

Myopic control of neural dynamics

David Hocker¹ , Il Memming Park^{* 1,2,3} ,

1 Department of Neurobiology and Behavior Stony Brook University, Stony Brook, NY, USA

2 Department of Applied Mathematics and Statistics, Stony Brook University, Stony Brook, NY, USA

3 Institute for Advanced Computational Science, Stony Brook University, Stony Brook, NY, USA

 These authors contributed equally to this work.

* memming.park@stonybrook.edu

Abstract

Manipulating the dynamics of neural systems through targeted stimulation is a frontier of research and clinical neuroscience; however, the control schemes considered for neural systems are mismatched for the unique needs of manipulating neural dynamics. An appropriate control method should respect the variability in neural systems, incorporating moment to moment “input” to the neural dynamics and behaving based on the current neural state, irrespective of the past trajectory. We propose such a controller under a nonlinear state-space feedback framework that steers one dynamical system to function as through it were another dynamical system entirely. This “myopic” controller is formulated through a novel variant of a model reference control cost that manipulates dynamics in a short-sighted manner that only sets a target trajectory of a single time step into the future (hence its myopic nature), which omits the need to pre-calculate a rigid and computationally costly neural feedback control solution. To demonstrate the breadth of this control’s utility, two examples with distinctly different applications in neuroscience are studied. First, we show the myopic control’s utility to probe the causal link between dynamics and behavior for cognitive processes by transforming a winner-take-all decision-making system to operate as a robust neural integrator of evidence. Second, an unhealthy motor-like system containing an unwanted beta-oscillation spiral attractor is controlled to function as a healthy motor system, a relevant clinical example for neurological disorders.

Author summary

Stimulating a neural system and observing its effect through simultaneous observation offers the promise to better understand how neural systems perform computations, as well as for the treatment of neurological disorders. A powerful perspective for understanding a neural system’s behavior undergoing stimulation is to conceptualize them as dynamical systems, which considers the global effect that stimulation has on the brain, rather than only assessing what impact it has on the recorded signal from the brain. With this more comprehensive perspective comes a central challenge of determining what requirements need to be satisfied to harness neural observations and then stimulate to make one dynamical system function as another one entirely. This could lead to applications such as neural stimulators that make a diseased brain behave

1
2
3
4
5
6
7
8
9
10

like its healthy counterpart, or to make a neural system previously capable of only hasty decision making to wait and accumulate more evidence for a more informed decision. In this work we explore the implications of this new perspective on neural stimulation and derive a simple prescription for using neural observations to inform stimulation protocol that makes one neural system behave like another one.

Introduction

Advances in recording technology are making it possible to gain real-time access to neural dynamics at different length and time scales [1, 2], allowing us to consider the structure of the brain’s operation in ways that were previously inaccessible. Central to that understanding of neural dynamics is the widely-held belief that *dynamical systems* underlie all of the core operations of neural systems [3–6]. Dynamical systems are systems of time-independent dynamics that drive the evolution of a set latent states that may or may not be directly observable, which in neural systems are proposed to account for motor function [7], cognitive processes [8–10], and sensory processing [11]. The controlled stimulation of neural systems offers not only a novel tool to perturbatively study the underlying dynamical systems; but also shows tremendous potential to treat a host of brain disorders, ranging from movement diseases such as Parkinson’s disease and essential tremor [12, 13], epilepsy [14, 15], and even mood disorders such as severe depression [16]. In particular, there has been recent success in combining real-time neural data acquisition with closed-loop stimulation for treating Parkinson’s disease [17, 18].

Unfortunately, the current framework for manipulating neural systems is not structured to deal with the unique challenges posed by controlling complex neural dynamics. One of the central goals of control theory is to manipulate a system to mimic some or all characteristics of a target system of dynamics, and nearly all control systems accomplish this by controlling the system state to either track a specified target trajectory or to regulate to a known set point [19]. Closed-loop control systems specifically designed for neural systems also operate under this paradigm [20–22], and clinical devices use even more simplistic open-loop or reactive protocols [15, 17, 23]. If neural systems function as a dynamical system by nonlinearly filtering signals [24, 25], then significant portions of the observed neural fluctuation would correspond to relevant exogenous input signals to the system such as volition, memory or sensory information. Such controls designed to move to or maintain a target state counteract any natural fluctuation in neural trajectories, and create a rigid system that is no longer dynamically computing. For example, when building neural prosthetics for an abnormal motor-related brain area, it is crucial for the controlled neural activity to be close to normal; however, simply controlling it to replay a fixed motor command would not allow flexibly changing one’s mind mid action. Therefore, any control objective that only considers externally set constraints through trajectory or set-point control would be limited both in their application for treating neurodynamic diseases as well as for studying neural computations in cases where preserving dynamic information processing capability is important.

Given this perspective, we propose a new control objective called *myopic control* that respects the unforeseeable variability in neural systems. The objective of myopic control is for the controlled system to behave as a target neural dynamical system. This is reminiscent of a well-developed field in control theory known as model reference control (MRC) [19], though MRC has been widely used for trajectory-tracking problems. Unlike MRC, myopic control is independent of the past trajectory and does not account for the far future—given the current state of the system, it tries to behave as the target dynamical system *instantaneously*. Additionally, our controller is constructed to be

agnostic about the ideal behavior of the neural system. Its only purpose is to generate the target dynamical system, and all of its emerging behaviors, as accurately as possible. The controller does not assume the role of performing the bulk action on the state, which is instead encompassed in the original dynamical system that presumably perform some form of related dynamics well. This is especially important in the context of neural dynamics performing a computation, where it would be undesirable for our controller to first perform the computation itself by tracing out a predefined trajectory. Instead, myopic control will assist that system’s natural ability to perform a neural computation.

The qualitative difference between our control scheme and trajectory-tracking methods is depicted in Fig. 1. Given some target dynamical system, utilizing trajectory control would force the neural system to follow a target trajectory, although not through the true target dynamics. Scenarios may arise where trajectory control and myopic control may be very similar (Fig. 1A), although there can be fundamental, qualitative differences in the presence of noise or large disturbances due to exogenous inputs (Fig. 1B). In that case, the trajectory resulting from trajectory control would not be generated from the target dynamics, and forces the state to evolve toward the pre-computed target state. In this way, our controller preserves the full neural variability of our target dynamics, ranging from potentially different trajectories towards the same fixed point to even allowing for potentially different behavior than expected.

Fig 1. Qualitative difference between our proposed *myopic* control of dynamics and trajectory control. Here $\mathcal{F}$ is controlled to perform an example target dynamics $\mathcal{G}$ (e.g., perform a motor command), where its gradient flow is given in gray and two attractors are denoted as circles. A precomputed target trajectory $\mathbf{x}_t$ through $\mathcal{G}$ is shown in black. **A)** In the presence of small disturbances, the evolution of trajectory control forces the system back to $\mathbf{x}_t$, whereas myopic control allows for natural deviations. **B)** A large disturbance away from $\mathbf{x}_t$ corresponding to an exogenous input that changes the target attractor mid-trajectory could lead to entirely different behavior between the two control methods. Only myopic control would capture the response of this disturbance through the true dynamics of $\mathcal{G}$, while trajectory control blindly follows $\mathbf{x}_t$.

The paper is organized as follows. First, we formulate the goals of our control objective for manipulating neural systems, then define myopic control for linear and nonlinear dynamics. Next, we discuss some design features of how to construct the target dynamics of a desired dynamical system to display expected or intended behavior, and what types of difficulties may arise when trying to define healthy or desired neural dynamics. We then demonstrate this control’s ability to make dynamical systems act as though they were another system entirely through two relevant examples. First, a winner-take-all decision-making model is transformed to operate as a robust neural integrator of information when shown a stimulus in a forced, two-choice decision-making task. Second, a “diseased” motor system containing an unwanted beta-oscillation state is controlled to function as a healthy motor system, which is a motivating example for the treatment of movement disorders or other diseases with an underlying neurological state.

