

Design Amortization for Bayesian Optimal Experimental Design

Noble Kenamer¹, Steven Walton², Alexander Ihler¹

¹Department of Computer Science, University of California Irvine

²Department of Computer Science, University of Oregon
nkenname@uci.edu

Abstract

Bayesian optimal experimental design is a sub-field of statistics focused on developing methods to make efficient use of experimental resources. Any potential design is evaluated in terms of a utility function, such as the (theoretically well-justified) expected information gain (EIG); unfortunately however, under most circumstances the EIG is intractable to evaluate. In this work we build off of successful variational approaches, which optimize a parameterized variational model with respect to bounds on the EIG. Past work focused on learning a new variational model from scratch for each new design considered. Here we present a novel neural architecture that allows experimenters to optimize a single variational model that can estimate the EIG for potentially infinitely many designs. To further improve computational efficiency, we also propose to train the variational model on a significantly cheaper-to-evaluate lower bound, and show empirically that the resulting model provides an excellent guide for more accurate, but expensive to evaluate bounds on the EIG. We demonstrate the effectiveness of our technique on generalized linear models, a class of statistical models that is widely used in the analysis of controlled experiments. Experiments show that our method is able to greatly improve accuracy over existing approximation strategies, and achieve these results with far better sample efficiency.

1 Introduction

Conducting experiments is often a resource-intensive endeavour, motivating experimenters to design their experiments to be maximally informative given the resources available. Optimal experimental design (OED) aims to address this challenge by developing approaches to define a utility, $U(d)$, of a possible designs, d , and algorithms for evaluating and optimizing this utility over all feasible designs \mathcal{D} . OED has been used widely across science and engineering, including systems biology (Liepe et al. 2013), geostatistics (Diggle and Lophaven 2006), manufacturing (Antony 2001) and more (Goos and Jones 2011).

In this work we focus on evaluating the expected information gain (EIG), a commonly used utility function in Bayesian optimal experiment design (BOED) (Chaloner and Verdinelli 1995; Ryan et al. 2016). We specify our model, composed of a likelihood and prior $p(y|\theta, d)p(\theta)$ for design

d , possible experimental outcomes y and latent variables θ . The EIG is then defined to be:

$$EIG(d) = \mathbb{E}_{p(y|d)} [H[p(\theta)] - H[p(\theta|y, d)]] \quad (1)$$

where $H[\cdot]$ is the entropy function. The experimenter then seeks the solution to $\operatorname{argmax}_{d \in \mathcal{D}} EIG(d)$, where \mathcal{D} is the set of all feasible designs. The EIG has sound theoretical justifications, proven to be optimal in certain settings (Sebastiani and Wynn 2000; Bernardo and Smith 2009).

While powerful, this framework is limited by the difficulty in evaluating the EIG due to the intractability of the posterior distribution $p(\theta|y, d)$. Foster et al. (2019) proposed four variational bounds for efficiently approximating the EIG. The method involves defining a variational distribution, $q_\phi(\cdot)$ that will approximate either the posterior distribution $p(\theta|y, d)$ or the marginal likelihood $p(y|d)$. The parameters ϕ of this variational distribution are optimized according to the proposed bounds. In principle their method allows for estimating the EIG for arbitrarily complex models using flexible variational models, however their implementations focused on simpler variational forms that required fitting a new variational model for every possible design. In this work we focus on design amortization by proposing a novel deep learning architecture based on conditional normalizing flows (NF) (Papamakarios et al. 2021; Kobyzev, Prince, and Brubaker 2020) and set invariant models (Zaheer et al. 2017) to define a flexible variational distribution $q_\phi(\cdot|d)$ that only needs to be trained once, but can then accurately estimate the EIG for potentially infinitely many designs. Our experiments will show how design amortization can dramatically improve computational efficiency and how our more flexible variational form can make much more accurate approximations to the EIG over competing methods. We provide our code here¹.

2 Background

Scientists consistently face the challenge of having to conduct experiments under limited resources, and must design their experiments to use these resources as efficiently as possible. BOED provides a conceptually clear framework for doing so. We assume we are given a model with design variables d , experimental outcomes y and latent parameters θ

¹https://github.com/NobleKenamer/amortized_boed

about which we wish to learn. We have prior information on the latent variables encoded in a prior distribution, $p(\theta)$, and a likelihood that predicts experimental outcomes from a design and latent variables, $p(y|\theta, d)$. Via Bayes rule, these two functions combine to give us the posterior distribution $p(\theta|y, d) \propto p(y|\theta, d)p(\theta)$ representing our state of knowledge about the latent variables after conducting an experiment with design d and observing outcomes y . For example the design variables, d , could represent the environmental conditions and chemical concentrations of a medium used to culture a strain of bacteria, which produces an important chemical compound. This design problem becomes more complex with increasing dimension of d , for example, if we have S petri dishes to work on (often called the *experimental units* or *subjects*). The experimental outcomes, y , would represent the amount of the chemical compound yielded from growing the culture in each of the conditions of d , and the latent variables θ represent parameters that define how the design variables d mediate the yield of the chemical compound y . After conducting the experiment and observing y , we can quantify our information gain (IG) as:

$$IG(y, d) = H[p(\theta)] - H[p(\theta|y, d)] \quad (2)$$

However, this gain cannot be evaluated before conducting the experiment, as it requires knowing the outcomes y . However, taking the expectation of the information gain with respect to the outcomes, $p(y|d)$, gives the EIG:

$$\begin{aligned} EIG(d) &= \mathbb{E}_{p(\theta, y|d)} \left[\log \left(\frac{p(\theta|y, d)}{p(\theta)} \right) \right] \\ &= \mathbb{E}_{p(\theta, y|d)} \left[\log \left(\frac{p(y|\theta, d)}{p(y|d)} \right) \right] \end{aligned} \quad (3)$$

