Getting What They Want Through Bayesian Modeling: Calibrating a Spinal Cord Stimulator for a Paraplegic Patient

Adam Kaplan, Thomas A. Murray
Division of Biostatistics, School of Public Health, University of Minnesota

Motivations

Compared to standard therapies, **spinal cord stimulation** (SCS) devices have been introduced as a promising alternative rehabilitation or therapy for chronic back pain and for spinal cord injury (SCI)-related paraplegia. These neurostimulators require the neurosurgeon to program combinations of frequency (in Hz) and pulse-widths (in micro-seconds) for the device, and to select where the device is implanted. One major constraint is that the device can store up to **8** parameter combinations for outside clinic use.

While past clinical trials have demonstrated the clinical efficacy of these devices as compared to standard physical therapies, they did not, however, *rigorously explore* the device's many parameter combinations, or device **configurations**, prior to the trial's commencement.

Therefore, we propose a year-long device **calibration phase** (i.e., an adaptive trial) for one patient with monthly follow-up visits; this allows for monthly reprogramming of 8 selected configurations for the patient to test in the following month.

Experimental Design Issues

With the restrictions that (i) the **patient cannot test all of the configurations** within a year and (ii) the device **stores** up to **8 configurations** between monthly clinic visits, the issues that the **calibration phase** needs to address are

(1) how to select 8 configurations to test each month,

(2) what data to collect about each tested configuration, and

(3) how to use that data at the end of the month to

i. evaluate whether to stop the trial early, and if not to

ii. pick the next month's 8 configurations,

all while balancing these two goals:

(a) finding **highly preferred** configurations, and

(b) exploring untested configurations

Calibration Phase Components

Part 1 - Monthly Outcomes

After the device is **reprogrammed with 8 configurations**,

(1) each configuration will be **randomly assigned** to a day with the *constraint* that a configuration cannot be observed on two consecutive days,

(2) and the **outcomes** are the patient-recorded answers to "did you **prefer** today's over yesterday's configuration?" If the patient did prefer the former over the latter configuration, the outcome is equal to 1, if not then 0.

What does the collected data look like?

$\{\boldsymbol{Y_{m,d}}, \boldsymbol{X_{m,d}}, \boldsymbol{X_{m,d-1}}\}$, where these are the patient-reported pairwise preferences between the current *(m,d)* and past *(m,d-1)* configurations.

Figure 1: An example of one month of patient-reported pairwise preferences between consecutively tested configurations. Configuration set is (A, B, C, D, E, F, G, H).

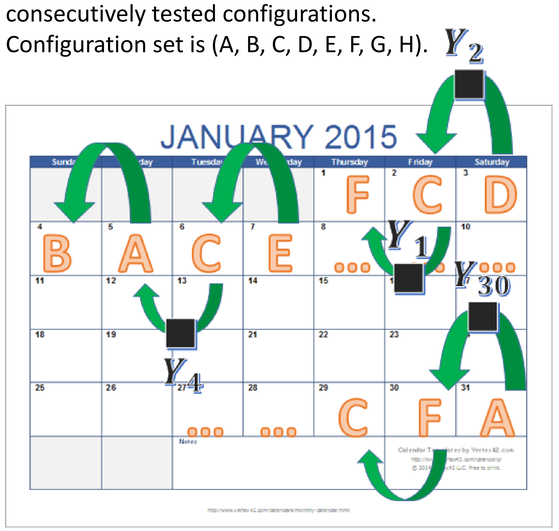

Part 2 - What Enables Calibration: Modeling Assumptions

(1) We assume that the preference of "today's" over "yesterday's" configuration follows a logistic regression model:

$$\boldsymbol{Prob}\left(\boldsymbol{Y_{m,d}} = 1 \,\middle|\, \boldsymbol{\alpha}, \boldsymbol{X_{m,d}}, \boldsymbol{X_{m,d-1}}\right) = \frac{\boldsymbol{exp}[\boldsymbol{\alpha}(\boldsymbol{X_{m,d}}) - \boldsymbol{\alpha}(\boldsymbol{X_{m,d-1}})]}{\mathbf{1} + \boldsymbol{exp}[\boldsymbol{\alpha}(\boldsymbol{X_{m,d}}) - \boldsymbol{\alpha}(\boldsymbol{X_{m,d-1}})]}$$

with latent *preferences* {α} for the configurations.

(2) We assume that the latent preferences ***{α}*** for the configurations follows a bi-directional spatial distribution (otherwise known as 2NRCAR). This assumption enforces that latent preferences for first-degree neighbors in **configuration indices** are related, thus informing preferences for **untested** configurations.