Materials and methods

Myopic dynamics control

Here we discuss the control problem of utilizing a dynamical system to behave as a separate dynamical system. Using a Bayesian state-space modeling framework [26], we are interested in the time evolution of a posterior distribution of time-dependent,

n -dimensional (latent) brain state x_t that are governed by (stochastic) dynamics $\mathcal{F}[x_t, u_t] \equiv \mathcal{F}_t$ with an m -dimensional control signal u_t ,

$$x_{t+1} = x_t + \mathcal{F}_t + w_t, \quad (1)$$

where $w_t \sim \mathcal{N}(0, Q)$ is the state noise upon the dynamics. A second set of *target* stochastic dynamics $\mathcal{G}[x_t] \equiv \mathcal{G}_t$ under which we would like our state to evolve, acts analogously on the state as

$$x_{t+1} = x_t + \mathcal{G}_t + w_t, \quad (2)$$

The noise in both dynamics is the same, as we are considering transforming $\mathcal{F}$ into $\mathcal{G}$ in the same physical neural system. In general, the control acts upon the dynamical system latent states x that may not be directly observable, and would need to be inferred from a set of observable variable to which the latent states are linked through an observation model. The influence of the controls would also be manifested in the observed *neural observations* that are relevant to experiments (*e.g.*, calcium image traces, local field potentials, etc.), though without loss of generality we have chosen to simplify our dynamics by omitting an observation model. While neural observations may contain much information about the system, generally speaking they are not dynamical states. States of a dynamical system require no time dependence to fully describe the system, unlike neural observations that may require a history to understand the dynamics. These dynamics are in general nonlinear, and we denote their Jacobians (linearization at the current state and stimulus) as

$$A_t = \left. \frac{\partial \mathcal{F}[x, u]}{\partial x} \right|_{x_t, u_t}, \quad \tilde{A}_t = \left. \frac{\partial \mathcal{G}[x]}{\partial x} \right|_{x_t}, \quad B_t = \left. \frac{\partial \mathcal{F}[x, u]}{\partial u} \right|_{x_t, u_0} \quad (3)$$

Arguably the most developed form of model-based control occurs for linear systems with quadratic costs on the state and control, known as linear quadratic gaussian (LQG) control [27]. Finite-time horizon LQG controllers are optimal for costs of the simplified form

$$J = \sum_{t=0}^T \mathbb{E}_x [\|x_t - \mathbf{x}_t\|^2] + \gamma u_t^T u_t, \quad (4)$$

with linear dynamics $\mathcal{F}_t = Ax_t + Bu_t + w_t$, and a regularization penalty factor γ is added onto the control power. The goal of minimizing (4) is to balance tracking along a target trajectory $\mathbf{x}_t$ with the cost of implementing a control. The optimal LQG controller form $u_t^* = K_t(x - \mathbf{x}_t)$ with gain K_t is found by solving the associated recursive Riccati equation from an end-point condition, and is a time-dependent controller through the time-dependence on K_t [27].

Generating target dynamics is similar in spirit to LQG-type costs, although instead we are interested in minimizing the difference between the effect of target dynamics and controlled dynamics alongside control costs. Requesting that the controlled dynamics of $\mathcal{F}_t$ act as through they are in fact $\mathcal{G}_t$ can be written in a regularized, stepwise quadratic form as

$$J_t = \mathbb{E}_x [(\mathcal{F}_t - \mathcal{G}_t)^T (\mathcal{F}_t - \mathcal{G}_t)] + \gamma u_t^T u_t. \quad (5)$$

Note that this cost is defined at each time point t , and depends on the current state (posterior) distribution over x_t . Utilizing control to track a defined trajectory that is generated from an uncontrolled set of target dynamics $\mathcal{G}$ is known as model reference control (MRC), [19] although the costs associated with this control design are traditionally limited to regulation of a controlled trajectory around a set point or tracking of a predefined target trajectory evolving under $\mathcal{G}$ over a long time horizon. Our cost in (5) is equivalent to MRC with a time horizon of $T = 1$, in which the control effectively recreates a *single* step of a target trajectory from $\mathcal{G}$. To our knowledge, this simplified form of MRC is a major departure from the typical use of model reference control. By weighting the difference between dynamics over a single time step, this myopic (*i.e.*, one-step) form negates the need to solve the Riccati equations, and the derivative $\partial J / \partial u_t$ can be straightforwardly calculated to identify the optimal *myopic control*.

Our work in this paper focuses primarily on designing a controller that optimizes eq. (5), which would be optimal for generating target dynamics over a single step. Since the controller would no longer contain any time dependence (the dynamics $\mathcal{F}_t$ and $\mathcal{G}_t$ are indexed by their current time, but are dependent upon the state x_t only), it would generate a dynamical system with the same state space. The qualitative advantages of myopic control are depicted in Figure 1, in which the evolution of a trajectory-controlled system tracking a defined trajectory $\mathbf{x}_t$ in a target dynamical system $\mathcal{G}$ is compared to the evolution of a myopically controlled system designed to perform the target dynamics. In a noiseless environment, both trajectories would be identical; however, in the presence of small disturbances away from $\mathbf{x}_t$, tracking control would correct the trajectory in a distinctly non-dynamical fashion, evolving not through $\mathcal{G}$ but instead forcing the system back onto $\mathbf{x}_t$ in an unnatural manner (Fig. 1A). Myopic control would instead lead trajectory through the natural dynamics of $\mathcal{G}$, which may lead to the same stable point, but through a distinctly different trajectory. Some disturbances may lead to different behavior between the two control methods, though. Figure 1B shows this scenario, in which a disturbance is corrected by trajectory control back toward $\mathbf{x}_t$, while myopic control followed the flow of $\mathcal{G}$, which lead it to a different attractor point. If this target dynamics were a decision-making computation, for example, myopic control may have lead to a “wrong” decision; however, allowing a controlled neural system to operate imperfectly in the perspective of modern control is precisely the type of flexibility that should be achieved to maintain its natural operation.

In the following sections we derive the form of our myopic controller. Ideally, the controller formulation will be distinct from the state estimator providing the feedback signal, and leads us to consider variants of the controller that rely upon different moments of the underlying state distribution. We first begin with the case of linear dynamics to demonstrate the simplified form of myopic control and its properties, then move the more applicable nonlinear case.

Linear dynamics

Here we demonstrate that the myopic controller for linear dynamics depends only upon the mean of the state distribution, and thus the state estimator and controller design are separable for myopic control.

Theorem 1. *If target and controlled dynamics are linear in state x and control u , then myopic control depends only upon state mean $\mathbb{E}[x]$.*

Proof. Let the linear dynamics under control and the target dynamics be

$$\mathcal{F}_t = Ax_t + Bu_t + w_t \quad (6)$$

$$\mathcal{G}_t = \tilde{A}x_t + w_t, \quad (7)$$

where the state distribution over x has first and second central moments $\mathbb{E}[x_t] = \mu_t$, $\mathbb{E}[(x_t - \mu_t)^2] = \Sigma_t$, and the state noise is normal with $w_t \sim \mathcal{N}(0, Q)$. Expanding the dynamics cost in (5) gives

$$\begin{aligned} J &= \mathbb{E}_x \left[|(A - \tilde{A})x_t + Bu_t|^2 \right] + \gamma u_t^T u_t \\ &= \text{Tr} \left[|(A - \tilde{A})|^2 \Sigma_t \right] + \mu_t^T (A - \tilde{A})^T B u_t + \\ &\quad u_t^T B^T (A - \tilde{A}) \mu_t + u_t^T (B^T B + \gamma I_m) u_t, \end{aligned} \quad (8)$$

where I_m is the $m \times m$ identity matrix. By examining (8) it is clear that regardless of the distribution over x , the cost depends only upon the first two moments of the distribution of x . Maximizing (8) yields the optimal linear myopic controller form $u_{t_{\text{lin}}}^*$, which depends only upon the state mean,

$$u_{t_{\text{lin}}}^* = -2(B^T B + \gamma I_m)^{-1} B^T (A - \tilde{A}) \mu_t. \quad (9)$$

□

Nonlinear dynamics controller with a moment expansion approximation

For nonlinear dynamics, simply differentiating (5) leads to an ill-suited expression for a controller, since there is an implicit dependence of the controller upon itself through $\mathcal{F}$. One approximation to alleviate this is to expand the nonlinear dynamics about null control ($u_0 = 0$) to first order, with the form

$$\begin{aligned} \mathcal{F}[x_t, u_t] &\approx \mathcal{F}[x_t, u_0] + B_t(u_t - u_0) \\ &\equiv f_t + B_t u_t \end{aligned} \quad (10)$$

where $f_t \equiv f[x_t] \equiv \mathcal{F}[x_t, u_0]$ and for the remainder of the work $B_t \equiv B[x_t, x_0]$ is the Jacobian of $\mathcal{F}$ as in eq. (3). This leads to an expression for the derivative of J and myopic controller as

$$\frac{\partial J}{\partial u_t} \approx \mathbb{E}_x [2(f_t^T + u_t^T B_t^T) B_t - 2\mathcal{G}_t^T B_t] + 2\gamma u_t^T \quad (11)$$

$$u_t^* = -(\mathbb{E}_x [B_t^T B_t] + \gamma I_m)^{-1} \mathbb{E}_x [B_t^T (f_t - \mathcal{G}_t)]. \quad (12)$$