Nested Monte Carlo: Typically $p(\theta|y, d)$ and $p(y|d)$ are intractable, making the EIG challenging to compute. One common approach to approximating EIG is to use a nested Monte Carlo (NMC) estimator (Myung, Cavagnaro, and Pitt 2013; Vincent and Rainforth 2017; Rainforth et al. 2018):

$$\hat{\mu}_{NMC} = \frac{1}{N} \sum_{n=1}^N \log \frac{p(y_n|\theta_{n,0}, d)}{\frac{1}{M} \sum_{m=1}^M p(y_n|\theta_{n,m}, d)} \quad (4)$$

where $\theta_{n,m} \sim p(\theta)$ and $y_n \sim p(y|\theta_{n,0}, d)$

Rainforth et al. (2018) showed that NMC is a consistent estimator converging as $N, M \rightarrow \infty$. They also showed that it is asymptotically optimal to set $M \propto \sqrt{N}$, resulting in the overall convergence rate of $\mathcal{O}(T^{-\frac{1}{3}})$, where T is the total number of samples drawn (i.e. $T = NM$ for NMC). However, this is much slower than the $\mathcal{O}(T^{-\frac{1}{2}})$ rate of standard Monte Carlo estimators (Robert and Casella 1999), in which the total number of samples is simply $T = N$.

The slow convergence of the NMC estimator can be limiting in practical applications of BOED. The inefficiencies can be traced to requiring an independent estimate of the marginal likelihood, $p(y_n|d)$, for each y_n (the denominator of Eq. (4)). Inspired by this, Foster et al. (2019) proposed employing techniques from variational inference by defining a functional approximation to either $p(\theta|y, d)$ or $p(y|d)$,

and allowing these estimators to amortize across the samples of y_n for more efficient estimation of the EIG. In this work we focus on two of the four estimators they proposed: the posterior estimator and variational nested Monte Carlo.

Posterior Estimator: The posterior estimator is an application of the Barber-Agakov bound to BOED, which was originally proposed for estimating the mutual information in noisy communication channels (Barber and Agakov 2003). It requires defining a variational approximation $q_\phi(\theta|y, d)$ to the posterior distribution, giving a lower bound to the EIG:

$$\begin{aligned} EIG(d) &\geq \mathcal{L}_{post}(d) \triangleq \mathbb{E}_{p(\theta, y|d)} \left[\log \left(\frac{q_\phi(\theta|y, d)}{p(\theta)} \right) \right] \\ &\approx \frac{1}{N} \sum_{n=1}^N \log \frac{q_\phi(\theta_n|y_n, d)}{p(\theta_n)} \end{aligned} \quad (5)$$

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

By maximizing this bound with respect to the variational parameters ϕ , we can learn a variational form that can efficiently estimate the EIG. A Monte Carlo estimate of this bound converges with rate $\mathcal{O}(T^{-\frac{1}{2}})$, and if the true posterior distribution is within the class of functions defined by the variational form q_ϕ , the bound can be made tight (dependent on the optimization) (Foster et al. 2019).

Variational Nested Monte Carlo: The second bound we discuss is variational nested Monte Carlo (VNMC). It is closely related to NMC, but differs by applying a variational approximation $q_\phi(\theta|y, d)$ as an importance sampler to estimate the marginal likelihood term in NMC:

$$\begin{aligned} EIG(d) &\leq \\ \mathcal{U}_{VNMC}(d, M) &\triangleq \mathbb{E} \left[\log \frac{p(y|\theta_0, d)}{\frac{1}{M} \sum_{m=1}^M \frac{p(y, \theta_m|d)}{q_\phi(\theta_m|y, d)}} \right] \end{aligned} \quad (6)$$

where the expectation is taken with respect to $y, \theta_{0:M} \sim p(y, \theta_0|d) \prod_{m=1}^M q_\phi(\theta_m|y, d)$.

By minimizing this upper bound with respect to the variational parameters ϕ , we can learn an importance distribution that allows for much more efficient computation of the EIG. Note that if $q_\phi(\theta|y, d)$ exactly equals the posterior distribution, the bound is tight and requires only a single nested sample ($M = 1$). Even if the variational form does not equal the posterior, the bound remains consistent as $M \rightarrow \infty$. Finally, it is worth noting that by taking $q_\phi(\theta|y, d) = p(\theta)$, the estimator simply reduces to NMC.

It was further shown by Foster et al. (2020) that VNMC can be easily made into a lower bound by including θ_0 (the sample from the prior) when estimating the marginal likelihood, a method we denote as contrastive VNMC (CVNMC):

$$\begin{aligned} EIG(d) &\geq \\ \mathcal{L}_{CoVNMC}(d, M) &\triangleq \mathbb{E} \left[\log \frac{p(y|\theta_0, d)}{\frac{1}{M+1} \sum_{m=0}^M \frac{p(y, \theta_m|d)}{q_\phi(\theta_m|y, d)}} \right] \end{aligned} \quad (7)$$

where the expectation is taken with respect to $y, \theta_{0:M} \sim p(y, \theta_0|d) \prod_{m=1}^M q_\phi(\theta_m|y, d)$. We can also employ this same technique to regular NMC to estimate both lower and upper bounds.