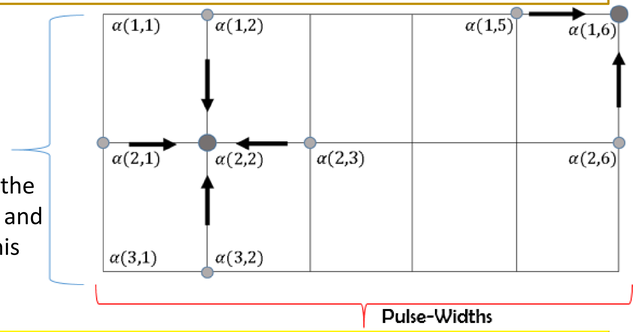

Figure 2: 3 x 6 rectangular grid illustrating the assumption on the latent preferences ***{α}***: preferences of neighbors in frequency and pulse-width inform the estimate of each latent preference. This enables estimation **untested** configurations' preferences.

Part 3 – Configuration "Batch" Selection

We balance configuration *exploration* and preference *maximization* through *Bayesian Optimization.* Using the current *latent preferences* ***{α}***, sequential optimization proposes the configuration with the highest *acquisition function,* $u^s(\boldsymbol{X})$, to test.

While these sequential methods propose **one** configuration, we need **a "batch" of 8 configurations** for the next month. *Bayesian Batch Optimization* penalizes $u^s(\boldsymbol{X})$ by each previously chosen configuration, and takes the max $u^b(\boldsymbol{X})$.

$$\underbrace{u^b(\boldsymbol{X}|\mathcal{H}^{j-1})}_{\text{Batch "Acquisition" for Configuration } \boldsymbol{X}} = \underbrace{h(\boldsymbol{X}|\mathcal{H}^{j-1})}_{\substack{\text{Sequential Penalty Onto} \\ \text{Configuration } \boldsymbol{X}}} \underbrace{u^s(\boldsymbol{X})}_{\substack{\text{"Old Style" Sequential} \\ \text{Acquisition Value for} \\ \text{Configuration } \boldsymbol{X}}}$$

Figure 3: Bayesian "Batch" Acquisition with penalizing a candidate configuration for each chosen batch element.

Part 4 – Early Stopping Rules

A year-long calibration phase is unreasonable if the patient shows (i) **no preference for a calibration** or (ii) **further calibration evidences no further improvement** on top of the current "best" configuration. Each "behavior" has a corresponding hypothesis test that captures it:

Preference Neutrality:

$$r[\boldsymbol{\alpha}] = \max \boldsymbol{\alpha} - \min \boldsymbol{\alpha}$$

$$H_0^N : r[\boldsymbol{\alpha}] < 1 \text{ versus } H_A^N : r[\boldsymbol{\alpha}] \geq 1$$

Calibration Convergence:

$$q[\boldsymbol{\alpha}] = \max_{\boldsymbol{X} \in \mathcal{X}} \alpha(\boldsymbol{X}) - \alpha(\hat{\boldsymbol{X}}^{opt})$$

$$H_0^C : q[\boldsymbol{\alpha}] \geq \mu \text{ versus } H_A^C : q[\boldsymbol{\alpha}] < \mu$$

with $\hat{\boldsymbol{X}}^{opt}$ being the month's estimated "best" configuration and **μ** is the "interval of convergence". At the end of a given month, **we stop the trial** if the preferences "behave" like one of the two cases **in red beyond a tolerable degree.**

Simulation Results

(1) If a patient ***doesn't prefer any*** of the configurations then their calibration *experience* is great because **any** configuration will do;

(2) whereas those that have ***strong preferences*** have to *explore poorer calibrations,* and therefore require more time to find a great calibration.

(3) Specifying smaller **μ,** or larger "stopping behavior" thresholds, requires a longer calibration time. Tweaking with these "tuning knobs" can yield desirable trial length at little expense in location of a patient's "near-best" configuration.

Works Cited

1. Reich, B. J., Hodges, J. S., and Carlin, B. P. (2007). Spatial Analyses of Periodontal Data Using Conditionally Autoregressive Priors having Two Classes of Neighbor Relations. Journal of the American Statistical Association 102, 44{55).
2. Ginsbourger, D., Le Riche, R., and Carraro, L. (2008). A Multi-Points Criterion for Deterministic Parallel Global Optimization based on Gaussian Processes. hal-00260579
3. Kaplan, A. and Murray, T. A. (2020). Batch Bayesian Optimization Design for Optimizing a Neurostimulator. Biometrics, *Submitted with revisions*.