The expectations in (12) depend upon the state distribution of x_t , although it would be desirable if akin to LQG that the controller was separated from state estimation, and only depended upon low-order moments of x . To construct such a controller we will expand $\mathbb{E}_x[\cdot]$ in terms of the mean and covariance of x_t ; in general, the terms in this expansion will contain Jacobian matrices, higher order derivatives, and state vectors that are all evaluated at the distribution mean μ_t , multiplied by the covariance Σ_t in some form. For example, the Jacobian B_t is expanded as

$$\begin{aligned}
B[x_t] &= B[\mu_t + (x_t - \mu_t)] \\
&\approx B[\mu_t] + B'[\mu_t](x_t - \mu_t) + \frac{1}{2}(x_t - \mu_t)^T B''[\mu_t](x_t - \mu_t),
\end{aligned} \tag{13}$$

and would follow similarly for the other terms in $\mathbb{E}_x[\cdot]$. Such an approximation is valid when the deviations from our estimated state μ_t are small, and in this regime only low-order moments are necessary. It is assumed that state estimation to obtain μ_t and Σ_t can be performed regularly enough in practice to operate in the regime such that (13) is valid, and we will consequently consider two forms of nonlinear myopic control. *First-order* myopic control will include only terms dependent upon state mean, just as in the linear dynamics case of the previous section. *Second-order* myopic control will analogously depend upon both μ_t and Σ_t . In each controller the terms $f, \mathcal{G}, B$ and derivatives $B'_t = \partial B_t / \partial x_t, B''_t = \partial^2 B_t / \partial x_t^2$, will all be evaluated at the distribution mean μ_t and null control $u_0 = 0$, so we will temporarily drop the functional dependence of these terms in the notation. The prime notation will indicate a derivative with respect to state.

The first expectation in (12) includes only terms relating to B . Expanding and keeping terms up to second order gives

$$\begin{aligned}
\mathbb{E}_x[B[x]^T B[x]] &= \mathbb{E}_x[B[\mu + (x - \mu)]^T B[\mu + (x - \mu)]] \\
&\approx B^T B + \frac{1}{2} B^T \text{Tr}_{3,4}[B''\Sigma] + \\
&\quad \frac{1}{2} \text{Tr}_{3,4}[B''^T \Sigma] B + \text{Tr}_{3,4}[B'^T B' \Sigma],
\end{aligned} \tag{14}$$

where $\text{Tr}_{3,4}$ denotes the partial trace over dimensions 3 and 4. For an $(n \times m \times n \times n)$ tensor T this operation maps to an $(n \times m)$ matrix $M = \text{Tr}_{3,4}[T]$ as

$$M_{j,k} = \sum_i T_{j,k,i,i}. \tag{15}$$

Similarly, expanding $\mathbb{E}_x[B_x^T(f[x] - \mathcal{G}[x])]$ up to second order yields

$$\begin{aligned}
&\mathbb{E}_x[B^T[\mu + (x - \mu)](f[\mu + (x - \mu)] - \mathcal{G}[\mu + (x - \mu)])] \\
&= B^T(f - \mathcal{G}) + \frac{1}{4} B^T \text{Tr}_{2,3}[(f'' - \mathcal{G}'')\Sigma] + \\
&\quad B^T \text{Tr}_{2,3}[B'^T(f' - \mathcal{G}')\Sigma] + \frac{1}{2} \text{Tr}_{3,4}[B''^T \Sigma](f - \mathcal{G}).
\end{aligned} \tag{16}$$

(12), (14), and (16) define our second-order nonlinear myopic controller $u_{2\text{nd}}$, and simply omitting the covariance-dependent terms gives first order controller expansions,

$$\mathbb{E}_x[B[x]^T B[x]]_{1\text{st order}} = B[\mu]^T B[\mu] \tag{17}$$

$$\mathbb{E}_x[B_x^T(f[x] - \mathcal{G}[x])]_{1\text{st order}} = B[\mu]^T(f[\mu] - \mathcal{G}[\mu]). \tag{18}$$

First-order control $u_{1\text{st}}$ is attractive for its simplicity, and it is important to ask under what circumstances would first-order control outperform second-order control?

First, if $\text{Tr}(\Sigma_t)$ is very small (*i.e.*, small uncertainty about the state x_t), then second-order terms are negligible. Second, by noting that nearly all second-order terms contain derivatives of B, f and $\mathcal{G}$, another regime in which first-order control may be superior is under “super smooth” dynamics in which the magnitude of successive derivatives is smaller than the previous one (*e.g.* $\|B\| > \|B'\| > \|B''\|$). Moreover, if the control portion of the Jacobian B is state-independent, then second-order control only has one covariance-dependent term in $\mathbb{E}_x[B_x^T(f[x] - \mathcal{G}[x])]$,

$$B[\mu]^T(f[\mu] - \mathcal{G}[\mu]) + \frac{1}{4}B[\mu]^T \text{Tr}_{2,3}[(f''[\mu] - \mathcal{G}''[\mu])\Sigma]. \quad (19)$$

Evaluating controller performance

The performance of a myopic controller is formally benchmarked by the regularized cost in (5), although it is important to isolate a cost describing the performance of only the dynamics. The cost of expected mean performance is denoted by $\tilde{J}_{\mu_t}$, and is given by

$$\tilde{J}_{\mu_t} = [(\mathcal{F}[\hat{\mu}_t, u_t^*] - \mathcal{G}[\hat{\mu}_t])^T (\mathcal{F}[\hat{\mu}_t, u_t^*] - \mathcal{G}[\hat{\mu}_t])]. \quad (20)$$

This cost is easy to compute and provides information about the mean behavior of the control, although it ignores the variability of the state distribution Σ_t . A more informative cost incorporates the impact of the entire distribution of x , denoted by $\tilde{J}_t$ as

$$\tilde{J}_t = \mathbb{E}_x[(\mathcal{F}[x_t, u_t^*] - \mathcal{G}[x_t])^T (\mathcal{F}[x_t, u_t^*] - \mathcal{G}[x_t])]. \quad (21)$$

This is simply the log-mean squared error of the controlled dynamics. We estimate (21) through Monte Carlo integration assuming the maximum entropy distribution at each time point given the first two moments, *i.e.*, a normal distribution $N(\hat{\mu}_t, \hat{\Sigma}_t)$.

Design principles for targeted dynamical systems

Myopic control omits the requirement of supplying a target neural trajectory or set point in the neural state space, which resonates with our design requirement of a future-agnostic controller that need not prescribe what the brain should be doing precisely.

Balancing the simplicity and ease of myopic control, though, is the relative complexity in designing a target dynamical system $\mathcal{G}_t$. At first glance, it may seem as though we have merely shifted complications of controlling neural dynamics. However, this perspective more clearly frames the goal of neural dynamics control, and we believe that it identifies a general design question yet to be seriously considered by the neural processing community: Given a rough sketch of neural dynamics and a desire to change them, what is an appropriate target dynamical system?

The choice of $\mathcal{G}$ can roughly be broken down into three design problems for dynamical systems: i) removal or avoidance of an unwanted feature, ii) addition of a desired feature, and iii) modification of an existing feature. For example, there may be attractors (representing macrostates) in $\mathcal{F}$ indicative of a dysfunctional behavior that should be avoided for healthy brain function, such as limit cycle attractors. Or, one may wish to introduce additional attractor macrostates in a decision-making system in order to support robust neural integration of evidence [28]. We will consider both of these scenarios in the following sections. Our ideal design approach used here is summarized in Figure 2, which is to use multiplicative filters upon the controlled dynamics $\mathcal{F}$ to

preserve desired features, with the addition of either barrier functions to remove undesirable aspects of $\mathcal{F}$ or to prevent access into that region of state space. Alternatively an additive function could be utilized to introduce new features. Care must be taken with the shape and positioning of the additive barrier or extra feature though, as any zero crossings of this additive term will introduce fixed points into the dynamics. In the example case in Figure 2, a barrier function is used to remove an undesirable feature of the dynamical system by producing a net rightward gradient flow in the low x_1 region of state space, where the zero crossing of the barrier function is aligned with other fixed points of the system that are denoted in blue. Under this strategy, we can also view modification of an existing feature as simply a removing it and replacing it with the desired one.

