

Efficient Real-world Testing of Causal Decision Making via Bayesian Experimental Design for Contextual Optimisation

Desi R. Ivanova[†]

Department of Statistics, University of Oxford

DESI.IVANOVA@STATS.OX.AC.UK

Joel Jennings

Microsoft Research Cambridge

JOELJENNINGS@MICROSOFT.COM

Cheng Zhang

Microsoft Research Cambridge

CHEZHA@MICROSOFT.COM

Adam Foster

Microsoft Research Cambridge

ADAM.E.FOSTER@MICROSOFT.COM

[†]Work done during an internship at Microsoft Research Cambridge

Abstract

The real-world testing of decisions made using causal machine learning models is an essential prerequisite for their successful application. We focus on evaluating and improving contextual treatment assignment decisions: these are personalised treatments applied to e.g. customers, each with their own contextual information, with the aim of maximising a reward. In this paper we introduce a model-agnostic framework for gathering data to evaluate and improve contextual decision making through Bayesian Experimental Design. Specifically, our method is used for the data-efficient evaluation of the *regret* of past treatment assignments. Unlike approaches such as A/B testing, our method avoids assigning treatments that are known to be highly sub-optimal, whilst engaging in some exploration to gather pertinent information. We achieve this by introducing an information-based design objective, which we optimise end-to-end. Our method applies to discrete and continuous treatments. Comparing our information-theoretic approach to baselines in several simulation studies demonstrates the superior performance of our proposed approach.

Keywords: Real-world testing, Bayesian experimental design, causal decision making evaluation, contextual optimisation, mutual information.

1. Introduction

Machine learning models, particularly causal machine learning models, can be used to make real-world decisions (Joyce, 1999; Geffner et al., 2022). We consider the problem of collecting experimental data to evaluate such models and improve decisions made in future. For example, suppose a company has developed a model describing the relationship between customer engagement and revenue. This model will be used to offer personalised promotions to each customer in order to maximise some reward, e.g. revenue. Before deploying it to all of its customers, the company may wish to collect experimental data to validate the model.

To formalise this, we consider the contextual optimisation problem in which a treatment or treatments must be assigned to each context to receive a reward. We focus on the experimental design problem of finding treatments that allow us to evaluate the regret of

past actions and to improve decisions made in the future. Our aim, then, is to choose actions to test the model; we are not directly optimising rewards in the experimental phase. In line with our real-world scenario, we focus on designing *large batch* experiments, in which we must select a large number of treatments for different contexts before receiving any feedback.

A standard experimental design procedure is random treatment assignment, which corresponds to A/B testing when treatment is binary (Fisher, 1936; Kohavi and Longbotham, 2017). However, this naive approach has significant drawbacks, as it may involve applying many sub-optimal actions, meaning that the *opportunity cost of experimentation* becomes high and experimental resources are wasted applying treatments that are known *a priori* to be sub-optimal. On the other hand, applying the best treatments based on existing beliefs—a pure exploitation strategy—may also be uninformative for model testing as it fails to account for uncertainty and the need to gather data where uncertainty is high.

To design an experiment in a more principled way, we turn to the framework of Bayesian experimental design (BED) (Lindley, 1956; Chaloner and Verdinelli, 1995; Foster et al., 2019). In contrast to standard BED, which maximises the information gained about all model parameters, we define the objective for experimental design as the *expected information gained about the maximum reward obtainable for a set of evaluation contexts (max-value EIG)*. This directly addresses the problem of gaining data to evaluate the regret of past decisions, because, given a past outcome, the unknown component of regret is *how much better the reward could have been*. The formulation is also relevant to the problem of gathering data (exploration) to obtain better rewards in later interactions (exploitation).

As a model-based framework, BED requires the specification of a Bayesian model (Gelman et al., 2013) relating contexts, treatments and rewards. Since the model makes predictions about interventional distributions, causality plays a central role in selecting a correct model, particularly when the model is pre-trained with observational data. One approach is to rely on assumptions about the causal graph (Pearl, 2009; Sharma et al., 2021). Alternatively, the model can incorporate Bayesian uncertainty about the causal structure (Heckerman et al., 1999; Annadani et al., 2021; Geffner et al., 2022). In this paper, we propose a model-agnostic approach to experimental design that accommodates this full range of Bayesian models; we only require *differentiability* of samples from the model.

To find the optimal experimental design for a given Bayesian model, we first show that the InfoNCE bound (van den Oord et al., 2018) can be used to estimate the max-value EIG objective. We then optimise the testing treatment assignment (experimental design) by gradients (Foster et al., 2020; Kleinegesse and Gutmann, 2020). Our method is applicable to multiple treatments, both continuous and discrete, using carefully tuned Gumbel–Softmax (Maddison et al., 2016; Jang et al., 2016) for the latter case. We show experimentally that our approach is effective on a number of synthetic models, learning intuitively correct large batch designs that outperform standard baselines by significant margins.

2. Problem set-up

We are interested in the efficient evaluation of the regret of past treatment assignment decisions. We formalise this by considering a contextual optimisation problem in which treatment, \mathbf{a} , is applied in context, \mathbf{c} , to receive reward y . For example, in the customer engagement scenario, \mathbf{c} would be the information we have about the customer that cannot

be directly acted upon, such as age, income and historical data of past interactions, \mathbf{a} would denote a particular choice of promotions that we offer to the customer, and the reward y would be the future revenue generated by the customer, net of promotion costs.

Expected regret We are interested in gathering data to evaluate the *expected regret* of the past action \mathbf{a} performed in a context of interest \mathbf{c}^* . This is defined as

$$R(\mathbf{a}, \mathbf{c}^*) = m(\mathbf{c}^*) - \mathbb{E}[y|\text{do}(\mathbf{a}), \mathbf{c}^*], \quad (1)$$

where $p(\cdot|\text{do}(\mathbf{a}), \mathbf{c}^*)$ denotes the interventional distribution when action \mathbf{a} is applied to context \mathbf{c}^* and $m(\mathbf{c}^*) = \max_{\mathbf{a}'} \mathbb{E}[y|\text{do}(\mathbf{a}'), \mathbf{c}^*]$ denotes the best obtainable outcome for context \mathbf{c}^* . Suppose that \mathbf{a} has been performed historically to \mathbf{c}^* , and we wish evaluate the regret of this treatment. The primary obstacle to evaluating this regret in the real world is that $m(\mathbf{c}^*)$ is not known, i.e. we do not know how much better the outcome could have been if a different treatment had been applied. We therefore wish to obtain experimental data, possibly from different contexts, that will help us to efficiently infer $m(\mathbf{c}^*)$. Gathering data to learn $m(\mathbf{c}^*)$ also aids in choosing the best treatment in *future* interactions with \mathbf{c}^* .

Bayesian model In general, we also do not know the true interventional distribution $p(y|\text{do}(\mathbf{a}), \mathbf{c})$. We therefore model the relationship between \mathbf{a}, \mathbf{c} and y in a Bayesian manner, by introducing a Bayesian parameter ψ with a prior $p(\psi)$. Different samples of ψ corresponds to different hypotheses $p(y|\text{do}(\mathbf{a}), \mathbf{c}, \psi)$, whilst marginalising over ψ summarises the total Bayesian uncertainty in $y|\text{do}(\mathbf{a}), \mathbf{c}$. Figure 1 illustrates an example with a binary treatment.