Note that the upper bound (6) and lower bound (7) are particularly useful when evaluating the performance of our method in settings where ground truth is not available. In these cases we can examine the bound pairs produced by NMC and by VNMC to assess which set more tightly constrains the true value.

Practical considerations. In this work, we apply the same flexible mathematical framework proposed in Foster et al. (2019). However, Foster et al. (2019) adopted a “classical” variational distribution setting, in which the variational form q_ϕ is selected to take a standard, parametric form. They found this approach effective, but tested only on very simple design problems, with only one experimental unit at a time. Their variational models only incorporate the design implicitly, requiring a separate optimization for every design to be considered². Unfortunately, as we show in the experiments this approach is not effective on more complex design problems. Instead, we propose a far more flexible, deep learning based distributional form that incorporates the design explicitly, allowing us to amortize training across and apply our trained model to evaluation of all (potentially continuously many) designs in our feasible set.

3 Method

We are interested in learning a parameterized function, $q_\phi(\theta|y, d)$, for approximating the posterior distribution. We now describe our proposed deep learning architecture for amortizing over designs, allowing practitioners to train a single model that is capable of evaluating the EIG for potentially infinitely many designs. We also discuss how we can efficiently train this model using the (simpler and cheaper) equation (5), then use the resulting approximation in the more accurate bounds provided by VNMC, (6)–(7). This advances the work in Foster et al. (2019) by providing a highly flexible variational form that can be used in a wide variety of contexts and an inexpensive procedure to train it.

Neural Architecture

Figure 1 shows a high level representation of our architecture. Broadly, it consists of two major components. The first is a learnable function for taking in the design variables d and simulated experimental outcomes from the model, y , and producing a design context vector, $c_{y,d}$, that will be used to define a conditional distribution. We focus on the common case where the experimental units lack any meaningful order and our learnable function must therefore be permutation invariant. We can incorporate this inductive bias into our model by making our function follow the general form of set functions proposed in Zaheer et al. (2017). In the sequel, we denote this component as our *set invariant*

²Although subsequent work (Foster et al. 2020) considered evolving both the design and distribution q simultaneously, even that work remains focused on a single (if evolving) design.

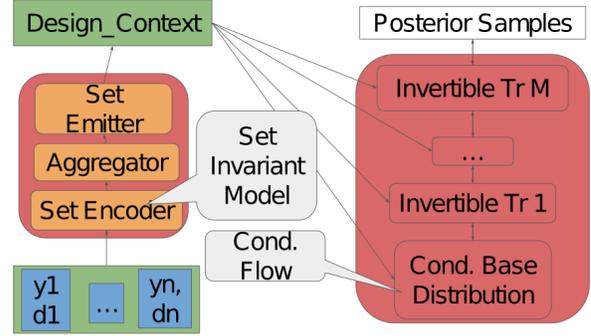


Figure 1: A high-level schematic of our architecture for amortizing over designs. The first component (left) takes in the design variables and simulated observations and produces a design context, $c_{y,d}$. In many experiments the individual units being experimented on are exchangeable, thus we use a set invariant architecture. The second (right) is a conditional normalizing flow, conditioned on the design context produced by the first component. Together, they define our variational posteriors $q_\phi(\theta|y, d)$, amortized over designs.

model. The second major component is a learnable distribution conditioned on the design context produced by the set invariant model. In this work we use conditional normalizing flows, which consist of a base distribution and sequence of invertible transformations with tractable Jacobian determinants to maintain proper accounting of the probability density through the transformations. Both the base distribution and transformations are learnable and conditioned on the design context.

Set Invariant Model. It is often the case that the individual units being experimented on do not possess an inherent ordering – for example, subjects in a randomized controlled clinical trial, or the petri dishes in our previous example. Suppose we would like to find the optimally informative design for an experiment with S experimental units, where d_i and y_i denote the design variables and simulated outcomes of unit i , respectively. In this setting we want our design context to be invariant to permutations in its inputs, e.g., re-ordering the individuals in the trial should not change our results. Learning permutation invariant functions is an active area of research (e.g., Bloem-Reddy and Teh 2020). In this work we follow the general form proposed by Zaheer et al. (2017), where our set invariant model is defined as,

$$c_{y,d} = \text{EMIT}_{\phi_{\text{EMIT}}} \left[\sum_{i=0}^S \text{ENC}_{\phi_{\text{ENC}}}(y_i, d_i) \right]. \quad (8)$$

In particular, we define two learnable functions. The set encoder $\text{ENC}_{\phi_{\text{ENC}}}(y_i, d_i)$ takes as input the design variables and simulated outcomes for each individual experimental unit. Its output is an intermediary representation for each experimental unit, which are aggregated together by summation; the permutation invariance of the sum ensures invariance of the overall function. The aggregated representation is then passed through the set emitter function

$\text{EMIT}_{\phi_{\text{EMIT}}}(\cdot)$, which creates the final design context used in the conditional normalizing flow. In our experiments we find substantially improved performance using attention layers (Vaswani et al. 2017) in the set encoder, which allows for interactions between the experimental units before aggregation. In this case we should denote our set encoder as $\text{ENC}_{\phi_{\text{encode}}}(y_i, d_i | y_{-i}, d_{-i})$ where y_{-i}, d_{-i} denote all simulated outcomes and design variables except for the i th unit.