Fig 2. Design strategy for creating target dynamics $\mathcal{G}$. Green and blue regions represent features of the original dynamics $\mathcal{F}$ that are to be maintained, while an undesirable feature is denoted with orange. A multiplicative filter removes the unwanted feature, while an additive barrier function prevents access to unwanted state space by enforcing a gradient flow toward desirable regions with well-behaved dynamics.

Numerical methods

In the follow sections we detail the dynamics of each dynamical system in the examples. Since the primary objective of this work is to understand the performance of the myopic controller, in both examples we used a simple state estimator to calculate $\hat{x}_t$ and $\hat{\Sigma}_t$, employing extended Kalman filtering (EKF) within Tensorflow assuming a noisy observation of the state as $y_t = x_t + v_t$, where $v_t \sim N(0, R)$ and R is a diagonal covariance matrix. For lags in observations and control signal calculation, state prediction was performed by propagating $\hat{x}_t$ via,

$$\hat{x}_{t+1} = \hat{x}_t + \mathcal{F}[\hat{x}_t, u_t], \quad (22)$$

and covariance was estimated through sampling of time-evolved state predictions $x_t^{(i)} \sim N(\hat{x}_t, \hat{\Sigma}_t)$ that also evolve in time using (22),

$$\hat{\Sigma}_{t+1} = \frac{1}{M-1} \left[\sum_{i=1}^M x_{t+1}^{(i)} - \mu_{t+1} \right] \left[\sum_{j=1}^M x_{t+1}^{(j)} - \mu_{t+1} \right]^T \quad (23)$$

$$\mu_{t+1} = \frac{1}{M} \sum_{k=1}^M x_{t+1}^{(k)}. \quad (24)$$

Beta oscillation disease states

The diseased state dynamics are a modification of a dynamical system used to describe linear integrate-and-fire neurons as a limit cycle attractor [6]. The specific limit cycle attractor is based upon the post-saddle-node bifurcation behavior at large current in the $I_{Na,p} + I_K$ model from Ch. 4 of [6] (eqs. 4.1-4.2). The high-threshold parameters in [6] were utilized to generate the attractor, and the external current was tuned to generate a beta oscillation. The beta oscillation dynamics of state $X = [X_1, X_2]^T$ (omitting state and observation noise, for succinctness) are

$$\mathcal{F}_{\text{limit cycle}} = \frac{\partial}{\partial t} \begin{bmatrix} X_1 \\ X_2 \end{bmatrix} \quad (25)$$

$$\begin{aligned} \frac{\partial \tilde{X}_1}{\partial t} = & \frac{C_1}{C_0} (I - g_l(\tilde{X}_1 - E_l) - \\ & g_{Na} m_\infty(\tilde{X}_1 - E_{Na}) - g_k X_2(\tilde{X}_1 - E_k)) \end{aligned} \quad (26)$$

$$\frac{\partial X_2}{\partial t} = C_2 \frac{n_\infty - X_2}{\tau} \quad (27)$$

$$p_\infty = \frac{1}{1 + \exp(\frac{V_{p_\infty} - \tilde{X}_1}{k_{p_\infty}})}, \quad p \in [m, n] \quad (28)$$

with parameters $C_0 = 1$, $I = 10$, $E_l = -80$, $E_{Na} = 60$, $E_k = -90$, $g_{Na} = 20$, $g_k = 10$, $g_L = 8$, $\tau = 1$, $V_{m,\infty} = -20$, $V_{n,\infty} = -25$, $k_{m,\infty} = 15$, $k_{n,\infty} = 5$. The original dynamics for X_2 corresponded to an activation variable in an integrate-and-fire model, and as such were scaled to operate at the order of magnitude $X_2 \in [0, 1]$; however, X_1 was originally a voltage variable, and we rescaled it such that $\tilde{X}_1 = 180X_1 - 80$. The magnitude of these dynamics were also scaled with $C_1 = 0.88$ and $C_2 = 160$ to reflect this change in X_1 .

The three stable points of the dynamics were added to the limit cycle attractor dynamics as two sets of gaussian-weighted Gabor functions centered at the three stable points $m_1 = [0.9, 0.25]$, $m_2 = [0.9, 0.50]$, $m_3 = [0.9, 0.75]$ with a width of the gaussian envelope $L = 0.2$. This portion of the dynamics was structured and scaled as

$$\mathcal{F}_a = \sum_{i=1}^3 - \left[\frac{200 \sin\left(\frac{\pi(X_1 - m_{i,1})}{L}\right)}{100 \sin\left(\frac{\pi(X_2 - m_{i,2})}{L}\right)} \right] e^{-\frac{2(X - m_i)^T (X - m_i)}{L^2}}. \quad (29)$$

To combine the stable attractors and limit cycle attractors in a smooth fashion we adopted our design strategy of filtering out regions of the limit cycle attractor in the stable attractors regions around m_1, m_2 , and m_3 , and added in the stable attractors. Finally, the state-independent control signal was added linearly to give the controlled dynamics with $\Delta t = 10^{-4}s$ as

$$\mathcal{F}_{\text{diseased}}[X, u] = \Delta t (\mathcal{F}_{\text{limit cycle}} \Pi_{i=1}^3 B_i(X) + \mathcal{F}_a) + u(t) \quad (30)$$

$$B_i(X) = 1 - \exp\left(-\frac{2(X - m_i)^T (X - m_i)}{(1.5L)^2}\right). \quad (31)$$

The healthy dynamics were designed by using the approach of Section to encourage the dynamics to stay near stable attractors, and avoid the limit-cycle attractor. We designed a hyperbolic tangent filter function FF preserve the stable points, and a barrier function $\mathcal{B}$ to encourage state movement away from the limit cycle

$$FF = \frac{1}{2} \begin{bmatrix} \tanh(a(x - m_{i,1})) + 1 & 0 \\ 0 & \tanh(a(x - m_{i,1})) + 1 \end{bmatrix} \quad (32)$$

$$\mathcal{B} = \frac{1}{2} \begin{bmatrix} \tanh(-a(x - m_{i,1})) + 1 \\ 0 \end{bmatrix}. \quad (33)$$

Both the filtering function and barrier function have their zeros at the intersection of the stable points, to avoid introducing additional unwanted stable points. The scale

factor $a = 20\pi$ creates a steep barrier. The healthy dynamics are then calculated by filtering on null-controlled unhealthy dynamics as

$$\mathcal{F}_{\text{healthy}} = FF(\mathcal{F}_{\text{diseased}}[X, u = 0]) + \mathcal{B}. \quad (34)$$

The state noise covariance $Q = 10^{-5}I_2$ was chosen to allow for noise-assisted departure of the uncontrolled dynamics from the stable points an into the limit cycle attractor. Observation noise covariances $R \in [10^{-6}I_2, 10^{-5}I_2, 10^{-4}I_2]$ were used.

Winner-take-all and robust neural integrator dynamics

The Winner-take-all dynamics are based upon the state-space description in [29], in which two sub-populations of excitatory neurons X_1 and X_2 have a reduced-state dynamical description for decision-making of the direction of a random moving dot visual stimulus. The dynamics for the two-dimensional state driven by control signals $u(t) = [u_1(t), u_2(t)]^T$ are given by

$$\frac{dX_i}{dt} = -\frac{X_i}{\tau_s} + \alpha(1 - X_i)H_i \quad (35)$$

$$H_i = \frac{ax_i - b}{1 - \exp[-d(ax_i - b)]} \quad (36)$$

$$\dot{x}_1 = J_{11}X_1 - J_{12}X_2 + I_0 + u_1(t) + I_1 \quad (37)$$

$$\dot{x}_2 = J_{22}X_2 - J_{21}X_1 + I_0 + u_2(t) + I_2. \quad (38)$$

The visual stimulus is represented as input current I_1 and I_2 to each population with stimulus strength $\mu_0 = 30Hz$ and directional percent of coherence c' ,

$$I_i = J_A\mu_0 \left(1 \pm \frac{c'}{100}\right). \quad (39)$$

High activity of a state X_1 corresponds to decision due to activity of that sub-population of neurons with positive-signed coherence of the stimulus, and X_2 alternative has high activity for negative-sign coherence of the stimulus, indicating the direction of the stimulus. Parameter values reported in [29] were used. The controls to the system $u_1(t)$ and $u_2(t)$ were modeled as additional input currents to the sub-populations.