The role of causality Causal considerations are central to building a correct model of the interventional distribution $p(y|\text{do}(\mathbf{a}), \mathbf{c})$, particularly when the prior $p(\psi)$ is fitted by conditioning on past *observational* data. By making certain assumptions about the causal graph and the (non-)existence of unobserved confounders, one can reduce the problem to that of learning specific functions (Pearl, 2009; Sharma et al., 2021). These assumptions are often made implicitly in the optimisation literature (Bareinboim et al., 2015). Through suitable assumptions, any Bayesian model that predicts y from \mathbf{a}, \mathbf{c} can be used as a model for $p(y|\text{do}(\mathbf{a}), \mathbf{c})$, so a causal treatment of the problem does not restrict the class of models that can be considered. Alternatively, the Bayesian parameter ψ can incorporate uncertainty in the causal graph (Heckerman et al., 1999; Annadani et al., 2021; Geffner et al., 2022). Our approach to experimental design is concordant with either of these modelling choices.

3. Bayesian Experimental Design for Contextual Optimisation

Whilst standard BED would design experiments to maximise the information gathered about the Bayesian parameter ψ , our goal is to efficiently evaluate the regret (1). This translates to estimating the maximum achievable rewards for a set of D^* evaluation contexts $\mathbf{C}^* = \mathbf{c}_1^*, \dots, \mathbf{c}_{D^*}^*$, denoted $\mathbf{m}^* = m(\mathbf{c}_1^*), \dots, m(\mathbf{c}_{D^*}^*)$. In the customer engagement example, \mathbf{C}^* represents a set of customers for which we are aiming to infer the best possible rewards \mathbf{m}^* . To learn about \mathbf{m}^* , we are allowed to perform a batch of D experiments at contexts $\mathbf{C} = \mathbf{c}_1, \dots, \mathbf{c}_D$. In the customer engagement example, \mathbf{C} represents the set of customers that will participate in the real-world test. The *experimental design* problem is to select the treatments $\mathbf{A} = \mathbf{a}_1, \dots, \mathbf{a}_D$ to apply to \mathbf{C} so that the experimental outcomes $\mathbf{y} = y_1, \dots, y_D$

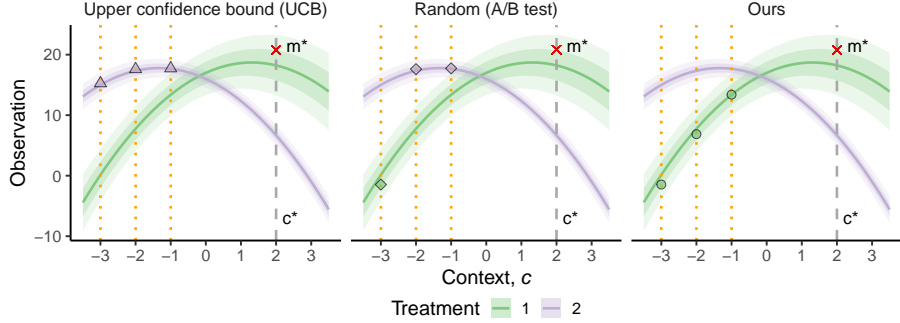


Figure 1: An example with a scalar context and a binary treatment $\mathbf{a} \in \{1, 2\}$. The dark (resp. light) shaded area shows our uncertainty about $y|\text{do}(\mathbf{a}), \mathbf{c}$ arising from uncertainty in ψ , measured by one (resp. two) st. dev. from the mean. We want to estimate $m^* = m(\mathbf{c}^*)$, the unknown maximum achievable reward at the evaluation context \mathbf{c}^* . We perform three experiments at $\mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_3$ (dotted orange lines). UCB chooses the reward-maximising treatments in each experimental context, bringing no information about m^* , as Treatment 2 is *a priori* suboptimal at \mathbf{c}^* . Random design selects Treatments 1 and 2 with equal probability; our method selects treatments whose outcomes will be most informative about \mathbf{m}^* given the prior.

that we get from the real-world test are maximally informative about \mathbf{m}^* . More formally, we design \mathbf{A} to maximize expected information gain between \mathbf{y} and \mathbf{m}^*

$$\mathcal{I}(\mathbf{A}; \mathbf{C}, \mathbf{C}^*) = \mathbb{E}_{p(\mathbf{y}, \mathbf{m}^* | \mathbf{C}, \mathbf{C}^*, \text{do}(\mathbf{A}))} \left[\log \frac{p(\mathbf{y} | \mathbf{m}^*, \text{do}(\mathbf{A}), \mathbf{C})}{p(\mathbf{y} | \text{do}(\mathbf{A}), \mathbf{C})} \right], \quad (2)$$

where $p(\mathbf{y} | \text{do}(\mathbf{A}), \mathbf{C}) = \mathbb{E}_{p(\psi)} [p(\mathbf{y} | \text{do}(\mathbf{A}), \mathbf{C}, \psi)]$, $p(\mathbf{m}^* | \mathbf{C}^*) = \mathbb{E}_{p(\psi)} [p(\mathbf{m}^* | \mathbf{C}^*, \psi)]$ and $p(\mathbf{y}, \mathbf{m}^* | \mathbf{C}, \mathbf{C}^*, \text{do}(\mathbf{A})) = \mathbb{E}_{p(\psi)} [p(\mathbf{m}^* | \mathbf{C}^*, \psi) p(\mathbf{y} | \mathbf{C}, \text{do}(\mathbf{A}), \psi)]$. We refer to $\mathcal{I}(\mathbf{A}; \mathbf{C}, \mathbf{C}^*)$ as the max-value EIG. The objective is doubly intractable (Rainforth et al., 2018) meaning that its optimisation, or even evaluation with a fixed \mathbf{A} , is a major challenge. The likelihoods in the expectation are *implicit* (not available analytically) making the problem more difficult.

3.1 Training stage: Objective estimation and optimization

Recently, gradient-based approaches for optimising EIG have been developed for experimental design for parameter learning (Foster et al., 2020; Kleingesse and Gutmann, 2020) as opposed to contextual optimisation, which is the focus of this work. Inspired by likelihood-free mutual information estimators, we show the InfoNCE bound (van den Oord et al., 2018) can be adapted to the contextual optimisation setting we consider here. Concretely,

$$\mathcal{L}(\mathbf{A}, U; \mathbf{C}, \mathbf{C}^*, L) = \mathbb{E}_{p(\psi) p(\mathbf{y}, \mathbf{m}_0^* | \text{do}(\mathbf{A}), \mathbf{C}, \mathbf{C}^*, \psi) p(\mathbf{m}_{1:L}^* | \mathbf{C}^*)} \left[\log \frac{\exp(U(\mathbf{y}, \mathbf{m}_0^*))}{\frac{1}{L+1} \sum_{\ell} \exp(U(\mathbf{y}, \mathbf{m}_{\ell}^*))} \right] \quad (3)$$

is a lower bound on $\mathcal{I}(\mathbf{A}; \mathbf{C}, \mathbf{C}^*)$ for any $L \geq 1$ and any critic function U . We provide a proof of this claim in Appendix A. To learn a suitable critic U , we use a neural network with trainable parameters ϕ . We then optimise the lower bound $\mathcal{L}(\mathbf{A}, U; \mathbf{C}, \mathbf{C}^*, L)$ simultaneously with respect to ϕ and \mathbf{A} , thereby improving the tightness of the bound and optimising the experimental design together. When the actions \mathbf{A} are continuous, we require that a sample

y of $p(y|\text{do}(\mathbf{a}), \mathbf{c}, \psi)$ is differentiable with respect to \mathbf{a} . In this case, a pathwise gradient estimator (Mohamed et al., 2020) for $\nabla_{\mathbf{A}, \phi} \mathcal{L}$ can be readily computed and we apply standard stochastic gradient ascent (SGA) (Robbins and Monro, 1951) to optimise our objective (3).

