Bayesian Optimization for Personalized Dose-Finding Trials with Combination Therapies

James Willard[1]
Shirin Golchi[1], Erica EM Moodie[1],
Bruno Boulanger[2], and Bradley P Carlin[2]

[1]McGill University
[2]PharmaLex Belgium

October 27, 2023

Disclaimer

PHARMALEX

The information provided during this presentation does not constitute legal advice. PharmaLex, and its parent Cencora, strongly encourage the audience to review available information related to the topics discussed during the presentation and to rely on their own experience and expertise in making decisions related thereto. Further, the contents of this presentation are owned by PharmaLex and reproduction of the slides used in today's presentation is not permitted without consent of PharmaLex.

Conflict of interest: Bruno Boulanger and Bradley P. Carlin are members of the local organizing committee for BAYES2023

Overview

1. Motivating Example
2. Bayesian Optimization
3. Simulation Study
4. Discussion and Future Work

Motivating Example

Motivating Example

- Motivated by a problem in industry
- Identify optimal combination of two agents $\mathbf{d}_{opt} = (d_1, d_2)_{opt}$ across continuous dose combination space $\mathbf{d}_{opt} \in \mathbb{D}$ where $\mathbb{D} \subset \mathbb{R}^2$
  - Optimality defined w.r.t. efficacy – continuous measure
- Both agents were approved individually – used for a long time
  - The agents are well tolerated
  - Assume minimal toxicity $\rightarrow$ do not consider dose escalation/de-escalation rules
- Response heterogeneity across binary covariate $\mathbf{Z}$ expected

Dose-Finding

- *Standard dose-finding* $\rightarrow$ ignore covariate information
  - Find optimal dose combination in "one size fits all" trials: $\mathbf{d}_{opt} \in \mathbb{D}$
- *Personalized dose-finding* $\rightarrow$ optimal dose combination may depend on covariates Z_p for $p = 1, 2, ..., P$
  - Cartesian product of the levels of these Z_p form $k = 1, ...K$ strata
  - Find optimal dose combination conditional on being in stratum k: $\mathbf{d}_{opt} \in \mathbb{D} \mid \mathbf{z}_k$
- Challenge: standard dose-finding methods not easily extended to personalized case
  - Parametric models require potentially large number of treatment-covariate interactions
  - Hard to estimate given sample size limitations in early phase trials

Proposed Method

- Small sample sizes lead to challenges in both standard and personalized dose-finding
- Proposed method has the following:
  1. Dose-response surface model is parsimonious, yet flexible – captures non-monotonic behavior
  2. Sequential design explores entire dose combination space in principled way – allows for early stopping
  3. Straightforward extension from standard to personalized dose-finding

- Method: response surface modeled by Gaussian process (GP) and Bayesian optimization methods utilized to guide sequential dose exploration

Bayesian Optimization

Bayesian Optimization

- Bayesian optimization globally optimizes *expensive-to-evaluate* objective functions $f(\mathbf{d})$ (Gramacy, 2020; Garnett, 2023)
  - $f(\mathbf{d})$: continuous efficacy or utility surfaces
  - Commonly assume minimization problem:

$$\underset{\mathbf{d} \in \mathbb{D}}{\operatorname{argmin}} f(\mathbf{d})$$

- Model $f(\mathbf{d})$ with GP and choose $\mathbf{d}^{(c+1)}$ as the one which maximizes acquisition function $\alpha(\tilde{\mathbf{d}} \mid \mathcal{D})$:

$$\mathbf{d}^{(c+1)} = \underset{\tilde{\mathbf{d}} \in \mathbb{D}}{\operatorname{argmax}}\, \alpha(\tilde{\mathbf{d}} \mid \mathcal{D})$$

- Repeat until sample size limits reached or early stopping criteria satisfied

Gaussian Process Regression

- GP is a stochastic process – considers any finite collection of observations as being distributed *multivariate normal*
- GP prior placed on the (latent) objective function

$$f(\mathbf{d}) \sim GP(m(\mathbf{d}), \mathcal{K}(\mathbf{d}, \mathbf{d}'))$$

  - $m(\mathbf{d})$: mean function
  - $\mathcal{K}(\mathbf{d}, \mathbf{d}')$: kernel function
- We use zero-mean functions $m(\mathbf{d}) = 0$ and flexible anisotropic squared exponential kernel function
- Importantly, posterior known in closed form $\rightarrow$ another multivariate normal (Williams and Rasmussen, 2006; Murphy, 2023)