Structuring our design encoder function in this way gives two major advantages. First, permutation invariance does not need to be learned by the function since it is already present by construction; this can make learning more efficient and reduces the total number of weights via weight sharing. Second, the function is able to encode designs with a variable number of experimental units, S , as long as the d_i and y_i have the same size for all units.

Note that not all experimental design problems are permutation invariant in the experimental units. For example, in some settings there could be a temporal component in the design variables, in which case we could replace our set invariant function with an order-based model such as a recurrent neural network.

Conditional Normalizing Flow. Normalizing flows define an expressive class of learnable probability distributions, which has been used in generative modeling and probabilistic inference (Papamakarios et al. 2021; Kobyzev, Prince, and Brubaker 2020). The main idea of normalizing flow based models is to represent a random variable θ as a transformation $\theta = T(x)$ of a random variable x sampled from a base distribution $p(x)$. The key property is that the transformation T must be invertible and differentiable. This allows us to obtain $p(\theta)$ via a change of variables,

$$p(\theta) = p(x) |\det J_T(x)|^{-1} \quad (9)$$

where $x = T^{-1}(\theta)$ and $\det J_T(x)$ is the determinant of the Jacobian at x . Both the transformations and base distribution may have learnable parameters. This provides a highly flexible class of distributions that can be both sampled and efficiently evaluated.

In our setting we would like to learn not just a single distribution, but rather a *conditional distribution* given design variables and experimental outcomes. This conditioning on d is key to allowing us to amortize over all possible designs. To this end, we learn a sequence of K conditional transformations $T_{\phi_i}(\cdot | c_{y,d})$ and a conditional base distribution $p_{\phi_0}(x | c_{y,d})$ which together define our variational approximation to the posterior distribution amortized over designs.

$$q_{\phi}(\theta | c_{y,d}) = p_{\phi_0}(x_0 | c_{y,d}) \prod_{i=1}^K |\det J_{T_{\phi_i}}(x_i | c_{y,d})|^{-1} \quad (10)$$

where $\theta = T_{\phi_K} \circ T_{\phi_{K-1}} \circ \dots \circ T_{\phi_1}(x_0)$, with all transformations and base distribution conditioned on the design context (8). The full architecture, including the set invariant model and conditional normalizing flow is trained end-to-end.

Variational Posterior Training

The posterior estimator and the (contrastive) VNMC bounds all require learning a variational approximation to the posterior distribution. In both Foster et al. (2019) and Foster

et al. (2020) each bound was trained separately, learning its own variational approximation. However in all cases the variational approximation produced by training one of the bounds *can* be used for evaluating any other bound, since they all only require an approximate posterior distribution. Ideally, we would like to train using only the posterior estimator since it is much cheaper – a total cost of only $\mathcal{O}(N)$ – whereas both VNMC bounds have a total cost of $\mathcal{O}(NM)$. However, it is not obvious that training on the (also less accurate) bound should still provide good EIG estimates when used in the VNMC bounds. Our experiments show that it is surprisingly effective across a broad range of models. Intuitively speaking, this is possible because all bounds share the same optimum – the true posterior distribution. Moreover, because the posterior estimator takes its expectation with respect to the model $p(y, \theta | d)$, the variational approximation $q_{\phi}(\theta | y, d)$ will in general be *wider* than the true posterior distribution, akin to variational inference using the “forward” Kullback-Liebler (KL) divergence, and in contrast to the more commonly used “reverse KL” variational optimization methods that result in underdispersed and mode-seeking optima. This property also makes the posterior estimator’s q_{ϕ} an excellent choice for importance sampling (as in VNMC), in which a too-narrow proposal distribution can lead to high variance in the importance weights, causing a small number of samples to dominate the estimator (Owen 2013, Chapter. 9).

4 Related Work

Our approach builds on the framework for BOED developed in Foster et al. (2019), which proposed four variational bounds for estimating EIG. The framework itself is quite flexible, capable of accommodating a wide variety of models (e.g., implicit vs explicit likelihoods), sequential experimentation (of arbitrary batch size) and arbitrary variational forms q . However, their experiments were limited to only single experimental units, and used simple variational forms that cannot amortize over designs (requiring a separate training procedure for each proposed design). Further recent work in Foster et al. (2021) has also studied an RL-based optimization approach to BOED, but it doesn’t produce EIG estimates, which is the main focus of this paper. In this work we propose a deep learning architecture which can easily be scaled to approximate arbitrarily complex distributions. In addition, our architecture can amortize over designs, allowing us to train a single variational model capable of estimating the EIG for potentially infinitely or continuously many designs. We also show that we can train our model using the cheaper posterior bound, then use its optimized approximate posterior within the VNMC bounds for a more accurate final approximation. We show that, using our proposed variational form, we can achieve highly accurate EIG estimates across a spectrum of complex design problems. While a few other EIG approximations have been proposed (see, e.g., Foster et al. (2019); Ryan et al. (2016)), in light of the experimental results of Foster et al. (2019) we mainly compare our experimental performance relative to NMC.