Discrete action space Previous gradient-based BED work has focused on fully differentiable models. Here we propose a practical way to deal with discrete designs. Rather than learning the treatments \mathbf{A} directly, we introduce a stochastic treatment policy π_{α} with trainable parameters α representing the probabilities of selecting each treatment. During training we use a Gumbel–Softmax (Maddison et al., 2016; Jang et al., 2016) relaxed form of \mathbf{A} and update α by gradients. Further detail is provided in Appendix A.

3.2 Deployment stage: Real-world testing

Once we optimise the objective (3), we perform the batch of experiments \mathbf{A} to obtain real-world outcomes \mathbf{y} and estimate $p(\psi|\mathcal{D})$ —the posterior of our Bayesian model, given the experimental data $\mathcal{D} = (\mathbf{y}, \mathbf{A}, \mathbf{C})$. We then calculate the posterior $p(\mathbf{m}^*|\mathbf{C}^*, \mathcal{D})$, and this is to estimate the regret (1) of a set of past treatments. Importantly, through our design of \mathbf{A} , the data contained in \mathcal{D} should lead to the *most accurate* estimate of \mathbf{m}^* , which translates to an accurate estimate of the regret and an accurate evaluation metric for the original decisions. Algorithm 1 in the Appendix gives a complete description of our method.

4. Related Work

The most closely related objective to our max-value EIG (2) is Max-value Entropy Search (MES, Wang and Jegelka, 2017). It was proposed as a computationally efficient alternative to Entropy Search (ES, Hennig and Schuler, 2012) and Predictive Entropy Search (PES, Hernández-Lobato et al., 2014). Our objective differs from MES in two essential ways: firstly, MES is applicable to Gaussian processes (GPs), whilst our objective is model-agnostic; secondly, we focus on *contextual* optimisation, rather than finding a single maximiser. For a more comprehensive discussion of related work, see Appendix B.

5. Empirical Evaluation

We test the efficacy of our method on several synthetic experiments. In all synthetic experiments, we assume the causal graph is as shown in Figure 2, but note that our method is not restricted to this graph structure. We compare our method against random treatment assignment and the Upper Confidence Bound (UCB) (Auer, 2002) algorithm. Our key evaluation metrics are **EIG** and **accuracy** of inferring \mathbf{m}^* . For the former we evaluate (3) with the trained critic and learnt designs. For the latter, we sample a realisation $\tilde{\psi}$ and $\tilde{\mathbf{m}}^*$ from the prior, treating it as the ground truth environment in the Deployment stage of our framework (§ 3.2). We report the mean squared error (MSE) between our estimate \mathbf{m}^* and the ground truth $\tilde{\mathbf{m}}^*$. We report three additional metrics—average regret, MSE of estimating $\tilde{\psi}$, and deviation from the true optimal treatment. Full details about the models, training, and further results are given in Appendix C.




Figure 2: Causal graph.

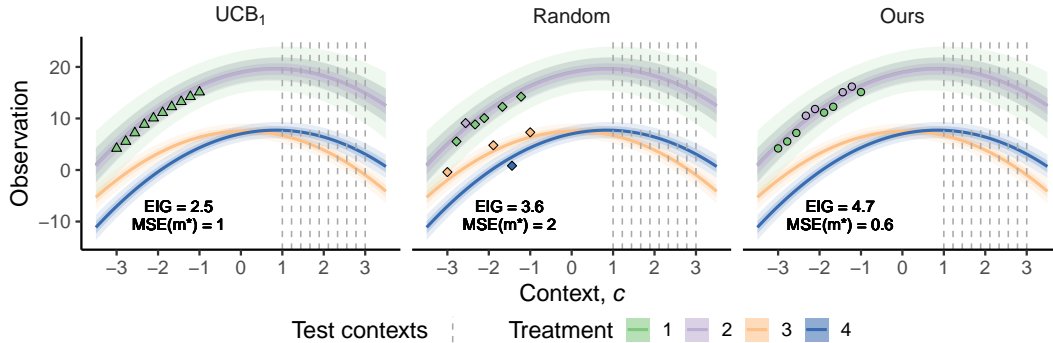


Figure 3: Discrete treatments: optimal designs and key metrics. More details in Appendix C.

Method	EIG estimate	MSE(\mathbf{m}^*)	MSE(ψ)	MSE(\mathbf{A})	Regret
Random _{0.2}	5.407 ± 0.003	0.0041 ± 0.0001	0.0114 ± 0.0003	0.544 ± 0.023	0.091 ± 0.002
Random _{1.0}	5.798 ± 0.004	0.0024 ± 0.0002	0.0053 ± 0.0002	0.272 ± 0.018	0.060 ± 0.002
Random _{2.0}	4.960 ± 0.004	0.0042 ± 0.0002	0.0102 ± 0.0003	0.450 ± 0.021	0.090 ± 0.002
UCB ₀	5.774 ± 0.003	0.0069 ± 0.0005	0.0154 ± 0.0007	0.747 ± 0.055	0.082 ± 0.002
UCB ₁	5.876 ± 0.003	0.0030 ± 0.0002	0.0102 ± 0.0004	0.338 ± 0.024	0.067 ± 0.002
UCB ₂	5.780 ± 0.004	0.0031 ± 0.0002	0.0093 ± 0.0003	0.378 ± 0.031	0.069 ± 0.002
Ours	6.527 ± 0.003	0.0014 ± 0.0001	0.0039 ± 0.0002	0.143 ± 0.018	0.044 ± 0.001

Table 1: Continuous treatments: evaluation of 40D design. More details in Appendix C.

Discrete treatments We consider a model with four treatment options, two of which are *a priori* sub-optimal; the other two have the same mean but Treatment 1 has higher variance than Treatment 2 (Figure 3). Intuitively, the optimal strategy would be to A/B test only the top 2 treatments, which is exactly what our method has learnt. The random method is wasteful as it queries the sub-optimal treatments, while UCB₁ only ever queries Treatment 1. As a result, our method is much more accurate at estimating the true \mathbf{m}^* .

Continuous treatments We test our method on a continuous treatment problem, where we design a batch of 40 experiments to learn about the max-values at 39 evaluation contexts. As Table 1 shows, our method outperforms the baselines on all metrics. Although our method is optimised to accurately estimate $\tilde{\mathbf{m}}^*$, it is also significantly better at estimating $\tilde{\psi}$, which in turn translates to choosing better treatments in the future, as the lower MSE(\mathbf{A}) indicates.

6. Discussion and Conclusion

By combining ideas from Bayesian experimental design, contextual bandits and Bayesian optimisation, we proposed a method for real-world model testing by regret evaluation of past context-dependent treatment assignments. Our method also applies to gathering data to improve future decision making, in fact, we saw that gaining information for regret evaluation and for model improvement are really two sides of the same coin. To compute large batch designs efficiently, we cast the problem in terms of implicit likelihood BED, resulting in a model agnostic approach that is able to handle both discrete and continuous designs.

Future work may introduce the ability to handle constraints, and explore the scaling of our method to high dimensional treatments and contexts, both continuous and discrete. Our method could also be extended to apply to sequential, adaptive experimentation.

