Towards Closed-Loop Deep Brain Stimulation via Optimal Control

Malvern Madondo[1], Deepanshu Verma[2], Lars Ruthotto[2], Nicholas Au Yong[3]

[1]Department of Computer Science, [2]Department of Mathematics, [3]Department of Neurosurgery

Emory University

Introduction

Research Goal: Apply Optimal Control (OC) theory to Deep Brain Stimulation (DBS).

- DBS is a neurosurgical procedure that involves delivering electrical pulses to the brain via surgically implanted electrodes.
- Conventional DBS approaches (open-loop) rely on clinicians manually tuning the pulse generator through trial-and-error.

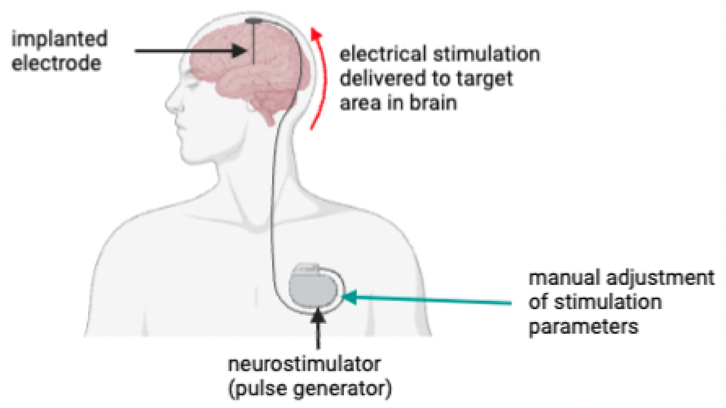

Fig. 1: Schematic of current state-of-the-art: open-loop DBS

Control Formulation

We consider the following dynamics

$$\frac{dz}{dt}(t) = f(t, z(t)) + e_1 u(t), \quad t \in [0, T], \; e_1 = [1, 0, 0, 0]^\top,$$

where

- $z(t) \in \mathbb{R}^4$ denotes the **state** of the system at time t
- $u(t)$ is the **control** (i.e., external current provided as input) applied at time t.
- f describes time evolution of state variable z according to the Hodgkin-Huxley neuronal model [1]:

$$f(t, z(t)) = \begin{pmatrix} -\big(I_{\mathrm{Na}}(t, z_0(t)) + I_{\mathrm{K}}(t, z_0(t)) + I_{\mathrm{L}}(t, z_0(t))\big) \\ \alpha_m(z_0(t))(1 - z_1(t)) - \beta_m(z_0(t))z_1(t) \\ \alpha_n(z_0(t))(1 - z_2(t)) - \beta_n(z_0(t))z_2(t) \\ \alpha_h(z_0(t))(1 - z_3(t)) - \beta_h(z_0(t))z_3(t) \end{pmatrix}.$$

where $I_{\mathrm{Na}}, I_{\mathrm{K}},$ and I_{L} are sodium, potassium, and leakage ion channel currents while α_x and β_x, for $x \in \{m, n, h\}$, are voltage-dependent rate constants.

Challenges & Promising Avenues

Large-scale Neuronal Dynamics

- Curse of Dimensionality arises as models increase in complexity and in solving HJB
  Difficulty: Makes direct solution intractable
  Remedy: Apply neural networks: scale well to high-dimensions. Or approximate large-scale dynamics with mean-field models.

Beyond Optimal Control

- Optimal Control approaches require full knowledge of system dynamics.
  Difficulty: Dynamics can be unknown/not fully accurate.
  Remedy: Consider combining with dynamics-agnostic approaches like Reinforcement Learning.

Optimal Control Problem

Goal: Find optimal control u_t^* that **minimizes** the **value function**

$$\Phi(t, z(t)) = \inf_u \left(\int_t^T L(s, z(s), u(s)) ds + G(z(T)) \right),$$

with running cost $L(t, z, u) = \frac{1}{2}\|u(t)\|^2$ and terminal cost $G(z(T)) = \frac{1}{2}\|z(T) - y\|^2$