Standard Dose Finding

- Collect data $\mathcal{D} \rightarrow$ fit GP
- Next dose combination: maximizer of acquisition function, $\alpha(\tilde{\mathbf{d}} \mid \mathcal{D})$:

$$\mathbf{d}^{(c+1)} = \underset{\widetilde{\mathbf{d}} \in \mathbb{D}}{\operatorname{argmax}}\, \alpha(\widetilde{\mathbf{d}} \mid \mathcal{D}).$$

- Expected Improvement (EI) balances trade-off between
  - *exploring* regions in $\mathbb{D}$ where $f(\mathbf{d})$ imprecisely estimated
  - *exploiting* regions in $\mathbb{D}$ with desirable values of $f(\mathbf{d})$
- EI of $\widetilde{\mathbf{d}}$ is expected improvement over current best observation f^*:

$$\alpha_{EI}(\widetilde{\mathbf{d}} \mid \mathcal{D}) = \mathbb{E}[\max(0, f^* - f(\widetilde{\mathbf{d}})) \mid \mathcal{D}, \widetilde{\mathbf{d}}].$$

- If GP used to model $f(\mathbf{d})$, EI is available in closed form[Jones et al., 1998]
  - Quick to calculate/optimize
- Repeat until algorithm termination

Personalized Dose Finding

- Incorporate additional covariate information $\mathbf{Z}$ into kernel function
- $f(\mathbf{d}, \mathbf{z}_k)$ may differ across strata, so there may exist stratum-specific current best observation values f_k^*
- Calculate stratum-specific EI, $\alpha_{EI}(\tilde{\mathbf{d}} \mid \mathcal{D}_\epsilon, \mathbf{Z}_k)$, known analytically:

$$\alpha_{EI}(\tilde{\mathbf{d}} \mid \mathcal{D}, \mathbf{z}_k) = \mathbb{E}[\max(0, f_k^* - f(\widetilde{\mathbf{d}}, \mathbf{z}_k)) \mid \mathcal{D}, \mathbf{z}_k]$$

- Next dose combination within stratum k:

$$\mathbf{d}_k^{(c_k+1)} \mid \mathbf{z}_k = \underset{\widetilde{\mathbf{d}} \in \mathbb{D}}{\operatorname{argmax}}\, \alpha_{EI}(\tilde{\mathbf{d}} \mid \mathcal{D}, \mathbf{z}_k)$$

- Repeat until algorithm termination

Recommended Dose Combination

- Algorithm termination $\rightarrow$ maximum sample size reached OR early stopping criteria satisfied
- Proposed stopping criteria
  - Stop when little improvement expected over current best $\rightarrow \max_{\widetilde{\mathbf{d}} \in \mathbb{D}} \alpha_{EI}(\widetilde{\mathbf{d}} \mid \mathcal{D}) < \delta$ AND
  - After some exploration of $\mathbb{D} \rightarrow$ when $n > n_{STOP}$
- At termination
  1. Sample from $p(\mathbf{d}_{opt} \mid \mathcal{D})$ OR
  2. Point estimate: $\widehat{\mathbf{d}}_{opt} \leftarrow \text{argmin}_{\widetilde{\mathbf{d}}} E[f \mid \mathcal{D}, \widetilde{\mathbf{D}}]$
     - We choose this $\rightarrow$ computational demands of simulations

Simulation Study

Performance Metrics

Goal: compare performance of standard and personalized dose-finding approaches under 1) no response heterogeneity and 2) response heterogeneity

1. How well algorithms capture location of true $\mathbf{d}_{opt}$
   - Expected Euclidean distance between recommended $\widehat{\mathbf{d}}_{opt}$ and true $\mathbf{d}_{opt}$
2. How well algorithms capture true optimal value f_{opt}
   - Expected root posterior squared error loss (RPSEL) of pointwise posterior $p(f \mid \mathcal{D}, \widehat{\mathbf{d}}_{opt})$
3. How well point estimates converge to true optimal value f_{opt}
   - Expected posterior mean estimates $E[f \mid \mathcal{D}, \widehat{\mathbf{d}}_{opt}]$