References

- Shipra Agrawal and Navin Goyal. Thompson sampling for contextual bandits with linear payoffs. In *International conference on machine learning*, pages 127–135. PMLR, 2013. 12
- Yashas Annadani, Jonas Rothfuss, Alexandre Lacoste, Nino Scherrer, Anirudh Goyal, Yoshua Bengio, and Stefan Bauer. Variational causal networks: Approximate bayesian inference over causal structures. *arXiv preprint arXiv:2106.07635*, 2021. 2, 3
- Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3(Nov):397–422, 2002. 5, 12
- Elias Bareinboim, Andrew Forney, and Judea Pearl. Bandits with unobserved confounders: A causal approach. *Advances in Neural Information Processing Systems*, 28, 2015. 3
- Eli Bingham, Jonathan P Chen, Martin Jankowiak, Fritz Obermeyer, Neeraj Pradhan, Theofanis Karaletsos, Rohit Singh, Paul Szerlip, Paul Horsfall, and Noah D Goodman. Pyro: Deep universal probabilistic programming. *Journal of Machine Learning Research*, 2018. 12
- Kathryn Chaloner and Isabella Verdinelli. Bayesian experimental design: A review. *Statistical Science*, pages 273–304, 1995. 2
- Ian Char, Youngseog Chung, Willie Neiswanger, Kirthevasan Kandasamy, Andrew O Nelson, Mark Boyer, Egemen Kolemen, and Jeff Schneider. Offline contextual bayesian optimization. *Advances in Neural Information Processing Systems*, 32, 2019. 12
- Wei Chu, Lihong Li, Lev Reyzin, and Robert Schapire. Contextual bandits with linear payoff functions. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 208–214. JMLR Workshop and Conference Proceedings, 2011. 12
- Ronald Aylmer Fisher. Design of experiments. *British Medical Journal*, 1(3923):554, 1936. 2
- Adam Foster, Martin Jankowiak, Elias Bingham, Paul Horsfall, Yee Whye Teh, Thomas Rainforth, and Noah Goodman. Variational Bayesian Optimal Experimental Design. In *Advances in Neural Information Processing Systems 32*, pages 14036–14047. Curran Associates, Inc., 2019. 2
- Adam Foster, Martin Jankowiak, Matthew O’Meara, Yee Whye Teh, and Tom Rainforth. A unified stochastic gradient approach to designing bayesian-optimal experiments. In *International Conference on Artificial Intelligence and Statistics*, pages 2959–2969. PMLR, 2020. 2, 4, 11, 12
- Tomas Geffner, Javier Antoran, Adam Foster, Wenbo Gong, Chao Ma, Emre Kiciman, Amit Sharma, Angus Lamb, Martin Kukla, Nick Pawlowski, et al. Deep end-to-end causal inference. *arXiv preprint arXiv:2202.02195*, 2022. 1, 2, 3
- Andrew Gelman, John B Carlin, Hal S Stern, David B Dunson, Aki Vehtari, and Donald B Rubin. *Bayesian data analysis*. Chapman and Hall/CRC, 2013. 2

- David Ginsbourger, Jean Baccou, Clément Chevalier, Frédéric Perales, Nicolas Garland, and Yann Monerie. Bayesian adaptive reconstruction of profile optima and optimizers. *SIAM/ASA Journal on Uncertainty Quantification*, 2(1):490–510, 2014. 12
- YanJun Han, Zhengqing Zhou, Zhengyuan Zhou, Jose Blanchet, Peter W Glynn, and Yinyu Ye. Sequential batch learning in finite-action linear contextual bandits. *arXiv preprint arXiv:2004.06321*, 2020. 12
- David Heckerman, Christopher Meek, and Gregory Cooper. A bayesian approach to causal discovery. *Computation, causation, and discovery*, 19:141–166, 1999. 2, 3
- Philipp Hennig and Christian J Schuler. Entropy search for information-efficient global optimization. *Journal of Machine Learning Research*, 13(Jun):1809–1837, 2012. 5
- José Miguel Hernández-Lobato, Matthew W Hoffman, and Zoubin Ghahramani. Predictive entropy search for efficient global optimization of black-box functions. In *Advances in neural information processing systems*, pages 918–926, 2014. 5
- Desi R Ivanova, Adam Foster, Steven Kleinegesse, Michael Gutmann, and Tom Rainforth. Implicit Deep Adaptive Design: Policy-Based Experimental Design without Likelihoods. In *Advances in Neural Information Processing Systems*, volume 34, pages 25785–25798. Curran Associates, Inc., 2021. URL <https://proceedings.neurips.cc/paper/2021/file/d811406316b669ad3d370d78b51b1d2e-Paper.pdf>. 12
- Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016. 2, 5
- James M Joyce. *The foundations of causal decision theory*. Cambridge University Press, 1999. 1
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 12
- Steven Kleinegesse and Michael Gutmann. Bayesian experimental design for implicit models by mutual information neural estimation. In *Proceedings of the 37th International Conference on Machine Learning*, Proceedings of Machine Learning Research, pages 5316–5326. PMLR, 2020. 2, 4, 12
- Ron Kohavi and Roger Longbotham. Online controlled experiments and a/b testing. *Encyclopedia of machine learning and data mining*, 7(8):922–929, 2017. 2
- Andreas Krause and Cheng Ong. Contextual gaussian process bandit optimization. *Advances in neural information processing systems*, 24, 2011. 12
- Dennis V Lindley. On a measure of the information provided by an experiment. *The Annals of Mathematical Statistics*, pages 986–1005, 1956. 2
- Chris J Maddison, Andriy Mnih, and Yee Whye Teh. The concrete distribution: A continuous relaxation of discrete random variables. *arXiv preprint arXiv:1611.00712*, 2016. 2, 5

- Shakir Mohamed, Mihaela Rosca, Michael Figurnov, and Andriy Mnih. Monte carlo gradient estimation in machine learning. *Journal of Machine Learning Research*, 21(132):1–62, 2020. 5
- Michael Pearce, Janis Klaise, and Matthew Groves. Practical bayesian optimization of objectives with conditioning variables. *arXiv preprint arXiv:2002.09996*, 2020. 12
- Judea Pearl. Causal inference in statistics: An overview. *Statistics surveys*, 3:96–146, 2009. 2, 3
- Ben Poole, Sherjil Ozair, Aäron van den Oord, Alex Alemi, and George Tucker. On variational bounds of mutual information. In *International Conference on Machine Learning*, pages 5171–5180, 2019. 12
- Tom Rainforth, Rob Cornish, Hongseok Yang, Andrew Warrington, and Frank Wood. On nesting monte carlo estimators. In *International Conference on Machine Learning*, pages 4267–4276. PMLR, 2018. 4
- Herbert Robbins and Sutton Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951. 5
- Amit Sharma, Vasilis Syrgkanis, Cheng Zhang, and Emre Kıcıman. Dowhy: Addressing challenges in expressing and validating causal assumptions. *arXiv preprint arXiv:2108.13518*, 2021. 2, 3
- Aäron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 2, 4
- Zi Wang and Stefanie Jegelka. Max-value entropy search for efficient bayesian optimization. In *International Conference on Machine Learning*, pages 3627–3635. PMLR, 2017. 5
- Andrea Zanette, Kefan Dong, Jonathan N Lee, and Emma Brunskill. Design of experiments for stochastic contextual linear bandits. *Advances in Neural Information Processing Systems*, 34:22720–22731, 2021. 12