The target neural dynamics of a robust neural integrator are based conceptually upon [28], and their state space description is modeled as a set of hyperbolic tangents that generate interwoven nullclines. The two states X_1 and X_2 have gradients given by

$$\begin{aligned} \frac{dX_i}{dt} &= \tau \left[\tanh\left(\frac{\pi}{a}[\pm u - 2(\tilde{X}_2 - 0.5)]\right) + BC_i \right] + I_i \\ u &= a \sin\left(\frac{(n-1)\pi}{L}\tilde{X}_1 + \pi\right), \quad i \in [1(+), 2(-)]. \end{aligned} \quad (40)$$

The nullcline shapes u are defined by $n - 1$ nodes over a length L , and have a hyperbolic tangent on each side in state space. (40) use a rotated set of state-space coordinates $\tilde{X}_i$ given by a rotation matrix $[\tilde{X}_1, \tilde{X}_2]^T = M([X_1, X_2]^T - R_0)$, where M rotates by the angle $\pi/4$ and $R_0 = [L/2, 1]^T$. The boundary conditions BC_1 and BC_2 enforce the final fixed points of the neural integrator line to be global attractors, and are given by additional hyperbolic tangents of the form

$$BC_1 = 10 \tanh \left(-\frac{(n-1)\pi}{b}(X_1 + c) \right) - 10 \tanh \left(-\frac{(n-1)\pi}{b}(X_1 - [c + L + a/4]) \right) \quad (41)$$

$$BC_2 = 10 \tanh \left(-\frac{\pi}{d}(X_2 + c) \right) + 10 \tanh \left(-\frac{\pi}{d}(X_2 - [L + c]) \right). \quad (42)$$

The parameters of the model were chosen to roughly match the magnitude and fixed-point locations of the winner-take-all dynamics. $\tau = 1e^{-3}$ (*i.e.*, 1ms timesteps), $a = 0.2$, $n = 7$, $L = 0.7$, $b = 4/3$, $c = 0.083$, $d = 1.2$. The stimuli to the robust neural integrator I_1 and I_2 were given by

$$[I_1, I_2] = \text{sign}(c')[\delta(1 + |c'|/100), -\delta(1 + |c'|/100)], \quad (43)$$

where $\delta = 7.5e - 4$. The state noise covariance $Q = 5 \times 10^{-5}I_2$ was chosen to allow the robust neural integrator to utilize the state noise to transition from one stable point to another, and observation noise covariances $R \in [10^{-6}I_2, 10^{-5}I_2, 10^{-4}I_2]$ were used.

Results

In the following examples we demonstrate the ability of myopic control to match the dynamics of several relevant dynamical systems for neural computations. Simulations to benchmark the performance of myopic control were conducted using Tensorflow (Python API). System details can be found in the methods section, and code for the myopic controller is available at <https://github.com/catnip-lab/myopiccontrol>.

Robust neural integration from winner-take-all dynamics

We first deal with controlling neural computations for decision making, and demonstrate how myopic control can be used to change a winner-take-all (WTA) decision-making dynamics and convert it into a robust neural integrator (RNI). WTA dynamics for a simple, forced two-choice decision-making process function through a dynamical system where stimulus modulates the dynamics to flow toward one of two stable attractors. As time progresses, the neural state is driven toward one of the two stable attractors, each comprising a separate decision. In contrast, a robust neural integrator has multiple fixed points in between those two final stable attractors that allow for a stable, intermediate representation of accumulated evidence—creating robustness against uncertainty in stimulus and small internal perturbations.

We implemented a well-known approximation of a WTA dynamical system underlying two populations of spiking, excitatory neurons connected through strong recurrent inhibitory neurons, and our control for this system is an external injected current into each excitatory population [29]. Our target dynamics embody a low-dimensional analogue of the robust neural integration model suggested by Koulakov and coworkers [28]. Our RNI dynamical system is conceptually quite simple: Two sinusoidal nullclines that are interwoven can generate alternating stable and unstable fixed points, and with the addition of boundary conditions on the final stable fixed points can generate a dynamical system with a line of stable fixed points. The phase portraits for each system are shown in Figure 3.

Fig 3. Phase portraits for the A) winner-take-all dynamics and B) robust neural integrator. The magnitude is plotted on a logarithmic scale for easier visualization, and arrows give gradient direction. Streamlines are depicted in cyan. The nullclines of each dynamical system are shown in black. The RNI system is formulated as an extension of the tangling of nullclines in the WTA dynamics, where additional crossings of the nullclines result in stable points.

Table 1. Accuracy of decision making for an uncertain stimulus .

Dynamics	Accuracy	% decisions made	control power P (mean $\pm$ std)
1. RNI	91 %	91 %	-
2. Myopic control	92 %	92 %	2.05 ± 0.87
3. Uncontrolled	41 %	89 %	-
4. Trajectory control	0 %	99 %	1.65 ± 0.42

A comparison between first-order myopic control and a trajectory control approach (*i.e.*, a control optimized for (4)) is presented in Figure 4. Specifically, we show that trajectory control possess shortcomings when dealing with particular decision-making tasks. Trajectory control approaches have an additional hurdle above myopic control in that there must be some policy in place to decide to which target $\mathbf{x}_t$ should evolve. The only way that trajectory control can help the neural system make an informed decision is by integrating the stimulus, itself, which assumes a role in the neural computation. Here, we allow the controller to observe and integrate 20ms of coherence at the beginning of the trial, and then use that information to prescribe a target point in state space (either the final + or - coherence decision points in Fig. 3). In this simulation the time-varying stimulus c' is initially uncertain, beginning with a small positive coherence for 500 ms and then changing to negative coherence for 500 ms, finally settling to a stronger negative coherence of $c' = -12\%$ for the remainder of a 1s trial, shown inset in Fig. 4. Both the target dynamics of robust neural integrator and myopic control to mimic it can adequately handle this “change-of-mind” in stimulus and eventually evolve its neural state to the negative coherence choice, but trajectory control system instead incorporates only the initial stimulus to incorrectly choose the positive coherence choice. Furthermore, it then holds it there with control, in spite of receiving new stimulus information that in some cases could have even been caught in the WTA system (Fig. 4, green).

An intuitive way to compare the performance of trajectory control vs. myopic control for decision making is to count the number of correct decisions made, which is summarized in Table 1, alongside the total power of the controls, calculated as

$$P = \sum_{i,m} |u_m(t_i)|^2. \quad (44)$$

Accuracy for each control type calculated as percentage of correct fixed points chosen (noted in Fig. 3), and percentage decided as number of trials in which the state evolved to within a close radius of a decision point (radius = 0.15). Examining the accuracy of each method in the table, it is clear that this is an extreme case of how poorly trajectory control can perform, where even uncontrolled dynamics sometimes was able to change its mind and choose the correct stimulus. This is a specific example of our qualitative arguments against trajectory control that were shown in Fig. 1, in which markedly different behavior can be artificially enforced. Moreover, the total power required by the control signals is comparable, indicating that myopic control did not require substantially more power to perform the target dynamics

Quantitative performance of myopic control of a sample of 500 trials is summarized

Fig 4. Example trajectories of controlled dynamics for an initially uncertain stimulus (inset). Trajectory controls with a policy to decide which decision to make by integrating the initial few ms of coherence signal not only steal the role of computational integration from the dynamical system, but in the presence of an initially uncertain stimulus perform poorly by making the wrong decision (red) where even a WTA system (green) can somewhat cope with a changing stimulus. Myopic control (blue) can integrate the incoming information just as the target RNI dynamics (black). inset: coherence for the change of mind stimulus. a final decision would be made by integrating the signal, signified by the gray region. Vertical line at 20ms indicates the portion of the signal integrated to decide on a target point for trajectory control.

in Figure 5. Sample trajectories for uncontrolled, first-order controlled, and healthy dynamics are shown in Fig. 5A under the influence of an increasingly stronger time-dependent stimulus, denoted by coherence c' . Both the RNI and controlled system linger at an intermediate stable nodes before coherence has increased enough to make a more informed decision, indicated by the progression to the decision node. In contrast, the uncontrolled trajectory evolves straight to the decision without any intermediate stability at low coherence.

Fig 5. A) Example trajectory of controlled dynamics for evidence accumulation. Nullclines and stable points (dots) of the neural integrator shown in black, with unstable nodes shown as diamonds. B) Average time dependence of log-performance of expected cost in eq. (20) for winner-take-all to robust neural integrator dynamics with first-order control, SNR of 7dB, and fixed coherence $c' = -6\%$. Example trials are shown in gray, and trial-averaged mean plotted in black. C) Control signals during the evidence accumulation with first-order control. D) Similar control signals for fixed-coherence trial. E) Violin plots showing distribution of time-averaged, log-performance for a lag in observations. F) Similar violin plots as in E), but for for varying observation noise strength, shown by signal-to-noise ratios (SNR) of system noise to observation noise, and including null control ($c_t = 0$). In both E) and F) * indicates $p < 0.001$ for a two-sample t-test. Unless otherwise noted, samples are not significantly different. ** : $p = 0.009$.