5 Experiments

We perform three types of experiments: *amortization*, *model* and *architecture* experiments. Our *amortization* experiment shows the dramatic increase in efficiency from amortization, and better EIG estimation provided by our more complex variational forms compared to those used in Foster et al. (2019). *Model* experiments examine how the benchmark method, NMC, breaks down as model complexity grows while our methods remain reliable for accurately estimating the EIG. *Architecture* experiments measure the impact of key components in our variational approximation and serve as a guide to using our method effectively.

In all experiments we focus on estimating designs for different types of generalized linear models (GLMs) (McCullagh and Nelder 1989). GLMs are a very common model class used to analyze controlled experiments and are regularly used in applications of optimal experimental design (Goos and Jones 2011). Our GLMs have the general pattern,

$$\begin{aligned} \theta &\sim \mathcal{N}(\mu_p, \Sigma_p) \\ r &= g^{-1}(D\theta) \\ y &\sim \text{Exponential Family Distribution}(r). \end{aligned} \quad (11)$$

Here, θ is a $N_p + 1$ dimensional parameter vector, where N_p is the number of predictors (+1 for the intercept term). D is a $N_E \times (N_p + 1)$ design matrix, where N_E is the number of experimental units. The inverse link function, g^{-1} , defines the type of GLM. Finally, μ_p and Σ_p are the prior mean and covariance of the parameters. Our experiments cover six GLMs: normal (known observation noise), normal unknown (unknown observation noise), logistic, binomial, categorical and multinomial. For the normal model with known observation noise we take $\sigma = 1$; for the normal model with unknown observation noise³ we use the prior $\sigma \sim \text{InverseGamma}(a_p, b_p)$ with $a_p = b_p = 3.5$. For the binomial model, we assume 10 random trials; we use 3 classes for the categorical model, and 10 trials and 3 classes for the multinomial model. All experiments in the sequel are run for designs with 5 experimental units. Our implementations made significant use of Pyro (Bingham et al. 2019) to implement the inference procedures and NFlows (Durkan et al. 2020) to construct our conditional normalizing flows. All training was done on a single Nvidia 2080TI and evaluation was done on an Intel I7-9800X with 64 GB of RAM.

Amortization Experiments

While providing an excellent framework, the variational forms used in Foster et al. (2019) are too simple to be effective on the GLM models we consider. Additionally, their work required training a new variational model for every design being approximated, while we propose a method that can amortize over designs. Our closest model to those tested in Foster et al. (2019) is our normal (linear) model, similar to their ‘‘AB model’’; we apply their variational form for the AB model (given in their appendix) in order to perform a

³In this case, the observation noise is included as the standard deviation in the normal distribution that samples y .

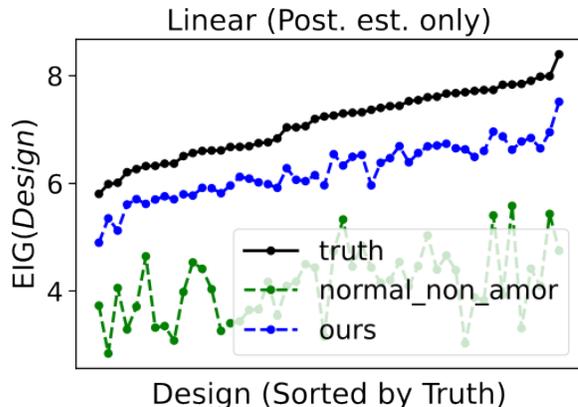


Figure 2: Comparing our method and the (non-amortized) variational form used by Foster et al. (2019) on the normal (linear) model with 5 predictors and 5 on the diagonal of the prior covariance (similar to the AB model in Foster et al. (2019)). For clarity we only show the posterior estimator values. Our method is $> 3\times$ faster (293s vs. 920s) and significantly more accurate. See text for further analysis.

comparison. We set $N_p = 5$ and $\Sigma_p = 5I$, i.e., diagonal with variance 5.

We generate 50 random designs with $N_E = 5$ experimental units and compare the quality of EIG approximations given by the posterior estimator as well as total wall clock time. The precise architecture and training procedure we use are described in the appendix. Figure 2 shows the results of this experiment: our method produces a much tighter lower bound that is highly correlated with the true values; selecting the highest estimate would pick the design with highest true EIG in the set. In contrast, the variational form used in Foster et al. (2019) yields a much looser and less correlated bound, which would select the 6th best design if used. Moreover, our method is more than $3\times$ faster (293 seconds compared to 920 seconds), showing the benefit of amortizing over designs. In fact, this speed-up understates how much more computationally efficient our method is, given that it leaves us with a model that can estimate the EIG for arbitrarily many designs without additional training. Training took 291 of our method’s 293 seconds; evaluating an additional 50 designs, then, would take virtually the same amount of time, compared to double the time required for the non-amortized approach. The non-amortized training is prohibitively slow, while moreover for our other GLM models it is often not clear what variational form from Foster et al. (2019) could be applied; for these reasons, in the rest of the experiments we compare only to standard NMC.