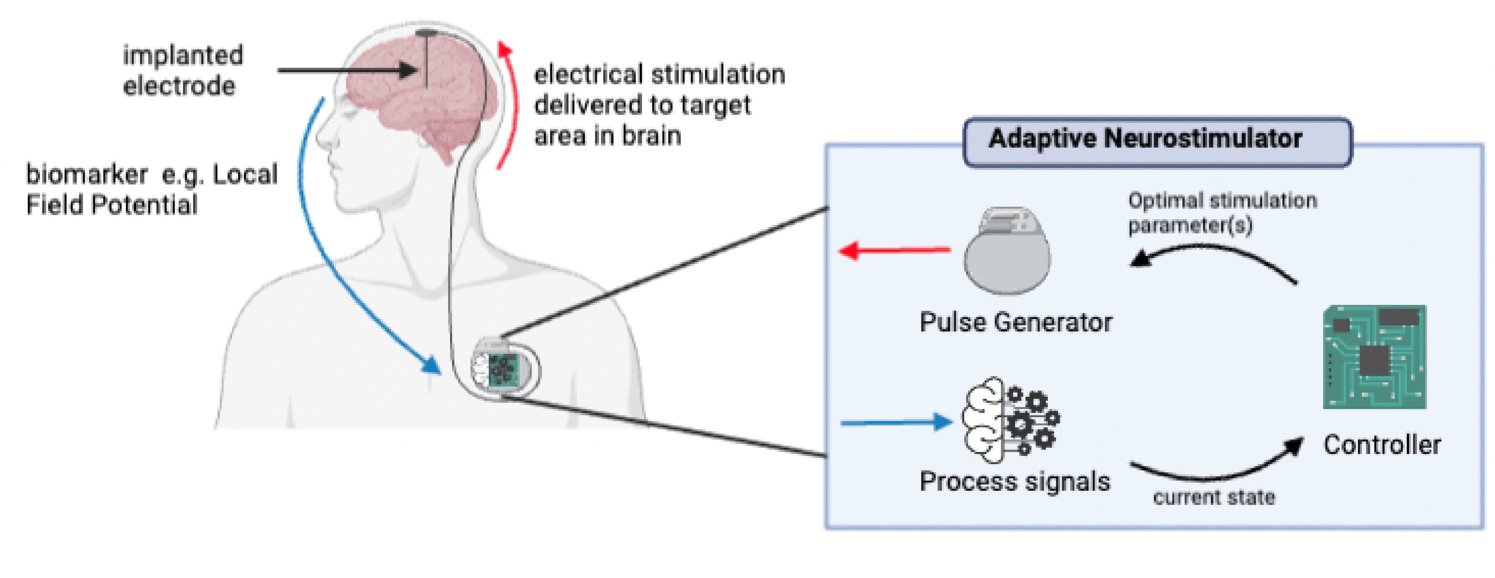

Fig. 2: Schematic of closed-loop DBS approach

Closed-loop DBS via OC

- Combine Pontryagin's Maximum Principle (PMP) and Hamilton Jacobi Bellman (HJB) approaches to yield a semi-global solution method as per [2, 3]
  - **Φ satisfies the HJB equation**

$$\begin{cases} -\partial_t \Phi(t, z(t)) + \sup_{u \in U} \mathcal{H}\big(t, z, \nabla_z \Phi(t, z(t)), u(t)\big) = 0, \\ \Phi(T, z(T)) = G(z(T)), \end{cases}$$

  where Hamiltonian $\mathcal{H}(t, z, p, u) = -\frac{1}{2}\|u(t)\|^2 - p^\top \cdot [f(z(t), t) + e_1 u(t)]$
  - **Optimality conditions**

$$\begin{cases} \partial_t p(s) = \nabla_z \mathcal{H}\big(s, z^*(s), p(s), u^*(s)\big), \\ \partial_s z^*(s) = -\nabla_p \mathcal{H}\big(s, z^*(s), p(s), u^*(s)\big) \\ z^*(t) = z, \; p(T) = \nabla_z G(z^*(T)) \end{cases}$$

  This system is solved by $p(s) = \nabla_z \Phi(s, z(s)), \quad t < s \leq T$. Optimal control recovered in **feedback form** as $u^*(s) \in \arg\max_u \mathcal{H}\big(s, z^*(s), \Phi(s, z^*(s)), u(s)\big)$
- Approximate Φ using a neural network parameterized by θ, denoted by Φ_θ

$$\Phi_\theta(x) = w^\top N(x; \theta) + \frac{1}{2} x^\top (A^\top A) x + b^\top x + c$$

  for space-time inputs $x = (t, z)$, weights $\theta = (w, A, b, c)$, and neural network $N(x; \theta)$
- Solve the **learning problem** i.e. find θ and weights of neural network by solving

$$\min_\theta \mathbb{E}_{x \sim \rho} \{\ell(T) + g(z(T)) + \beta_1 c_{\mathrm{HJ}}(T) + \beta_2 |\Phi_\theta(T, z(T)) - g(z(T))|\}$$

  subject to

$$\partial_s \begin{pmatrix} z(s) \\ \ell(s) \\ c_{\mathrm{HJ}}(s) \end{pmatrix} = \begin{pmatrix} -\nabla_p \mathcal{H}\big(s, z(s), \nabla_z \Phi(s, z(s)), u^*(s)\big) \\ L(s, z, u^*(s)) \\ \big| -\partial_t \Phi(s, z(s)) + \mathcal{H}\big(s, z, \nabla_z \Phi_\theta(s, z(s)), u^*(s)\big) \big| \end{pmatrix},$$

  initialized with $z(0) = x, \ell(0) = c_{\mathrm{HJ}}(0) = 0$ and for hyperparameters β_1 and β_2.

