and Peter Brown. Adaptive deep brain stimulation in advanced Parkinson disease. *Annals of Neurology*, 74(3):449–457, 2013. doi: 10.1002/ana.23951.
- Simon Little, Elina Tripoliti, Martje Beudel, Alek Pogosyan, James J. FitzGerald, Gerd Tinkhauser, Wesley Thevathasan, Jonathan A. Hyam, Tom Foltynie, Patricia Limousin, Ludvic Zrinzo, Marwan Hariz, Alexander L. Green, Tipu Z. Aziz, and Peter Brown. Adaptive deep brain stimulation for Parkinson’s disease demonstrates reduced speech side effects compared to conventional stimulation in the acute setting. *Journal of Neurology, Neurosurgery & Psychiatry*, 87(12):1388–1389, 2016. doi: 10.1136/jnnp-2016-313518.
- Andres M. Lozano, Nir Lipsman, Hagai Bergman, Peter Brown, Stephan Chabardes, Jin Woo Chang, Keith Matthews, Cameron C. McIntyre, Thomas E. Schlaepfer, Michael Schulder, Yasin Temel, Rashid Bhatt, Sameer Bhatt, et al. Deep brain stimulation: current challenges and future directions. *Nature Reviews Neurology*, 15(3):148–160, 2019. doi: 10.1038/s41582-018-0128-2.
- Bethany Lusch, J. Nathan Kutz, and Steven L. Brunton. Deep learning for universal linear embeddings of nonlinear dynamics. *Nature Communications*, 9:4950, 2018. doi: 10.1038/s41467-018-07210-0.
- Mi Lv, Chunni Wang, Guodong Ren, Jun Ma, and Xinlin Song. Model of electrical activity in a neuron under magnetic flow effect. *Nonlinear Dynamics*, 85(3):1479–1490, 2016. doi: 10.1007/s11071-016-2773-6.
- Jun Ma and Jun Tang. A review for dynamics in neuron and neuronal network. *Nonlinear Dynamics*, 89(3):1569–1578, 2017. doi: 10.1007/s11071-017-3565-3.
- Alexandre Mauroy and Igor Mezić. Global stability analysis using the eigenfunctions of the Koopman operator. *IEEE Transactions on Automatic Control*, 61(11):3356–3369, 2016. doi: 10.1109/TAC.2016.2518918.
- Alexandre Mauroy, Igor Mezić, and Yoshihiko Susuki. The Koopman operator in systems and control: Concepts, methodologies, and applications. *Lecture Notes in Control and Information Sciences*, 484, 2020. doi: 10.1007/978-3-030-35713-9.

- Cameron C. McIntyre, Warren M. Grill, David L. Sherman, and Nitish V. Thakor. Cellular effects of deep brain stimulation: model-based analysis of activation and inhibition. *Journal of Neurophysiology*, 91(4): 1457–1469, 2004. doi: 10.1152/jn.00989.2003.
- Igor Mezić. Spectral properties of dynamical systems, model reduction and decompositions. *Nonlinear Dynamics*, 41(1-3):309–325, 2005. doi: 10.1007/s11071-005-2824-x.
- Bharat Monga, Dan Wilson, Tim Matchen, and Jeff Moehlis. Phase reduction and phase-based optimal control for biological systems: a tutorial. *Biological Cybernetics*, 113(1–2):11–46, 2019. doi: 10.1007/s00422-018-0780-z.
- J. Nagumo, S. Arimoto, and S. Yoshizawa. An active pulse transmission line simulating nerve axon. *Proceedings of the IRE*, 50(10):2061–2070, 1962. doi: 10.1109/JRPROC.1962.288235.
- Zeric Tabekoueng Njitacke, Jacques Kengne, and Hilaire Bertrand Fotsin. Coexistence of multiple stable states and bursting oscillations in a 4D Hopfield neural network. *Circuits, Systems, and Signal Processing*, 39: 3424–3444, 2020. doi: 10.1007/s00034-019-01324-6.