Importantly, the controlled dynamics demonstrate the intermediate stability behavior found in robust neural integrators. Figure 5B summarizes the log-cost of 500 trials of first-order control with a fixed stimulus coherence of $c' = -6\%$, where prototypical trials are shown in gray alongside the trial average in black. The control signals for the increasing coherence demonstration (Fig. 5A) and for the benchmark trajectories (Fig. 5B) are plotted in 5C and 5D, respectively. Again, promising and modest control amplitudes are observed in both cases. Finally, figures 5E and 5F show the time-averaged, log-performance for varying observation lag and noise strengths. Comparable to the previous section we see that second order control performs equivalently to first order across increasing observation lag and noise. However, the time-averaged distributions at low observation lag have quite long-tailed, unimodal distributions, and have negligible performance at a lag of 50 steps (note that Δt is an order of magnitude higher for this system, which corresponds to a 50ms-ahead prediction). There is some change to a bimodal distribution for increasing observation noise in this system, but the notable feature is the increasingly long distribution tail for second-order control, which gives the opportunity for inferior performance as compared to first-order control.

Avoiding beta-oscillation disease states

Here, we aim to preserve an original set of dynamical features in $\mathcal{F}$ while avoiding an unwanted regime of state space containing undesirable dynamics. This paradigm can act as the basis of state-space control for neurological disorders, where regions of state space may be associated with disease symptoms [30,31]. Utilizing myopic control as a therapy for neurological disorders lends itself to considering which features of neural dynamics are *undesirable*, rather than discerning which features of the dynamical system are lacking. For example, tremors in Parkinson’s disease (PD) are associated with a characteristic beta oscillation (*i.e.*, 13-30 Hz) of the local field potential in the subthalamic nucleus, and state-of-the art feedback control strategies use this signal to trigger deep brain stimulation (DBS) until the beta oscillation subsides [18]. Similar neural signatures are also present for epilepsy [14,15]. A model “diseased” system with three stable fixed points representing three possible voluntary movement command was constructed with an additional, unwanted spiral attractor representing the beta oscillation macrostate. Difficulty in initiating voluntary motion (bradykinesia) in PD patients could be due to strong attractive macrostate [32,33]. Using myopic control, we manipulated the dynamics to match the target dynamics of a healthy system structured using the design principles from section to avoid the avoid beta oscillation state while preserving the fixed points of the system. The phase portrait of the target dynamics are shown in Figure 6.

Fig 6. Phase portraits for diseased and healthy dynamical systems. Color denotes the magnitude of the dynamics $\mathcal{F}$ and $\mathcal{G}$, and the direction is shown by the arrows. Streamlines are shown in cyan. The diseased system contains three stable points, but also a spiral attractor at small values of X_1 and X_2 . The target healthy dynamics has been designed to contain slightly repulsive dynamics in the original spiral attractor region, but still maintains its stable points.

The overall performance of myopic control is summarized in Fig. 7. A sample of 500 trials points was initialized in the asymptotic distribution of the PD limit-cycle attractor, and state estimation was performed for 100ms in the absence of control before the control was switched on. Monitoring $\log[\tilde{J}_t]$ in Figure 7A, individual trials reflect the initial oscillatory behavior of being in the disease state before sharply declining, whereas the trial-averaged behavior shows the overall improvement due to control. This remarkable removal of disease-state behavior is further demonstrated in state-space trajectory of a typical trial (Fig. 7B). Once control is switched on the target dynamics successfully lead it out of the limit cycle and into a stable attractor point. A spectrogram of the state X_1 for an analogous, longer simulation is shown in Fig. 7 D. There, a beta oscillation endured for 1s, and then myopic control was switched on to evolve to a stable point. The spectrogram reflects the oscillations during the uncontrolled period, and once the control is switched on it subsides and leaves only low-frequency components as it moves toward the stable point. The optimal control signal for the colored trajectory in Fig7A in shown in Fig. 7C. it is modest in amplitude relative to the magnitude of the dynamics, and has a straightforward waveform, demonstrating that given only minimal additional consideration to constraints on the control signal that myopic control could feasibly, efficiently, and safely be implemented in living subjects.

Finally, we benchmarked the performance of first- and second-order control as compared to uncontrolled dynamics by calculating distributions of the time-averaged log-cost $\sum_{i=1}^T \log(\tilde{J}_t)$ for varying lags between state observation and control signal calculation (Fig. 7E), and for different observation noise strength (Fig. 7F). While there is a interesting trend in the stretching of bimodal distribution into a near unimodal one

Fig 7. A) Time-dependent performance of first-order control for small observation noise (SNR 10dB). Instances of simulations are shown in gray, while the mean behavior is given in black. A representative trial has been singled out in blue and red for additional analysis. Initial trajectories are uncontrolled (blue), and allowed to fall into the asymptotic distribution of the limit cycle before control is switched on (red). Costs are plotted as time-locked 100ms before the control is switched on. B) Trajectory of typical trial through state space shown before and after the control is implemented, demonstrating a move back towards healthy state space. C) Control signal of a representative trial. D) Spectrogram of X_1 for an analogous, longer simulation of uncontrolled evolution for 1s (noted by vertical line), and myopic control for final 1s . E) Violin plots showing distribution of time-averaged log-performance for a lag in observations, requiring state prediction. Horizontal line corresponds to the average log-loss of null control. F) Similar violin plots as in E), but for for varying observation noise strength, and including null control ($u_t = [0, 0]$). In both E) and F) * indicates $p < 0.001$ for a two-sample t-test. Unless otherwise noted, samples are not significantly different.

at high observation lag, we observed no impactful difference between first- and second-order control with an increasing delay of observations. Similarly, there is a transition to a distinct bimodal distribution at large signal to noise, though both controllers perform similarly.

Discussion

Here we developed a perspective on what features are necessary for a flexible control of any dynamical system underlying neural computation. The controller should function to assist the dynamical system performing the computation, not taking on the role of a dynamical system, itself. In order for the controlled dynamics to function as a separate dynamical system on its own, we proposed a myopic control scheme that alternatively manipulates the dynamics to function as a set of target dynamics over a single time step, as opposed to trajectory-tracking controllers that function over a finite time horizon and must first perform the neural computation on their own. We developed an approximation of this control for nonlinear dynamics that is separable from state estimation, provided direction about design principles for how to construct a targeted dynamical system, and demonstrated its application in two varied scenarios. In both examples, first order control performed comparably to second-order control, showing the potential to generate feasible control signals that function under practical conditions.

The base of our controller formulation is reminiscent of the model-based control and model reference (adaptive) control (MRAC). Utilizing model-based control alongside quality state estimation [30] to manipulate neural dynamics is an attractive strategy that can harness machine learning methods to build effective, patient-specific statistical models of the brain by using real-time patient data, which could then be used as precision medical treatment [34, 35]. The initial efforts of MRAC focused heavily on adaptive update rules for estimating the parameters of different forms of target plants (dynamics, in our work), and predominantly one adaptive controller form was utilized: strictly positive real (SPR) Lyapunov design. This form of controller depended on a SPR transfer function formulation of its plant dynamics, as was designed to guarantee bounded control signals that can track target trajectories or regulate to a fixed point from a target plant. Our controller structure is similar in form to SPR control and could benefit from the similar extensions that took place in MRAC, such as analysis of the Lyapunov stability to rigorously establish safe bounds on the control [19] and the

use of neural networks capable of handling nonlinear plant dynamics [36, 37].

This is not the first neural controller to consider neural variability as an important component to preserve in neural systems. Todorov and Jordan suggested a “minimal intervention principle” for neural systems that allows for deviations from a target trajectory, provided that they do not interfere with the target task [38]. The target was considered as a single point in state space, and their formulation allowed for high redundancy in the number of optimal trajectories that reached the target with the same cost. Their controller only corrects the trajectory when failing to act would result in a worse-than-optimal cost. While this is the only instance of control that acknowledges and respects neural variability during control, even prescribing a single point in state space as a target falls short of the general goals accomplished by myopic control to generate an entire target dynamics. For example, returning to the qualitative operation of myopic control in Fig. 1B, minimal intervention control would perform comparably to trajectory control by forcing state evolution in a non-dynamical fashion, while also restricting the neural variability that lead to an alternative fixed point.