Model Experiments

In our model experiments we vary the GLMs in two ways: the number of predictors (not including the intercept) and the diagonal components of the prior covariance (all off-diagonal terms are zero). The number of predictors N_p is varied from 1 to 5 and the diagonal of the prior covariance in $\{1, 5, 25\}$. For the neural network architecture we use atten-

tion layers for the set encoder, a residual network for the set emitter (He et al. 2016), a full rank Gaussian distribution for the conditional base of the normalizing flow and four affine coupling layers each parameterized with a residual network. Additional details of the architecture and training parameters can be found in the appendix. In all experiments we train the variational approximation using the posterior estimator. During training, new designs are generated randomly from a multivariate normal distribution with identity covariance in dimension $N_p + 1$. For final evaluation we generate 50 new random designs and estimate the posterior bound with $N = 5000$ samples, while the VNMC bounds are estimated with $N = 1000$ and $M = 31$ samples and nested samples, and the NMC bounds are estimated with $N = 30000$ and $M = 173$ samples and nested samples. The number of samples for VNMC and NMC were selected based on the maximum number of samples that fit into memory (64 GB RAM) for the largest model (multinomial with 5 predictors).

Figure 3 shows EIG evaluations for 50 randomly generated designs for the 5 estimators: posterior, VNMC upper, VNMC lower (contrastive VNMC), NMC upper and NMC lower (contrastive NMC). Since this figure pertains to the linear model we can calculate the ground truth EIG exactly, shown in solid black. For visual clarity we sort the designs in order of ground truth EIG value. We see that all estimators perform reasonably well on the easiest form of the model (1 predictor with unit prior covariance). However, even in this case the VNMC bounds (upper and lower) more tightly constrain the ground truth – in fact both are nearly exact. In addition the posterior bound (a lower bound) is consistently above the NMC lower bound and closer to the truth. These trends become magnified as the prior covariance and number of predictors increase. In all cases the VNMC bounds are nearly exact, while the performance of NMC degrades rapidly with problem difficulty. Again, the posterior estimator remains above and closer to truth than the NMC lower bound.

Figure 4 shows exactly the same set of experiments, but for the linear unknown model. In this case we can calculate a high-quality Monte Carlo estimate of the ground truth thanks to conjugacy, and sort the designs to be evaluated in order of this ground truth. The results are largely consistent with those from the linear model with known observation noise: the VNMC bounds constrain the ground truth much more tightly than the NMC bounds. However in this case the posterior estimator is only above the NMC lower bound in the two hardest cases (5 predictors and 5 or 25 diagonal covariance).

Figures 6, 7, 8, 9, in the appendix, show the same set of experiments but for the logistic, binomial, categorical and multinomial models. In none of these cases can we calculate ground truth, so all plots order their designs by the benchmark NMC (upper). Even without ground truth we still clearly see the lower and upper bounds of VNMC are much closer together and below/above their NMC counterparts in practically all cases. One exception is the logistic model, where as long as the model only contains one predictor variable, the NMC bounds are as tight as the VNMC bounds; however by 5 predictors the VNMC bounds are

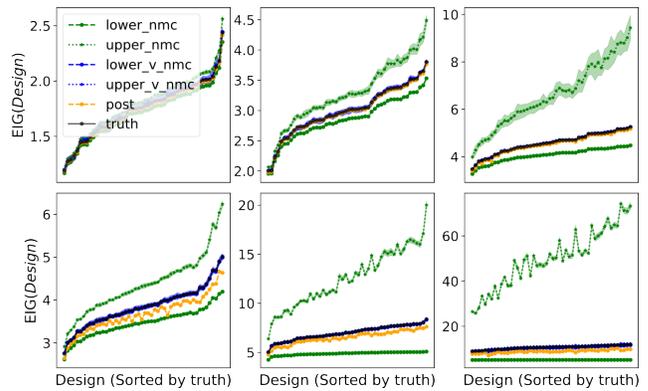


Figure 3: Results for estimating EIG in the linear model with known observation noise. The x-axis ranges over the index of 50 randomly selected designs, each with 5 experimental units. Due to the dimensionality, the designs lack a meaningful order; for visual clarity we plot them in the sorted order of the true EIG. The rows vary the number of predictors ($N_p \in \{1, 5\}$) while columns show changes in the diagonal of the prior covariance matrix, $\{1, 5, 25\}$, from informative to uninformative. NMC and VNMC methods can estimate both upper and lower bounds, while the posterior estimator only provides a lower bound. Our proposed method gives much tighter bounds on the truth than the competing NMC (with $167\times$ fewer samples). The shading shows one standard deviation of our estimates over 20 runs.

tighter. In fact the VNMC lower and upper bounds agree with each other in all cases, suggesting they are closely estimating the true EIG. We also see that the VNMC lower and upper bounds are touching across all settings of the binomial model and all but the hardest in the categorical and multinomial models (where they are still much tighter than NMC), again suggesting that our VNMC estimators are nearly exact. The results for the posterior estimator are more mixed – sometimes it is above the NMC lower bound and sometimes below. Nevertheless, the variational posterior learned using the posterior estimator can be used to compute VNMC bounds on the EIG that are much tighter than NMC.

Our experiments show that training a *single* variational posterior, amortizing over designs, we can calculate the EIG much more accurately than the competing NMC benchmark, nearly calculating ground truth exactly. Moreover, the experiments show that training on the posterior estimator can provide a variational distribution that remains effective for estimation using the more costly VNMC bounds (see Section 3). Not only does VNMC provide far more accurate estimates, it does so with many fewer samples – more than two orders of magnitude ($167\times$) fewer samples than NMC.