Algorithm 1: Data-efficient regret evaluation

Input: Bayesian simulator $p(y, m^* | \text{do}(\mathbf{a}), \mathbf{c}, \mathbf{c}^*, \psi)$, batch size D of experimental contexts \mathbf{C} , batch size D^* of evaluation contexts \mathbf{C}^* , initial \mathbf{A} , initial U_ϕ , number of contrastive samples $L \geq 1$

Output: Optimal treatments \mathbf{A} for experimental contexts \mathbf{C} to be tested

Training stage:

while Computational training budget not exceeded **do**

 Sample $\psi \sim p(\psi)$

 Sample $\mathbf{y} \sim p(\mathbf{y} | \mathbf{C}, \text{do}(\mathbf{A}), \psi)$ and $\mathbf{m}^* \sim p(\mathbf{m}^* | \mathbf{C}^*, \psi)$

 Compute $\mathcal{L}(\mathbf{A}, U_\phi; L)$ (3) and update the parameters (\mathbf{A}, ϕ) using a SGA

end

Deployment stage:

 Run a batch of experiments \mathbf{A} to obtain real-world observations \mathbf{y} and estimate a posterior $p(\psi | \mathcal{D})$, $\mathcal{D} = \{\mathbf{y}, \mathbf{A}, \mathbf{C}\}$.

 Use the updated model to get an estimate of \mathbf{m}^* and evaluate the regret (1) associated with some (past) treatment assignment in contexts \mathbf{C}^* .

Appendix A. Method

For completeness we formalise the claim made in Section 3.1 in Proposition 1, and provide a proof, following standard arguments.

Proposition 1 (InfoNCE lower bound) *For any critic function $U : \mathcal{Y} \times \mathbb{R} \rightarrow \mathbb{R}$ and number of contrastive samples $L \geq 1$ we have $\mathcal{L}(\mathbf{A}; \mathbf{C}, \mathbf{C}^*, L) \leq \mathcal{I}(\mathbf{A}; \mathbf{C}, \mathbf{C}^*)$, where*

$$\mathcal{L}(\mathbf{A}, U; \mathbf{C}, \mathbf{C}^*, L) = \mathbb{E}_{p(\psi)p(\mathbf{y}, \mathbf{m}_0^* | \mathbf{C}, \mathbf{C}^*, \text{do}(\mathbf{A}), \psi)p(\mathbf{m}_{1:L}^*, \psi_{1:L} | \mathbf{C}^*)} \left[\log \frac{\exp(U(\mathbf{y}, \mathbf{m}_0^*))}{\frac{1}{L+1} \sum_{\ell} \exp(U(\mathbf{y}, \mathbf{m}_\ell^*))} \right] \quad (4)$$

The bound is tight for the optimal critic $U^*(\mathbf{y}, \mathbf{m}^*) = \log p(\mathbf{y} | \mathbf{m}^*, \text{do}(\mathbf{A}), \mathbf{C}) + c(\mathbf{y})$ as $L \rightarrow \infty$, where $c(\mathbf{y})$ is an arbitrary function depending only on the outcomes \mathbf{y} .

Proof Let $U : \mathcal{Y} \times \mathbb{R} \rightarrow \mathbb{R}$ be any function (critic) and introduce the shorthand

$$g(\mathbf{y}, \mathbf{m}_{0:L}^*) := \frac{\exp(U(\mathbf{y}, \mathbf{m}_0^*))}{\frac{1}{L+1} \sum_{\ell} \exp(U(\mathbf{y}, \mathbf{m}_\ell^*))} \quad (5)$$

Multiply the mutual information objective of Equation (2) by $g(\mathbf{y}, \mathbf{m}_{0:L}^*)/g(\mathbf{y}, \mathbf{m}_{0:L}^*)$ to get

$$\mathcal{I}(\mathbf{A}; \mathbf{C}, \mathbf{C}^*) = \mathbb{E}_{p(\mathbf{y}, \mathbf{m}^* | \mathbf{C}, \mathbf{C}^*, \text{do}(\mathbf{A}))} \left[\log \frac{p(\mathbf{y}, \mathbf{m}^* | \mathbf{C}, \mathbf{C}^*, \text{do}(\mathbf{A}))}{p(\mathbf{y} | \mathbf{C}, \text{do}(\mathbf{A}))p(\mathbf{m}^* | \mathbf{C}^*)} \right] \quad (6)$$

$$= \mathbb{E}_{p(\psi)p(\mathbf{y}, \mathbf{m}_0^* | \mathbf{C}, \mathbf{C}^*, \text{do}(\mathbf{A}), \psi)} \left[\log \frac{p(\mathbf{y}, \mathbf{m}_0^* | \mathbf{C}, \mathbf{C}^*, \text{do}(\mathbf{A}))}{p(\mathbf{y} | \mathbf{C}, \text{do}(\mathbf{A}))p(\mathbf{m}_0^* | \mathbf{C}^*)} \right] \quad (7)$$

$$= \mathbb{E}_{p(\psi)p(\mathbf{y}, \mathbf{m}_0^* | \mathbf{C}, \mathbf{C}^*, \text{do}(\mathbf{A}), \psi)p(\mathbf{m}_{1:L}^*, \psi_{1:L} | \mathbf{C}^*)} \left[\log \frac{p(\mathbf{y} | \mathbf{m}_0^*, \text{do}(\mathbf{A}), \mathbf{C})}{p(\mathbf{y} | \text{do}(\mathbf{A}), \mathbf{C})} \right] \quad (8)$$

$$= \mathbb{E}_{p(\psi)p(\mathbf{y}, \mathbf{m}_0^* | \mathbf{C}, \mathbf{C}^*, \text{do}(\mathbf{A}), \psi)p(\mathbf{m}_{1:L}^*, \psi_{1:L} | \mathbf{C}^*)} \left[\log \frac{p(\mathbf{y} | \mathbf{m}_0^*, \text{do}(\mathbf{A}), \mathbf{C})g(\mathbf{y}, \mathbf{m}_{0:L}^*)}{p(\mathbf{y} | \text{do}(\mathbf{A}), \mathbf{C})g(\mathbf{y}, \mathbf{m}_{0:L}^*)} \right] \quad (9)$$

which we can split into two expectations—one that does not contain the implicit likelihoods and another that is equal to $\mathcal{L}(\mathbf{A}, U; \mathbf{C}, \mathbf{C}^*, L)$

$$\begin{aligned}
&= \mathbb{E}_{p(\psi)p(\mathbf{y}, \mathbf{m}^* | \mathbf{C}, \mathbf{C}^*, \text{do}(\mathbf{A}), \psi)p(\mathbf{m}_{1:L}^* | \psi_{1:L}, \mathbf{C}^*)} \left[\log \frac{p(\mathbf{y} | \mathbf{m}_0^*, \text{do}(\mathbf{A}), \mathbf{C})}{p(\mathbf{y} | \text{do}(\mathbf{A}), \mathbf{C})g(\mathbf{y}, \mathbf{m}_{0:L}^*)} \right] \\
&\quad + \mathbb{E}_{p(\psi)p(\mathbf{y}, \mathbf{m}^* | \mathbf{C}, \mathbf{C}^*, \text{do}(\mathbf{A}), \psi)p(\mathbf{m}_{1:L}^* | \psi_{1:L}, \mathbf{C}^*)} [\log g(\mathbf{y}, \mathbf{m}_{0:L}^*)] \\
&= \text{KL}(p_1 \parallel p_2) + \mathcal{L}(\mathbf{A}, U; \mathbf{C}, \mathbf{C}^*, L) \\
&\geq \mathcal{L}(\mathbf{A}, U; \mathbf{C}, \mathbf{C}^*, L)
\end{aligned} \tag{10}$$