An important feature of myopic control was its modular design. We studied a form in which only low-order moments of the state distribution were used in the controller, which decoupled the controller form from state estimation and allowed for any state estimator to be implemented. First-order control is considerably more straightforward to use because of its lack of higher-order derivatives on the dynamics, which may also come with a benefit being a more robust controller during practical instances in which the dynamics must be inferred from data. Operating in the perturbative regime $x - \hat{x}_t$ through regular state estimation small would ensure that first-order methods are successful.

A key issue that dictates myopic control’s qualitative success is the choice or design of target dynamics. Target dynamics could be designed by modifying the current dynamics through either addition, removal, or modification of specific features in the state space. There is an appeal to omitting features, as this approach resembles lesioning studies that aim to infer causal importance to behavior. In our beta oscillation example, omitting the limit cycle appeared to be the safest and most practical approach. However, if more experimentally accurate understanding of the Parkinsonian dynamics suggests that omitting a limit cycle could introduce unwanted behavior, it may be more prudent to modify the limit cycle with an exit pathway. Still, one may wish to study the extent of a neural system’s computational flexibility by adding features as we did with our decision-making example.

It should be noted the adding in new features can be a tedious process in practice, as it took considerable parameter tuning to create our robust neural integrator system. Additionally, we considered only two-dimensional systems, but interesting dynamical systems may in fact lie in higher dimensions [8]. Our design approach of filtering functions is general enough to extend to high dimensions, though implementing them in practice may take additional care. Removing features with smoothed versions of step functions would still work, as would adding stable points with gabor functions, though they would need to be high-dimensional analogues. Visualization and analysis tools for high-dimensional spaces could help determine the hypervolumes to omit, and how (an)isotropic the features must be.

Myopic control is certainly not the only form of controlled stimulation of neural systems, and it is important to note how these methods differ from our perspective. One of most successful applications of neural stimulation is in the field of neuroprosthetics, where implants mimic afferent sensory inputs such as cochlear or retinal implants, or translate efferent outputs into motor actions for artificial limbs [39]. The control strategies behind these technologies are complex and varied compared to myopic control [20], though this is required in part because the goal of neuroprosthetics

is distinctly different: the controlled (de)coding of these neural signals do not constitute a dynamical system, but rather interacts with the pre-existing normal neural dynamics of the area as inputs. Neural prosthetics for cognitive function, for example, memory processing in hippocampus [40], are much more amenable to myopic control scheme, since the normal function of the neural system constitutes a dynamical system.

Deep brain stimulation (DBS) for neurological disorders (*e.g.*, Parkinson’s disease) is a control application within the scope of myopic control, as we demonstrated with our first example study. A recent approach to DBS that harnesses neural recordings uses a model-free method to simply reduce beta-band oscillations seen in local field potential recordings in the basal ganglia [18], a potential neural signal related to PD symptoms [41,42]. The disadvantage to such a heuristic approach is that the link between beta oscillations in basal ganglia and cortex, let alone its relationship to actual PD symptoms, is still not fully understood. Moreover, other feedback targets are being actively considered as well [17,31]. Myopic control allows us to causally investigate the role of neural signatures correlated with the disease—we can specifically target fixes to the abnormal dynamics for beta oscillations, for example, and improve our understanding of the disease and also improve treatments.

Our first example was motivated from a position of understanding neural dynamical systems for evidence accumulation and decision making, and more generally to demonstrate its application as a tool to causally investigate cognitive processes. Several models of evidence accumulation have been considered in the context of using variability in spiking dynamics of lateral intraparietal cortex (LIP) in monkeys [43–45], and one future experiment could attempt myopic control using different models for the control systems to produces a given target system, say RNI. Performing myopic control in that context would be a more powerful approach than perturbative, random stimulation of the system to simply infer parameters of an underlying dynamical model represented in LIP. Additionally, a more sophisticated experiment could attempt to utilize controlled stimulation to force the opposite decision of a target dynamics; the success of which would not only provide evidence that the controller is operating based upon the correct dynamical systems model, but would also constitute a substantial advance in the control of cognitive dynamics.

The history of advances in model reference adaptive control (MRAC) provides a strong template for how myopic controllers for neural dynamics control could be developed. Our work here assumed a known model for the controlled dynamics, and future work should integrate adaptive estimation of the controlled dynamics, themselves, into the controller. In particular, as an extension of the initial neural network structures used to perform MRAC [36,37], there is opportunity to utilize deep networks that accomplish adaptive estimation these dynamics and their states within a neural-network myopic controller architecture [46–48].

A larger and more immediate question is what steps must be taken to implement myopic control experimentally? The most important underlying component is access to quality neural measurements. That is why recent work combining neural stimulation and observation as in [49] is so vital. In our work we assumed that the ground-truth neural dynamics for both the controlled dynamics $\mathcal{F}$ and the target dynamics $\mathcal{G}$ were known, but in practice they must estimated from neural measurements. We demonstrated that first-order myopic control can function well, which necessitates estimation of only the state mean over higher order moments, but first order control also requires estimating the full dynamics in (18).

The timescale of the underlying neural computation also suggests practical consideration. Since myopic control is designed as an online control, the state estimations and estimation of the dynamics must be fast in order to implement in real time. Longer time constants for processes that are characterized by a smaller total

dynamics $\mathcal{F}_t$ lead to slower changes in neural state, which allow for more accurate online state estimation, and thus a more accurate control signal. Akin to a slower moving target in state space, the less the dynamics have progressed, the more up-to-date that state information will be, and the better the control performance for slower dynamical processes. This motivated our demonstration that myopic control can still function well with a lag between neural observations and control implementation.

Moreover, estimating latent state dynamics is a difficult task altogether [3], and would likely need to be performed prior to control use, with adaptive updates to the dynamics estimation occurring online. Taking into consideration a generic framework i) signal processing (*e.g.*, spike sorting), ii) control signal calculation, and iii) delivery of stimulation; it seems reasonable to assume ~ 5 ms of time required for myopic control, which is comparable to other closed-loop control estimates [50]. In this regime of time lags of less than 5 ms, myopic control was demonstrated to perform well, which is promising for its implementation.

Acknowledgments

This work was supported by the Thomas Hartman Center for Parkinson's Research (64249).

References

1. Brain Research through Advancing Innovative Neurotechnologies (BRAIN) Working Group Report to the Advisory Committee to the Director, NIH. National Institute of Health; 2014.
2. Jun JJ, Steinmetz NA, Siegle JH, Denman DJ, Bauza M, Barbarits B, et al. Fully integrated silicon probes for high-density recording of neural activity. *Nature*. 2017;551(7679):232–236. doi:10.1038/nature24636.
3. Breakspear M. Dynamic models of large-scale brain activity. *Nat Neurosci*. 2017;20(3):340–352.
4. Fairhall A, Machens C. Editorial overview: Computational neuroscience. *Current Opinion in Neurobiology*. 2017;doi:https://doi.org/10.1016/j.conb.2017.09.009.
5. Sussillo D. Neural circuits as computational dynamical systems. *Current Opinion in Neurobiology*. 2014;25:156 – 163. doi:https://doi.org/10.1016/j.conb.2014.01.008.
6. Izhikevich EM. *Dynamical Systems in Neuroscience: The Geometry of Excitability and Bursting*. MIT Press; 2006.
7. Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, et al. Neural population dynamics during reaching. *Nature*. 2012;487(7405):51–56.
8. Park IM, Meister MLR, Huk AC, Pillow JW. Encoding and decoding in parietal cortex during sensorimotor decision-making. *Nat Neurosci*. 2014;17(10):1395–1403.
9. Mante V, Sussillo D, Shenoy KV, Newsome WT. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*. 2013;503(7474):78–84.