Preliminary Results

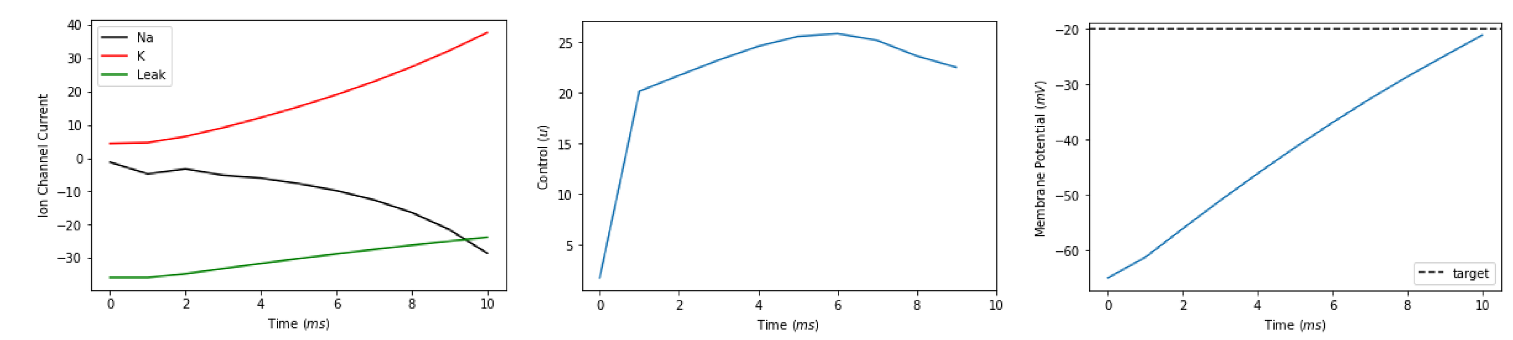

Fig. 3: Ion channel currents, membrane potential, and controls from local solution method with L-BGFS-B

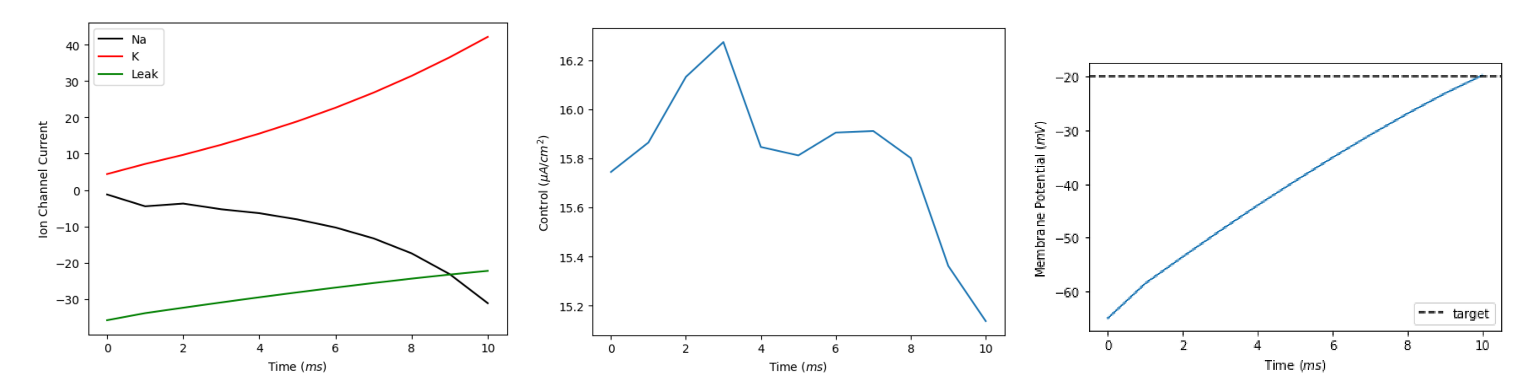

Fig. 4: Ion channel currents, membrane potential, and controls from global solution method from [2, 3]

Conclusion

- Formulated the problem of finding an optimal neurostimulation strategy as a control problem.
- Derived an optimal value function which satisfies the HJB equation and from which the optimal (stimulation) control can be recovered in feedback form.
- Established a concrete link between the learning problem and optimal control, specified by the PMP and HJB equation.

Funding Acknowledgements

Research supported by the 2021 Google PhD Fellowship award in Computational Neural and Cognitive Sciences, AFOSR: FA9550-20-1-0372 and DOE RISE: ASCR 20-023231

References

[1] Alan L Hodgkin and Andrew F Huxley. "A quantitative description of membrane current and its application to conduction and excitation in nerve". In: *The Journal of physiology* 117.4 (1952).

[2] Derek Onken et al. *A Neural Network Approach for Real-Time High-Dimensional Optimal Control.* 2021. arXiv: `2104.03270 [math.OC]`.

[3] Lars Ruthotto et al. "A machine learning framework for solving high-dimensional mean field game and mean field control problems". In: *Proceedings of the National Academy of Sciences* 117.17 (2020).

Fig. 5: Please visit `mmadondo.github.io` for updates and more info!