Architecture Experiments

We next investigate the importance of architectural decisions for the neural networks defining $q_\phi(\theta|y, d)$. We compare using attention layers vs. residual layers in the set encoder, and transform type and number in the normalizing flow. We compare using 4 vs. 8 transforms, and test affine cou-

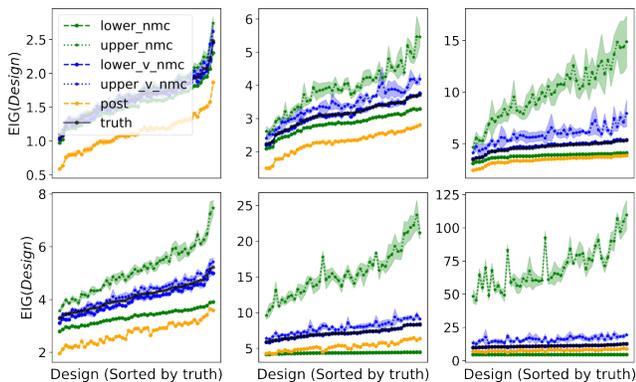


Figure 4: Same as Figure 3, but for the linear model with unknown observation noise.

pling transforms (Dinh, Sohl-Dickstein, and Bengio 2017), rational quadratic (RQ) splines (Durkan et al. 2019), and no transform (just the conditional base distribution). We run all combinations for the linear unknown and binomial model with 5 predictors and 25 on the diagonal of the prior covariance. Note that the true posterior for the linear unknown model is t -distributed, while the binomial is not analytically expressible, so we expect the use of normalizing flows to be advantageous over just the normal base distribution. The rest of the architecture components are the same as the Model Experiments and full details can be found in the appendix.

Figure 5 shows the loss curves of the experiments for the linear unknown model. The inset plot on the top right shows the loss curves across all epochs, while the main plot shows a detail of the last 50 optimization steps. Each optimization step is run on a batch of 50 designs, so this plot indicates final performance on 2500 randomly generated designs. Specifically the loss is $-\sum_{i=1}^N \log q_\phi(\theta_i|y_i, d)$ where $y_i, \theta_i \sim p(y, \theta|d)$ and $N = 50$ – the cross entropy of the variational posterior. Empirically, we see that using attention layers in the set encoder is the most important architectural decision (lower 5 curves vs upper); all networks using attention layers achieved superior performance to all networks using ResNets regardless of the other architectural settings. Beyond this, we see that using an affine coupling layer is also important, but see little difference between 4 and 8 transform layers. Surprisingly, the RQ transforms perform no better than having no transform. This is because the RQ transforms are restricted to the range of $(0, 1)$, with linear tails outside. Even after training, almost all parameter samples are outside this range by the time they reach the spline; the linear tails effectively skip the flow layers completely, explaining why its performance is comparable to models with no transform layers. In the appendix, Figure 11 shows the performance of a subset of these models at evaluating the EIG for 50 random designs, highlighting that the loss curves’ values are directly related to the accuracy and sample efficiency of the EIG estimate. Figures 10 and 12 show the same plots for the binomial model and further support these conclusions.

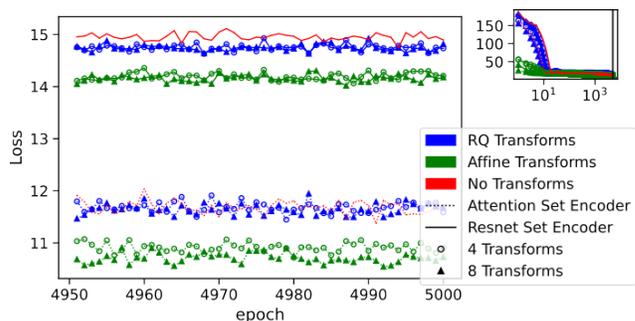


Figure 5: Results for our architecture experiments on the linear unknown model with 5 predictors and diagonal prior covariance 25. We vary the architecture of the set encoder (attention vs. resnet), the normalizing flow transform type (affine coupling, affine spline, or no transform), and the number of transforms (4 or 8). The main plot shows the loss for the posterior estimator over the last 50 steps of training; each step is performed over a batch of 50 random designs (2500 designs total). The inset plot shows the loss curves over all 5000 training steps, indicating all architectures have converged. Further discussion is given in the text.

6 Conclusion

In this paper we expand on the work of Foster et al. (2019), which proposed variational bounds for estimating the EIG for Bayesian optimal experimental design. In particular we propose a deep learning architecture incorporating set invariance and conditional normalizing flows that allows us to train a single model capable of estimating the EIG across the design space. Our experiments show that this architecture is highly effective at estimating EIG, and that design amortization provides significant computational speed ups. For cases where ground truth can be calculated, our model’s VNMC bounds are nearly exact, while in cases without ground truth our VNMC upper and lower bounds are often sufficiently tight to suggest they are exact. These estimates are significantly more accurate than those of standard NMC while requiring far ($167\times$) fewer samples, as well as far more accurate and efficient than the simpler, non-amortized variational forms used in (Foster et al. 2019). We also demonstrate that we can train our model using the much cheaper posterior estimator bound, with cost $\mathcal{O}(N)$, then evaluate using this fitted model within the more accurate but costly VNMC bounds, $\mathcal{O}(NM)$. Together, we provide a method for faster and more accurate approximation of the EIG across many possible designs. In future work, we plan to extend our approach to design optimization tasks; on this point, we observe that our variational form, $q_\phi(\theta|y, d)$, is differentiable with respect to the designs, d , suggesting it can generalize and potentially improve on the gradient-based design optimization objectives proposed by Foster et al. (2020).