where

$$\text{KL}(p_1 \parallel p_2) = \mathbb{E}_{p(\psi)p(\mathbf{y}, \mathbf{m}_0^* | \mathbf{C}, \mathbf{C}^*, \text{do}(\mathbf{A}), \psi)p(\psi_{1:L})p(\mathbf{m}_{1:L}^* | \psi_{1:L}, \mathbf{C}^*)} \left[\log \frac{p(\mathbf{y} | \mathbf{m}_0^*, \text{do}(\mathbf{A}), \mathbf{C})}{p(\mathbf{y} | \text{do}(\mathbf{A}), \mathbf{C})g(\mathbf{y}, \mathbf{m}_{0:L}^*)} \right] \tag{11}$$

$$= \mathbb{E} \left[\log \frac{p(\mathbf{y}, \mathbf{m}_0^* | \mathbf{C}, \mathbf{C}^*, \text{do}(\mathbf{A}))p(\mathbf{m}_{1:L}^* | \mathbf{C}^*)}{p(\mathbf{m}_0^* | \mathbf{C}^*)p(\mathbf{y} | \text{do}(\mathbf{A}), \mathbf{C})g(\mathbf{y}, \mathbf{m}_{0:L}^*)p(\mathbf{m}_{1:L}^* | \mathbf{C}^*)} \right] \tag{12}$$

where the expectation is with respect to $p(\psi)p(\mathbf{y} | \mathbf{C}, \text{do}(\mathbf{A}), \psi)p(\mathbf{m}_0^* | \psi, \mathbf{C}^*)p(\psi_{1:L})p(\mathbf{m}_{1:L}^* | \psi_{1:L}, \mathbf{C}^*)$. This is a valid KL divergence (since $p_2 = p(\mathbf{m}_0^* | \mathbf{C}^*)p(\mathbf{y} | \text{do}(\mathbf{A}), \mathbf{C})g(\mathbf{y}, \mathbf{m}_{0:L}^*)p(\mathbf{m}_{1:L}^* | \mathbf{C}^*)$ integrates to one by symmetry argument) and hence non-negative.

Finally, substituting $U^*(\mathbf{y}, \mathbf{m}^*) = \log p(\mathbf{y} | \mathbf{m}^*, \text{do}(\mathbf{A}), \mathbf{C}) + c(\mathbf{y})$ in the definition of \mathcal{L}

$$\mathcal{L}(\mathbf{A}, U^*; \mathbf{C}, \mathbf{C}^*, L) = \mathbb{E}_{p(\psi)p(\mathbf{y}, \mathbf{m}_0^* | \text{do}(\mathbf{A}), \mathbf{C}, \mathbf{C}^*, \psi)p(\mathbf{m}_{1:L}^* | \psi_{1:L}, \mathbf{C}^*)} \left[\log \frac{p(\mathbf{y} | \mathbf{m}^*, \text{do}(\mathbf{A}), \mathbf{C})}{\frac{1}{L+1} \sum_{\ell} p(\mathbf{y} | \mathbf{m}_{\ell}^*, \text{do}(\mathbf{A}), \mathbf{C})} \right] \tag{13}$$

which is monotonically increasing in L and tight in the limit as $L \rightarrow \infty$ since

$$\frac{1}{L+1} \sum_{\ell} p(\mathbf{y} | \mathbf{m}_{\ell}^*, \text{do}(\mathbf{A}), \mathbf{C}) \rightarrow p(\mathbf{y} | \text{do}(\mathbf{A}), \mathbf{C}) \text{ a.s.} \tag{14}$$

(See e.g. Foster et al., 2020, Theorem 1 for more detail). ■

Discrete action space Given $K \geq 2$ possible treatments, we learn a (non-deterministic) policy π with parameters α

$$\pi_{d,i} = \frac{\exp((\log \alpha_{d,i} + g_{d,i})/\tau)}{\sum_{j=1}^K \exp((\log \alpha_{d,j} + g_{d,j})/\tau)}, \quad i = 1, \dots, K, \quad g_{d,i} \sim \text{Gumbel}(0, 1), \quad \tau > 0, \tag{15}$$

where the parameter τ is called the temperature, which we anneal during optimisation—starting the optimisation at high temperatures, when gradients exhibit low variance, and gradually reducing it towards 0, at which stage gradients are high variance, but we (should) have already learnt good treatments. We optimise the parameters α and those of the critic network ϕ jointly with SGA. Once the policy is trained, the optimal design for experiment d in the batch is $\mathbf{a}_d = \arg \max(\pi_{d,1}, \dots, \pi_{d,K})$.

Appendix B. Additional Related Work

Our setting is related to contextual Bayesian optimisation; some of the more closely related methods include Profile Expected Improvement (PEI, Ginsbourger et al., 2014), Multi-task Thompson Sampling (MTS, Char et al., 2019) and conditional Bayesian optimization (ConBO, Pearce et al., 2020). All of these methods are, however, restricted to GPs and the criteria they use to choose optimal designs are not information-based.

Contextual bandits is another broad framework that our work is related to. An extensive line of research is focused on *online linear* bandits and discrete actions chosen using (variations of) UCB, Thompson sampling or ϵ -greedy strategy (Auer, 2002; Chu et al., 2011; Agrawal and Goyal, 2013; Han et al., 2020). Krause and Ong (2011) instead model the reward as a GP defined over the context-action space and develop CGP-UCB. More recently, Zanette et al. (2021) proposed designing a batch of experiments *offline* to collect a good dataset from which to learn a policy. Although this is similar to our setting, our approach is not restricted to linear rewards and uses an information-theoretic design objective.

Similar variational EIG objectives have been used in implicit likelihood BED methods for *parameter learning*, but not for contextual optimisation. Gradient-based methods for large batch experimentation include SG-BOED (Foster et al., 2020) and MINEBED (Kleinegesse and Gutmann, 2020), while the policy-based iDAD (Ivanova et al., 2021) applies to batch and adaptive settings. Our ability to handle discrete designs is another important distinction of our method.

Appendix C. Experiments

C.1 Training details

We implement all experiments in Pyro (Bingham et al., 2018). All experiments baselines ran for 50K gradient steps, using a batch size of 2048. We used the Adam optimiser (Kingma and Ba, 2014) with initial learning rate 0.001 and exponential learning rate annealing with coefficient 0.96 applied every 1000 steps. We used a separable critic architecture (Poole et al., 2019) with simple MLP encoders with ReLU activations and 32 output units.

For the discrete treatment example: we added *batch norm* to the critic architecture, which helped to stabilise the optimisation. We had one hidden layer of size 512. Additionally, for the Gumbel–Softmax policy, we started with a temperature $\tau = 2.0$ and `hard=False` constraint. We applied temperature annealing every 10K steps with a factor 0.5. We switch to `hard=True` in the last 10K steps of training.

For the continuous treatment example: We used MLPs with hidden layers of sizes [design dimension \times 2, 412, 256] and 32 output units.

Note: In order to evaluate the EIG of various baselines, we train a critic network for each one of them with the same hyperparameters as above.

C.2 Posterior inference details

After completing the training stage of our method (Algorithm 1), we need to deploy the learnt optimal designs in the real world in order to obtain rewards \mathbf{y} . This experimental data is then used to fit a posterior $p(\psi|\mathcal{D})$.

