Modern Bayesian Experimental Design

Tom Rainforth, Adam Foster, Desi R. Ivanova and Freddie Bickford Smith

Presenter: Alonso Granados

Bayesian Experimental Design

Key idea: better experiments lead to better data. Hence, we want to choose decisions that maximize the amount of information.

- ξ: controllable aspect (design).
- y: experiment's outcome.
- θ: the quantity we aim to gather information about.
- $p(\theta)$: prior
- $p(y|\theta, \xi)$: simulator.

$$\text{InfoGain}_\theta(\xi, y) = H(p(\theta)) - H(p(\theta|y,\xi))$$

where $p(\theta|y,\xi) \propto p(\theta)p(y|\theta,\xi)$.

$$\begin{aligned}\text{EIG}_\theta(\xi) &= \mathbb{E}_{p(y|\xi)}[\text{InfoGain}_\theta(\xi, y)] \\ &= \mathbb{E}_{p(\theta)p(y|\xi,\theta)}[\log p(\theta|y,\xi) - \log p(\theta)] \\ &= \mathbb{E}_{p(\theta)p(y|\xi,\theta)}[\log p(y|\theta,\xi) - \log p(y|\xi)]\end{aligned}$$

Bayesian Adaptive Design

BED can be extended to adaptive settings by iteratively incorporating information — referred as Bayesian adaptive design (BAD). The designs and outcomes can be broken down into a $\xi = \{\xi_1, ..., \xi_T\}$ and $y = \{y_1, ..., y_T\}$. Then the incremental EIG is defined as:

$$\mathrm{EIG}_\theta(\xi_t | h_{t-1}) = \mathbb{E}_{p(\theta|h_{t-1})p(y_t|\theta,\xi_t,h_{t-1})}\left[\log \frac{p(y_t|\theta, \xi_t, h_{t-1})}{p(y_t|\xi_t, h_{t-1})}\right]$$

where $h_{t-1} = \{(\xi_k, y_k)\}_{k=1}^{t-1}$.

Nested estimation

We will assume for now that our model is explicit — we can evaluate the densities $p(\theta)$ and $p(y|\theta, \xi)$.
First idea, estimate EIG using a MC estimator:

$$\text{EIG}_\theta(\xi) \approx \frac{1}{N} \sum_{n=1}^{N} \log p(y_n|\theta_n, \xi) - \log p(y_n|\xi),$$

where $\theta_n, y_n \sim p(\theta)p(y|\theta, \xi)$.

Another option is to use a Rao-Blackwellized estimator:

$$\hat{\mu}_N = \sum_{y \in \mathcal{Y}} \frac{1}{N} \sum_{n=1}^{N} p(y|\theta_n, \xi) \log p(y|\theta_n, \xi) - \hat{p}(y|\xi) \log \hat{p}(y|\xi),$$

where $\hat{p}(y|\xi) = \frac{1}{N} \sum_{n=1}^{N} p(y|\theta_n, \xi)$.

Nested MC estimator:

$$\hat{\mu}_{N,M} = \frac{1}{N} \sum_{n=1}^{N} \log \frac{p(y_n|\theta_n, \xi)}{\frac{1}{M} \sum_{m=1}^{M} p(y_n|\theta'_m, \xi)}.$$

Can be improved by importance sampling:

$$\hat{\mu}_{N,M,q} = \frac{1}{N} \sum_{n=1}^{N} \log \frac{p(y_n|\theta_n, \xi)}{\frac{1}{M} \sum_{m=1}^{M} \frac{p(\theta'_m)p(y_n|\theta'_m,\xi)}{q(\theta'_m|y_n,\xi)}}.$$

Debiasing Schemes

$$\mathrm{EIG}_\theta(\xi) = \mathbb{E}[\log \frac{p(y_1|\theta_n,\xi)}{\lim_{m\to\infty} \frac{1}{M}\sum_{m=1}^{M} \frac{p(\theta'_m)p(y_1|\theta'_m,\xi)}{q(\theta'_m|y_1,\xi)}}] = \mathbb{E}[\sum_{l=0}^{\infty} \Delta_l]$$

$$= \mathbb{E}[\frac{\Delta_l}{r(l)}]$$

Nested MC is very inefficient:

$$\hat{\mu}_{N,M,q} = \frac{1}{N}\sum_{n=1}^{N} \log \frac{p(y_n|\theta_n,\xi)}{\frac{1}{M}\sum_{m=1}^{M}\frac{p(\theta'_m)p(y_n|\theta'_m,\xi)}{q(\theta'_m|y_n,\xi)}}.$$

Use samples instead to learn $q(y|\xi) \approx p(y|\xi)$.

We have discussed these bounds in previous presentations:

$$\mathrm{EIG}_\theta(\xi) \leq \mathbb{E}_{p(\theta)p(y|\xi,\theta)}[\log p(y|\theta,\xi) - \log q(y|\xi)]$$

$$\mathrm{EIG}_\theta(\xi) \geq \mathbb{E}_{p(\theta)p(y|\xi,\theta)}[\log q(\theta|y,\xi) - \log p(\theta)]$$

$$\mathrm{EIG}_\theta(\xi) \leq \mathbb{E}_{p(\theta)p(y|\xi,\theta)}[\hat{\mu}_{1,M,q}].$$

Implicit models: estimate ratio $\frac{p(y|\theta,\xi)}{p(y|\xi)}$, use variational bounds, and ABC techniques.

Optimization methods

We have learned a few ways to evaluate $\text{EIG}_\theta(\xi)$, but our main goal is to select the design that maximizes information. We can attempt to do this directly by using $\nabla_\xi \text{EIG}_\theta(\xi)$.

- Gradients on Nested-MC.
- Gradients nested laplace approximations.
- Gradient optimization w.r.t. contrastive bound.

Policies methods

Despite these innovations, BAD can still be a prohibitively expensive challenge. Traditional BAD methods make greedy actions that failed to account future steps. One solution to this problem is to learn a policy rather than optimizing designs directly.

Linking with Related Areas/ Future directions

- Bayesian active learning: can deal with high-dimensional datasets, but is constrained to classifications problems.
- Reinforcement learning: exploration is underpinned by the idea of information gain.
- Model Misspecification: analysis is limited on both theoretical and empirical implications.
- Model improvements: improvements in implicit models have significant potential given that it is often easier to define an accurate simulator.