- Fatemeh Parastesh, Sajjad Jafari, Hamed Azarnoush, Zahra Shahriari, Zhouchao Wang, Stefano Boccaletti, and Matjaž Perc. Chimeras. *Physics Reports*, 898:1–114, 2021. doi: 10.1016/j.physrep.2020.10.003.
- Oleksandr V. Popovych and Peter A. Tass. Desynchronizing electrical and sensory coordinated reset neuromodulation. *Frontiers in Human Neuroscience*, 6:58, 2012. doi: 10.3389/fnhum.2012.00058.
- Joshua L. Proctor, Steven L. Brunton, and J. Nathan Kutz. Dynamic mode decomposition with control. *SIAM Journal on Applied Dynamical Systems*, 15(1):142–161, 2016. doi: 10.1137/15M1013857.
- Steven J. Schiff. *Neural Control Engineering: The Emerging Intersection Between Control Theory and Neuroscience*. MIT Press, Cambridge, MA, 2012. doi: 10.7551/mitpress/8436.001.0001.
- Peter J. Schmid. Dynamic mode decomposition of numerical and experimental data. *Journal of Fluid Mechanics*, 656:5–28, 2010. doi: 10.1017/S0022112010001217.
- Robert V. Shannon. A model of safe levels for electrical stimulation. *IEEE Transactions on Biomedical Engineering*, 39(4):424–426, 1992. doi: 10.1109/10.126616.
- Scott Stanslaski, Pedram Afshar, Peng Cong, Jon Giftakis, Paul Stypulkowski, Dave Carlson, David Linde, David Ullestad, Al-Thaddeus Avestruz, and Tim Denison. Design and validation of a fully implantable, chronic, closed-loop neuromodulation device with concurrent sensing and stimulation. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 20(4):410–421, 2012. doi: 10.1109/TNSRE.2012.2183617.
- Peter A. Tass. A model of desynchronizing deep brain stimulation with a demand-controlled coordinated reset of neural subpopulations. *Biological Cybernetics*, 89(2):81–88, 2003. doi: 10.1007/s00422-003-0425-7.
- Jonathan H. Tu, Clarence W. Rowley, Dirk M. Luchtenburg, Steven L. Brunton, and J. Nathan Kutz. On dynamic mode decomposition: Theory and applications. *Journal of Computational Dynamics*, 1(2):391–421, 2014. doi: 10.3934/jcd.2014.1.391.
- Peter J. Uhlhaas and Wolf Singer. Neural synchrony in brain disorders: relevance for cognitive dysfunctions and pathophysiology. *Neuron*, 52(1):155–168, 2006. doi: 10.1016/j.neuron.2006.09.020.
- A. Velisar, J. Syrkin-Nikolau, Z. Blumenfeld, M. H. Trager, M. F. Afzal, V. Prabhakar, and H. M. Bronte-Stewart. Dual threshold neural closed loop deep brain stimulation in Parkinson disease patients. *Brain Stimulation*, 12(4):868–876, 2019. doi: 10.1016/j.brs.2019.02.020.
- Matthew O. Williams, Ioannis G. Kevrekidis, and Clarence W. Rowley. A data-driven approximation of the Koopman operator: Extending dynamic mode decomposition. *Journal of Nonlinear Science*, 25(6):1307–1346, 2015. doi: 10.1007/s00332-015-9258-5.

- Dan Wilson and Jeff Moehlis. Isostable reduction of periodic orbits. *Physical Review E*, 94(5):052213, 2016. doi: 10.1103/PhysRevE.94.052213.
- Xinglong Zhang, Wei Pan, Riccardo Scattolini, Shuyou Yu, and Xin Xu. Robust tube-based model predictive control with Koopman operators. *Automatica*, 137:110114, 2022. doi: 10.1016/j.automatica.2021.110114.