10. Jazayeri M, Afraz A. Navigating the Neural Space in Search of the Neural Code. *Neuron*. 2017;93(5):1003–1014. doi:10.1016/j.neuron.2017.02.019.
11. Li HH, Rankin J, Rinzel J, Carrasco M, Heeger DJ. Attention model of binocular rivalry. *Proceedings of the National Academy of Sciences*. 2017;114(30):E6192–E6201. doi:10.1073/pnas.1620475114.
12. Deuschl G, Schade-Brittinger C, Krack P, Volkmann J, Schäfer H, Bötzel K, et al. A Randomized Trial of Deep-Brain Stimulation for Parkinson's Disease. *New England Journal of Medicine*. 2006;355(9):896–908. doi:10.1056/NEJMoa060281.
13. Lyons KE, Pahwa R. Deep brain stimulation and essential tremor. *J Clin Neurophysiol*. 2004;.
14. Handforth A, DeGiorgio CM, Schachter SC, Uthman BM, Naritoku DK, Tecoma ES, et al. Vagus nerve stimulation therapy for partial-onset seizures: A randomized active-control trial. *Neurology*. 1998;51(1):48–55. doi:10.1212/WNL.51.1.48.
15. Morrell MJ. Responsive cortical stimulation for the treatment of medically intractable partial epilepsy. *Neurology*. 2011;77(13):1295–1304. doi:10.1212/WNL.0b013e3182302056.
16. Ineichen C, Baumann-Vogel H, Christen M. Deep Brain Stimulation: In Search of Reliable Instruments for Assessing Complex Personality-Related Changes. *Brain Sciences*. 2016;6(3):40. doi:10.3390/brainsci6030040.
17. Rosin B, Slovik M, Mitelman R, Rivlin-Etzion M, Haber SN, Israel Z, et al. Closed-Loop Deep Brain Stimulation Is Superior in Ameliorating Parkinsonism. *Neuron*. 2011;72(2):370 – 384. doi:http://dx.doi.org/10.1016/j.neuron.2011.08.023.
18. Malekmohammadi M, Herron J, Velisar A, Blumenfeld Z, Trager MH, Chizeck HJ, et al. Kinematic Adaptive Deep Brain Stimulation for Resting Tremor in Parkinson's Disease. *Movement Disorders*. 2016;31(3):426–428. doi:10.1002/mds.26482.
19. Ioannou P, Sun J. *Robust Adaptive Control*. Dover; 2012.
20. Schiff SJ. *Neural Control Engineering: The Emerging Intersection between Control Theory and Neuroscience*. MIT Press; 2011.
21. Yang Y, Shanechi MM. An adaptive and generalizable closed-loop system for control of medically induced coma and other states of anesthesia. *Journal of Neural Engineering*. 2016;13(6):066019.
22. Newman JP, Fong Mf, Millard DC, Whitmire CJ, Stanley GB, Potter SM. Optogenetic feedback control of neural activity. *eLife*. 2015;4:e07192. doi:10.7554/eLife.07192.
23. Little S, Pogosyan A, Neal S, Zavala B, Zrinzo L, Hariz M, et al. Adaptive deep brain stimulation in advanced Parkinson disease. *Ann Neurol*. 2013;74(3):449–457. doi:10.1002/ana.23951.
24. Freeman WJ. *Mass Action in the Nervous System*. Academic Press; 1975. Available from: <http://www.amazon.com/exec/obidos/redirect?tag=citeulike07-20&path=ASIN/0124120474>.

25. Haykin S, Principe J. Making sense of a complex world [chaotic events modeling]. *IEEE Signal Processing Magazine*. 1998;15(3):66–81. doi:10.1109/79.671132.
26. Ogata K. *Modern control engineering*. Prentice Hall; 2010. Available from: <http://www.worldcat.org/isbn/9780136156734>.
27. Stengel E. *Optimal Control and Estimation*,. Dover; 1994.
28. Koulakov AA, Raghavachari S, Kepecs A, Lisman JE. Model for a robust neural integrator. *Nature Neuroscience*. 2002;5(8):775–782. doi:10.1038/nn893.
29. Wong KF, Wang XJ. A Recurrent Network Mechanism of Time Integration in Perceptual Decisions. *The Journal of Neuroscience*. 2006;26(4):1314–1328. doi:10.1523/jneurosci.3733-05.2006.
30. Watter M, Springenberg J, Boedecker J, Riedmiller M. Embed to Control: A Locally Linear Latent Dynamics Model for Control from Raw Images. In: Cortes C, Lawrence ND, Lee DD, Sugiyama M, Garnett R, editors. *Advances in Neural Information Processing Systems 28*. Curran Associates, Inc.; 2015. p. 2746–2754. Available from: <http://papers.nips.cc/paper/5964-embed-to-control-a-locally-linear-latent-dynamics-model-for-control-pdf>.
31. Little S, Brown P. What brain signals are suitable for feedback control of deep brain stimulation in Parkinson’s disease? *Annals of the New York Academy of Sciences*. 2012;1265(1):9–24. doi:10.1111/j.1749-6632.2012.06650.x.
32. Nini A, Feingold A, Sloviter H, Bergman H. Neurons in the globus pallidus do not show correlated activity in the normal monkey, but phase-locked oscillations appear in the MPTP model of parkinsonism. *Journal of Neurophysiology*. 1995;74(4):1800–1805.
33. Little S, Brown P. Focusing Brain Therapeutic Interventions in Space and Time for Parkinson’s Disease. *Current Biology*. 2014;24(18):R898–R909. doi:10.1016/j.cub.2014.08.002.
34. Collins FS, Varmus H. A New Initiative on Precision Medicine. *New England Journal of Medicine*. 2015;372(9):793–795. doi:10.1056/NEJMp1500523.
35. Ozomaro U, Wahlestedt C, Nemeroff CB. Personalized medicine in psychiatry: problems and promises. *BMC Medicine*. 2013;11(1):132. doi:10.1186/1741-7015-11-132.
36. Narendra KS, Parthasarathy K. Identification and control of dynamical systems using neural networks. *IEEE Transactions on Neural Networks*. 1990;1(1):4–27. doi:10.1109/72.80202.
37. Hagan MT, Demuth HB, Jesús OD. An introduction to the use of neural networks in control systems. *International Journal of Robust and Nonlinear Control*. 2002;12(11):959–985. doi:10.1002/rnc.727.
38. Todorov E, Jordan MI. A Minimal Intervention Principle for Coordinated Movement. In: Becker O Thurn, editor. *Advances in Neural Information Processing Systems 15*; 2013. p. 27–34. Available from: <https://homes.cs.washington.edu/~todorov/papers/TodorovNIPS03.pdf>.
39. Sanchez J. *Neuroprosthetics: Principles and Applications*. CRC Press; 2015.

40. Berger TW, Hampson RE, Song D, Goonawardena A, Marmarelis VZ, Deadwyler SA. A cortical neural prosthesis for restoring and enhancing memory. *Journal of Neural Engineering*. 2011;8(4):046017+. doi:10.1088/1741-2560/8/4/046017.
41. Quinn EJ, Blumenfeld Z, Velisar A, Koop MM, Shreve LA, Trager MH, et al. Beta oscillations in freely moving Parkinson's subjects are attenuated during deep brain stimulation. *Movement Disorders*. 2015;30(13):1750–1758. doi:10.1002/mds.26376.
42. Trager MH, Koop MM, Velisar A, Blumenfeld Z, Nikolau JS, Quinn EJ, et al. Subthalamic beta oscillations are attenuated after withdrawal of chronic high frequency neurostimulation in Parkinson's disease. *Neurobiology of Disease*. 2016;96:22–30. doi:10.1016/j.nbd.2016.08.003.
43. Churchland AK, Kiani R, Chaudhuri R, Wang XJ, Pouget A, Shadlen MN. Variance as a Signature of Neural Computations during Decision Making. *Neuron*. 2011;69(4):818 – 831. doi:http://dx.doi.org/10.1016/j.neuron.2010.12.037.
44. Huk AC, Shadlen MN. Neural Activity in Macaque Parietal Cortex Reflects Temporal Integration of Visual Motion Signals during Perceptual Decision Making. *Nature Neuroscience*. 2005;25(45):10420–10436. doi:10.1523/JNEUROSCI.4684-04.2005.
45. Resulaj A, Kiani R, and Wolpert DM, Shadlen MN Changes of mind in decision-making. *Nature*. 2009;461(8):263+. doi:https://doi.org/10.1038/nature08275.
46. Zhao Y, Park IM. Recursive Variational Bayesian Dual Estimation for Nonlinear Dynamics and Non-Gaussian Observations; 2017.
47. Zhao Y, Park IM. Variational Latent Gaussian Process for Recovering Single-Trial Dynamics from Population Spike Trains. *Neural Computation*. 2017;29(5).
48. Sussillo D, Jozefowicz R, Abbott LF, Pandarinath C. LFADS - Latent Factor Analysis via Dynamical Systems. 2017;.
49. Tan H, Debarros J, Poghosyan A, Aziz TZ, Huang Y, Wang S, et al. Decoding voluntary movements and postural tremor based on thalamic LFPs for closed-loop stimulation for essential tremor. *bioRxiv*. 2018;doi:10.1101/436709.
50. Ciliberti D, Frederic M, Kloosterman F Real-time classification of experience-related ensemble spiking patterns for closed-loop applications. *eLife* 2018;7:e36275. doi:10.7554/eLife.36275.