No response heterogeneity - no early stopping

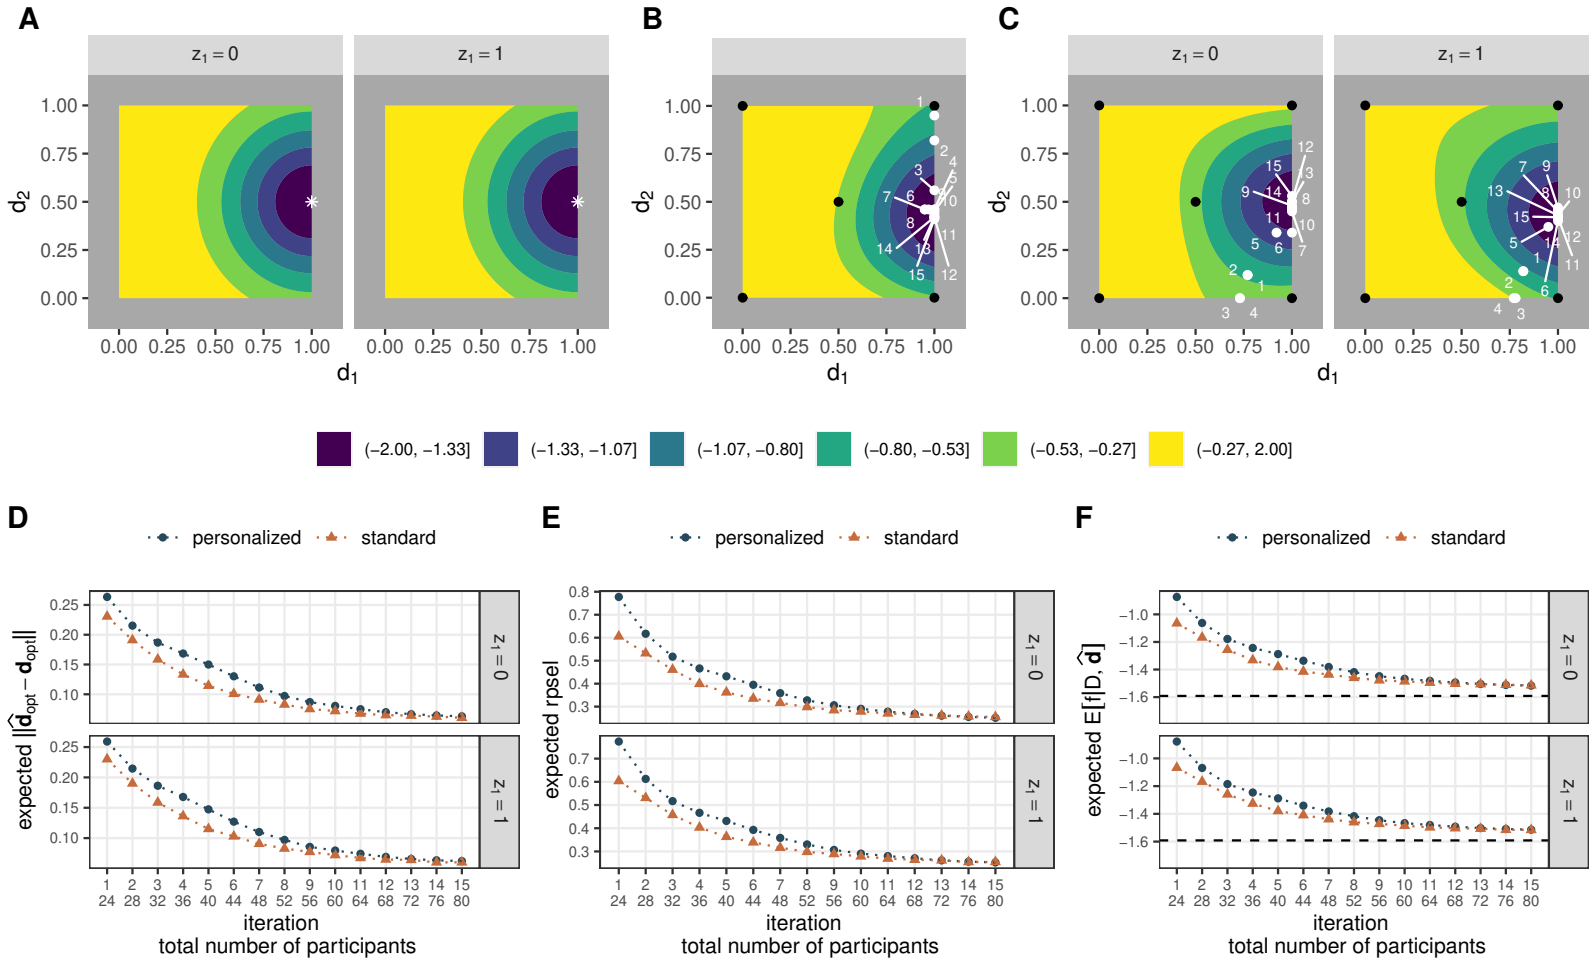

Response heterogeneity - no early stopping

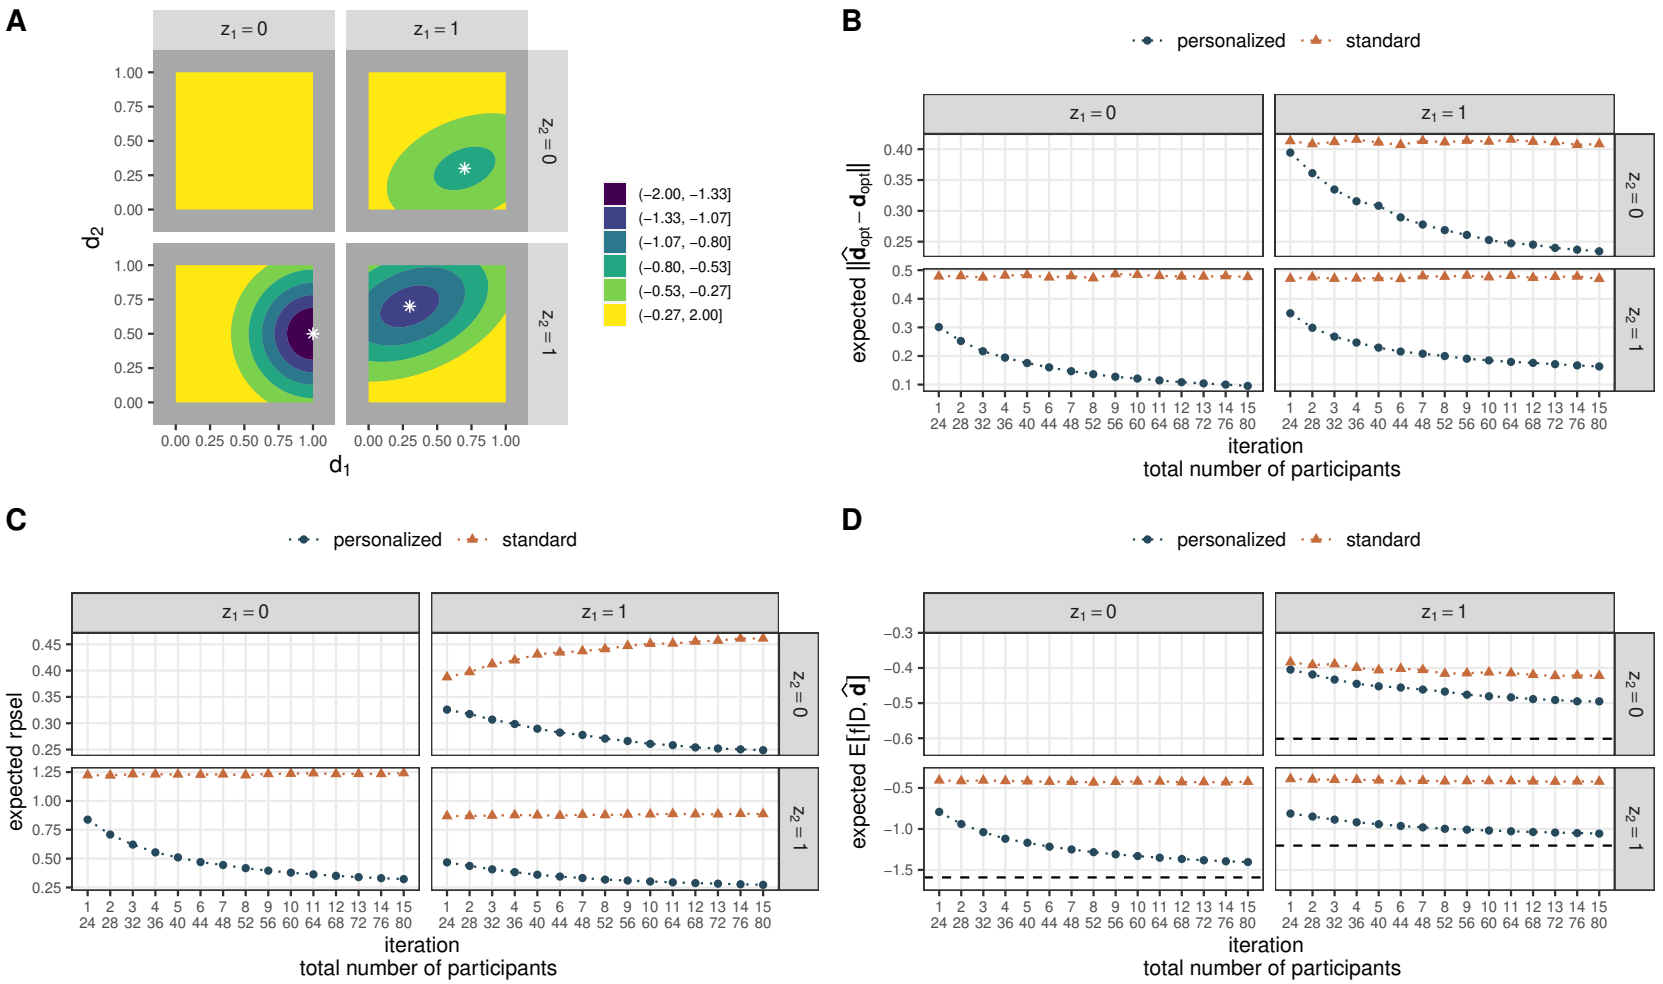

Motivating Example

- Identify $\mathbf{d}_{opt} \in \mathbb{D}$ where $\mathbb{D} \subset \mathbb{R}^2$
- Agents well tolerated - assume minimal toxicity setting
- Objective function $f(\mathbf{d})$ corresponds to continuous efficacy response
- Expect response heterogeneity across binary covariate $\mathbf{Z}$
- New doses expensive to engineer $\rightarrow$ stop early if possible

Goal: compare performance of standard and personalized dose-finding approaches under a variety of scenarios