Method	EIG estimate	MSE(\mathbf{m}^*)	MSE(ψ)	Hit rate(\mathbf{A})	Regret
UCB _{0.0}	1.735 \pm 0.005	2.541 \pm 0.104	3.129 \pm 0.054	0.513 \pm 0.01	1.170 \pm 0.036
UCB _{1.0}	2.514 \pm 0.006	1.003 \pm 0.043	1.521 \pm 0.021	0.496 \pm 0.01	1.119 \pm 0.035
UCB _{2.0}	2.504 \pm 0.006	0.965 \pm 0.045	1.486 \pm 0.021	0.497 \pm 0.01	1.169 \pm 0.037
Random	3.573 \pm 0.006	1.953 \pm 0.070	1.347 \pm 0.024	0.503 \pm 0.01	1.150 \pm 0.036
Ours	4.729 \pm 0.009	0.594 \pm 0.025	1.326 \pm 0.020	0.501 \pm 0.01	1.152 \pm 0.035

Table 2: Discrete treatments: evaluation metrics of 10D design.

There are many ways to do the posterior inference and quality of the results will crucially depend on the accuracy of the fitted posteriors. In both of our examples and for all baselines we use Pyro’s self-normalised importance sampling (SNIS). Samples from this posterior are used for the evaluation metrics.

We validate the accuracy of estimated posteriors by running various sanity checks, including diagnostic plots such as Figure 4, showing the standard deviation of our posterior mean estimate (measure of how uncertainty about the parameter) and L_2 error to the true parameter. The red line shows the rolling mean over 200 points of the latter, and the gray band—the 2 standard deviations. For this plot we used the continuous treatments example with $D = 20$ experimental contexts.




Figure 4: Posterior checks

C.3 Evaluation metrics details

As discussed in the main text, we evaluate how well we can estimate \mathbf{m}^* by sampling a ground truth $\tilde{\psi}$ from the prior and obtaining a corresponding ground truth $\tilde{\mathbf{m}}^*$. We approximate the max-values \mathbf{m}^* empirically using 2000 posterior samples of ψ . We similarly estimate ψ using 2000 posterior samples. We define the optimal treatment under the posterior model to be average (with respect to that posterior) optimal treatment when treatments are continuous or UCB₀ when treatments are discrete. Finally the regret is computed as the average difference between the true max value (from the true environment and the true optimal treatment) and the one obtained by applying the estimated optimal treatment. We used 4000 (resp. 2000) true environment realisation for the continuous (resp. discrete) example.

C.4 Discrete treatments example

Model We first give details about the toy model we consider in Figure 3. Each of the four treatments $\mathbf{a} = 1, 2, 3, 4$ is a random function with two parameters $\psi_k = (\psi_{k,1}, \psi_{k,2})$ with the

Method	EIG estimate	MSE(\mathbf{m}^*)	MSE(ψ)	Hit rate(\mathbf{A})	Regret
UCB _{0.0}	1.740 \pm 0.003	2.709 \pm 0.058	3.197 \pm 0.018	0.500 \pm 0.005	1.150 \pm 0.017
UCB _{1.0}	2.508 \pm 0.002	0.993 \pm 0.016	1.529 \pm 0.007	0.498 \pm 0.004	1.140 \pm 0.007
UCB _{2.0}	2.505 \pm 0.006	0.991 \pm 0.023	1.518 \pm 0.012	0.497 \pm 0.003	1.145 \pm 0.015
Random	3.573 \pm 0.333	2.369 \pm 0.382	1.756 \pm 0.269	0.502 \pm 0.003	1.166 \pm 0.008
Ours	4.769 \pm 0.048	0.628 \pm 0.025	1.369 \pm 0.014	0.502 \pm 0.005	1.160 \pm 0.021

Table 3: Discrete treatments example: 10D design, stability across training seeds

following Gaussian priors (parameterised by mean and covariance matrix):

$$\psi_1 \sim \mathcal{N}\left(\begin{pmatrix} 5.00 \\ 15.0 \end{pmatrix}, \begin{pmatrix} 9.00 & 0 \\ 0 & 9.00 \end{pmatrix}\right) \quad \psi_2 \sim \mathcal{N}\left(\begin{pmatrix} 5.00 \\ 15.0 \end{pmatrix}, \begin{pmatrix} 2.25 & 0 \\ 0 & 2.25 \end{pmatrix}\right) \quad (16)$$

$$\psi_3 \sim \mathcal{N}\left(\begin{pmatrix} -2.0 \\ -1.0 \end{pmatrix}, \begin{pmatrix} 1.21 & 0 \\ 0 & 1.21 \end{pmatrix}\right) \quad \psi_4 \sim \mathcal{N}\left(\begin{pmatrix} -7.0 \\ 3.0 \end{pmatrix}, \begin{pmatrix} 1.21 & 0 \\ 0 & 1.21 \end{pmatrix}\right) \quad (17)$$

and reward (outcome) likelihoods:

$$y|\mathbf{c}, \mathbf{a}, \psi \sim \mathcal{N}(f(\mathbf{c}, \mathbf{a}, \psi), 0.1) \quad (18)$$

$$f(\mathbf{c}, \mathbf{a}, \psi) = -\mathbf{c}^2 + \beta(\mathbf{a}, \psi)\mathbf{c} + \gamma(\mathbf{a}, \psi) \quad (19)$$

$$\gamma = (\psi_{\mathbf{a},1} + \psi_{\mathbf{a},2} + 18)/2 \quad (20)$$

$$\beta = (\psi_{\mathbf{a},2} - \gamma + 9)/3 \quad (21)$$

Intuition about the parameterisation: The first component of each ψ_i defines the mean reward at context $\mathbf{c} = -3$, while the second one defines the mean reward at context $\mathbf{c} = 3$. The reward is then the quadratic equation that passes through those points and leading coefficient equal to -1 .

Experimental and evaluation contexts We use experimental and evaluation contexts of the same sizes. The experimental context, \mathbf{c} is an equally spaced grid of size 10 between -3 and -1 . We set the evaluation context $\mathbf{c}^* = -\mathbf{c}$. Figure 3 in the main text visually illustrates this: the x -axis of the points in each plot are the experimental contests, while the dashed gray lines are the evaluation contexts.

Further results Table 2 shows all the evaluation metrics for the discrete treatment example from the main text. Our method achieves substantially higher EIG and lower MSE of estimating the max-rewards. On all other metrics all methods perform similarly. This is to be expected, since Treatments 1 and 2 have exactly the same means and due to the way the model was parameterised (by the value of the quadratic at contexts 3 and -3), the probability of the optimal treatment being 1 or 2 is exactly 50% (the hit rate all baselines achieve). Note that UCB₁ and UCB₂ achieve statistically identical results, which is expected given they select the same designs.

Training stability We perform our method with the same hyperparameters but different training seeds and report the mean and standard error in Table 3.