Acknowledgments

This work was supported by the Hasso Plattner Institute and the National Science Foundation grant numbers IIS-1816365 and IIS-2008516.

References

- Antony, J. 2001. Improving the manufacturing process quality using design of experiments: a case study. *International Journal of Operations & Production Management*.
- Barber, D.; and Agakov, F. V. 2003. The IM Algorithm: A Variational Approach to Information Maximization. In *NIPS*, 201–208.
- Bernardo, J. M.; and Smith, A. F. 2009. *Bayesian theory*, volume 405. John Wiley & Sons.
- Bingham, E.; Chen, J. P.; Jankowiak, M.; Obermeyer, F.; Pradhan, N.; Karaletsos, T.; Singh, R.; Szerlip, P. A.; Horsfall, P.; and Goodman, N. D. 2019. Pyro: Deep Universal Probabilistic Programming. *J. Mach. Learn. Res.*, 20: 28:1–28:6.
- Bloem-Reddy, B.; and Teh, Y. W. 2020. Probabilistic Symmetries and Invariant Neural Networks. *J. Mach. Learn. Res.*, 21: 90–1.
- Chaloner, K.; and Verdinelli, I. 1995. Bayesian experimental design: A review. *Statistical Science*, 273–304.
- Diggle, P.; and Lophaven, S. 2006. Bayesian geostatistical design. *Scandinavian Journal of Statistics*, 33(1): 53–64.
- Dinh, L.; Sohl-Dickstein, J.; and Bengio, S. 2017. Density estimation using Real NVP. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Durkan, C.; Bekasov, A.; Murray, I.; and Papamakarios, G. 2019. Neural Spline Flows. In Wallach, H.; Larochelle, H.; Beygelzimer, A.; d'Alché-Buc, F.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Durkan, C.; Bekasov, A.; Murray, I.; and Papamakarios, G. 2020. nflows: normalizing flows in PyTorch. Zenodo.
- Foster, A.; Ivanova, D. R.; Malik, I.; and Rainforth, T. 2021. Deep adaptive design: Amortizing sequential bayesian experimental design. In *International Conference on Machine Learning*, 3384–3395. PMLR.
- Foster, A.; Jankowiak, M.; Bingham, E.; Horsfall, P.; Teh, Y. W.; Rainforth, T.; and Goodman, N. 2019. Variational Bayesian Optimal Experimental Design. In Wallach, H.; Larochelle, H.; Beygelzimer, A.; d'Alché-Buc, F.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Foster, A.; Jankowiak, M.; O'Meara, M.; Teh, Y. W.; and Rainforth, T. 2020. A unified stochastic gradient approach to designing bayesian-optimal experiments. In *International Conference on Artificial Intelligence and Statistics*, 2959–2969. PMLR.
- Goos, P.; and Jones, B. 2011. *Optimal design of experiments: a case study approach*. John Wiley & Sons.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Kobyzev, I.; Prince, S. J.; and Brubaker, M. A. 2020. Normalizing flows: An introduction and review of current methods. *IEEE transactions on pattern analysis and machine intelligence*, 43(11): 3964–3979.
- Liepe, J.; Filippi, S.; Komorowski, M.; and Stumpf, M. P. 2013. Maximizing the information content of experiments in systems biology. *PLoS computational biology*, 9(1): e1002888.
- McCullagh, P.; and Nelder, J. 1989. *Generalized Linear Models, Second Edition*. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. Taylor & Francis. ISBN 9780412317606.
- Myung, J. I.; Cavagnaro, D. R.; and Pitt, M. A. 2013. A tutorial on adaptive design optimization. *Journal of mathematical psychology*, 57(3-4): 53–67.
- Owen, A. B. 2013. Monte Carlo theory, methods and examples. Forthcoming.
- Papamakarios, G.; Nalisnick, E. T.; Rezende, D. J.; Mohamed, S.; and Lakshminarayanan, B. 2021. Normalizing Flows for Probabilistic Modeling and Inference. *J. Mach. Learn. Res.*, 22(57): 1–64.
- Rainforth, T.; Cornish, R.; Yang, H.; Warrington, A.; and Wood, F. 2018. On nesting monte carlo estimators. In *International Conference on Machine Learning*, 4267–4276. PMLR.
- Robert, C. P.; and Casella, G. 1999. *Monte Carlo statistical methods*, volume 2. Springer.
- Ryan, E. G.; Drovandi, C. C.; McGree, J. M.; and Pettitt, A. N. 2016. A review of modern computational algorithms for Bayesian optimal design. *International Statistical Review*, 84(1): 128–154.
- Sebastiani, P.; and Wynn, H. P. 2000. Maximum entropy sampling and optimal Bayesian experimental design. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 62(1): 145–157.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L. u.; and Polosukhin, I. 2017. Attention is All you Need. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Vincent, B. T.; and Rainforth, T. 2017. The DARC Toolbox: automated, flexible, and efficient delayed and risky choice experiments using Bayesian adaptive design. *PsyArXiv*.
- Zaheer, M.; Kottur, S.; Ravanbakhsh, S.; Póczos, B.; Salakhutdinov, R. R.; and Smola, A. J. 2017. Deep Sets. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.