Motivating Example - early stopping

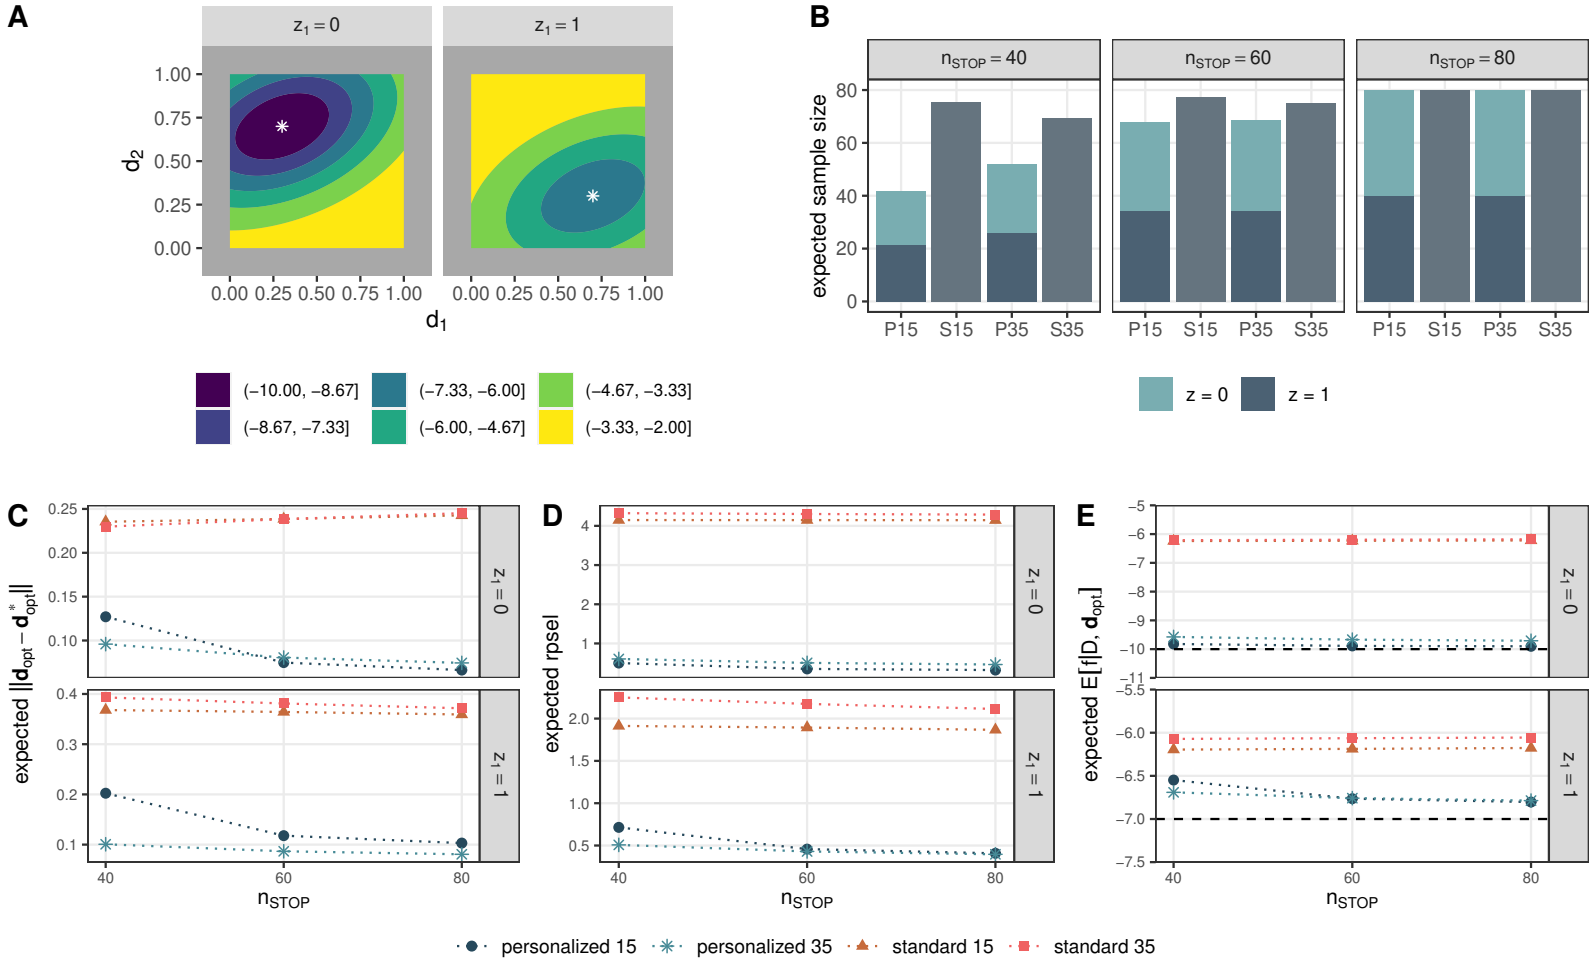

Discussion and future work

Takeaways

- No-response heterogeneity
  - personalized algorithm slightly less efficient
  - equivalent to standard algorithm by end of trial
- Response heterogeneity
  - poor performance of standard algorithm
  - incapable of separately modeling strata
- Personalized algorithm feasible even with small sample sizes
  - Final design of motivating example - around 50 participants for two strata
  - More strata may require more data

Future Work

- Extend to higher-grade toxicity setting $\rightarrow$ incorporate dose escalation/de-escalation rules
- Extend to non-categorized continuous covariates
- Investigate use of different kernel and acquisition functions

References

Garnett, R. (2023). *Bayesian Optimization*. Cambridge University Press, Cambridge, UK.

Gramacy, R. B. (2020). *Surrogates: Gaussian Process Modeling, Design and Optimization for the Applied Sciences*. Chapman Hall/CRC, Boca Raton, FL. `http://bobby.gramacy.com/surrogates/`.

Jones, D. R., Schonlau, M., and Welch, W. J. (1998). Efficient Global Optimization of Expensive Black-Box Functions. *Journal of Global Optimization*, 13(4):455–492.

Murphy, K. P. (2023). *Probabilistic Machine Learning: Advanced Topics*. MIT Press, Cambridge, MA.

Williams, C. K. and Rasmussen, C. E. (2006). *Gaussian Processes for Machine Learning*. MIT Press, Cambridge, MA.