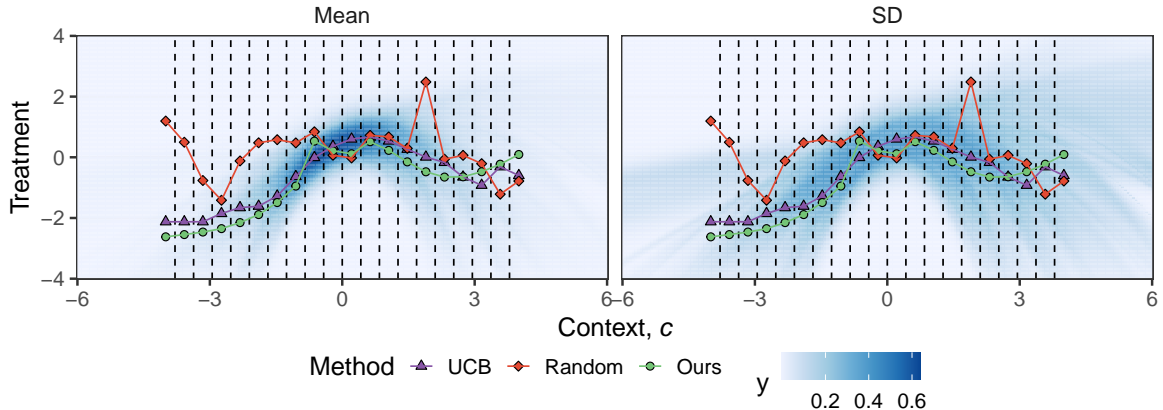


Figure 5: Continuous treatment example: 20D design to learn about 19 evaluation contexts (dashed gray lines). The random baseline samples treatments from $\mathcal{N}(0, 1)$.

Method	EIG estimate	MSE(\mathbf{m}^*)	MSE(ψ)	MSE(\mathbf{A})	Regret
Random _{0.2}	5.548 \pm 0.044	0.0037 \pm 0.0002	0.0101 \pm 0.0008	0.451 \pm 0.033	0.083 \pm 0.004
Random _{1.0}	5.654 \pm 0.128	0.0031 \pm 0.0004	0.0065 \pm 0.0008	0.343 \pm 0.044	0.069 \pm 0.006
Random _{2.0}	5.118 \pm 0.163	0.0045 \pm 0.0003	0.0096 \pm 0.0010	0.498 \pm 0.032	0.086 \pm 0.004
UCB _{0.0}	5.768 \pm 0.002	0.0066 \pm 0.0002	0.0148 \pm 0.0002	0.729 \pm 0.022	0.082 \pm 0.001
UCB _{1.0}	5.892 \pm 0.006	0.0031 \pm 0.0001	0.0097 \pm 0.0002	0.354 \pm 0.013	0.068 \pm 0.001
UCB _{2.0}	5.797 \pm 0.004	0.0030 \pm 0.0001	0.0090 \pm 0.0001	0.343 \pm 0.011	0.071 \pm 0.001
Ours	6.538 \pm 0.008	0.0013 \pm 0.0001	0.0038 \pm 0.0001	0.131 \pm 0.006	0.042 \pm 0.0001

Table 4: Continuous treatments example: 40D design training stability. Mean and standard error are reported across 6 different training seeds.

C.5 Continuous treatment example

Model For the continuous treatment example we use the following model:

$$\text{Prior: } \psi = (\psi_0, \psi_1, \psi_2, \psi_3), \quad \psi_i \sim \text{Uniform}[0.1, 1.1] \text{ iid} \quad (22)$$

$$\text{Likelihood: } y|\mathbf{c}, \mathbf{a}, \psi \sim \mathcal{N}(f(\psi, \mathbf{a}, \mathbf{c}), \sigma^2), \quad (23)$$

where

$$f(\psi, \mathbf{a}, \mathbf{c}) = \exp\left(-\frac{(a - g(\psi, \mathbf{c}))^2}{h(\psi, \mathbf{c})}\right) \quad g(\psi, \mathbf{c}) = \psi_0 + \psi_1 \mathbf{c} + \psi_2 \mathbf{c}^2 \quad h(\psi, \mathbf{c}) = \psi_3 \quad (24)$$

Experimental and evaluation contexts The experimental context, \mathbf{c} is an equally spaced grid of size $D = 40$ (or 20 or 60 in Further Results below) between -3.5 and 3.5 . The evaluation context \mathbf{c}^* is of size $D^* = D - 1$ and consists of the midpoints of the experimental context (see Figure 5 for an illustration).

Method	EIG estimate	MSE(\mathbf{m}^*)	MSE(ψ)	MSE(\mathbf{A})	Regret
Random _{0,2}	4.262 \pm 0.004	0.0086 \pm 0.0003	0.0176 \pm 0.0004	1.046 \pm 0.041	0.120 \pm 0.002
Random _{1,0}	4.264 \pm 0.004	0.0068 \pm 0.0003	0.0158 \pm 0.0004	0.799 \pm 0.033	0.114 \pm 0.002
Random _{2,0}	4.116 \pm 0.003	0.0083 \pm 0.0003	0.0198 \pm 0.0005	1.002 \pm 0.044	0.127 \pm 0.003
UCB _{0,0}	5.093 \pm 0.004	0.0074 \pm 0.0004	0.0186 \pm 0.0006	0.800 \pm 0.047	0.097 \pm 0.002
UCB _{1,0}	5.040 \pm 0.004	0.0072 \pm 0.0004	0.0180 \pm 0.0006	0.764 \pm 0.041	0.097 \pm 0.002
UCB _{2,0}	5.038 \pm 0.004	0.0048 \pm 0.0003	0.0127 \pm 0.0004	0.573 \pm 0.033	0.086 \pm 0.002
Ours	5.642 \pm 0.003	0.0034 \pm 0.0002	0.0073 \pm 0.0003	0.065 \pm 0.002	0.344 \pm 0.027

Table 5: Continuous treatment example: 20D design to learn about 19 evaluation contexts.

Method	EIG estimate	MSE(\mathbf{m}^*)	MSE(ψ)	MSE(\mathbf{A})	Regret
Random _{0,2}	6.033 \pm 0.003	0.0026 \pm 0.0001	0.0068 \pm 0.0002	0.307 \pm 0.019	0.068 \pm 0.002
Random _{1,0}	5.877 \pm 0.004	0.0025 \pm 0.0002	0.0058 \pm 0.0002	0.310 \pm 0.023	0.064 \pm 0.002
Random _{2,0}	6.153 \pm 0.003	0.0022 \pm 0.0002	0.0046 \pm 0.0002	0.226 \pm 0.019	0.055 \pm 0.002
UCB _{0,0}	6.106 \pm 0.003	0.0056 \pm 0.0004	0.0134 \pm 0.0006	0.586 \pm 0.045	0.077 \pm 0.002
UCB _{1,0}	6.200 \pm 0.003	0.0027 \pm 0.0003	0.0086 \pm 0.0003	0.305 \pm 0.028	0.063 \pm 0.002
UCB _{2,0}	6.234 \pm 0.003	0.0024 \pm 0.0002	0.0069 \pm 0.0002	0.252 \pm 0.024	0.064 \pm 0.002
Ours	6.932 \pm 0.003	0.0007 \pm 0.0001	0.0027 \pm 0.0001	0.069 \pm 0.009	0.033 \pm 0.001

Table 6: Continuous treatment example: 60D design to learn about 59 evaluation contexts.

Baselines Since we have a continuous treatment, for the random baseline we consider sampling designs at random from $\mathcal{N}(0, 0.2)$, $\mathcal{N}(0, 1)$ or $\mathcal{N}(0, 2)$, which we denote by Random_{0,2}, Random₁ and Random₂, respectively.

Training stability We perform our method with the same hyperparameters but different training seeds and report the mean and standard error in Table 4.

Further results We report the results of the same experiment, but with a smaller and larger batch sizes of experimental and evaluation contexts. Table 5 shows results for an experimental batch size of 20 contexts to learn about 19 evaluation contexts, while Figure 5 visually illustrates the model and the designs. Finally, Table 6 shows results for an experimental batch size of 60 contexts to learn about 59 evaluation contexts.