Bayesian Experimental Design for Implicit Models by Mutual Information Neural Estimation

Steven Kleinegesse [1]   **Michael U. Gutmann** [1]

Abstract

Implicit stochastic models, where the data-generation distribution is intractable but sampling is possible, are ubiquitous in the natural sciences. The models typically have free parameters that need to be inferred from data collected in scientific experiments. A fundamental question is how to design the experiments so that the collected data are most useful. The field of Bayesian experimental design advocates that, ideally, we should choose designs that maximise the mutual information (MI) between the data and the parameters. For implicit models, however, this approach is severely hampered by the high computational cost of computing posteriors and maximising MI, in particular when we have more than a handful of design variables to optimise. In this paper, we propose a new approach to Bayesian experimental design for implicit models that leverages recent advances in neural MI estimation to deal with these issues. We show that training a neural network to maximise a lower bound on MI allows us to jointly determine the optimal design and the posterior. Simulation studies illustrate that this gracefully extends Bayesian experimental design for implicit models to higher design dimensions.

1. Introduction

Many processes in nature can be described by a parametric statistical model from which we can simulate data. When the corresponding data-generating distribution is intractable we refer to it as an implicit model. These models are abundant in science and engineering, having been used, for instance, in high-energy physics (Agostinelli et al., 2003), cosmology (M. Schafer & Freeman, 2012), epidemi-

ology (Corander et al., 2017), cell biology (Vo et al., 2015) and pharmacokinetics (Donnet & Samson, 2013). Usually, we wish to infer the free parameters $\boldsymbol{\theta}$ of the implicit model in order to understand the underlying natural process better or to predict some future events. Since the likelihood function is intractable for implicit models, we have to revert to likelihood-free inference methods such as approximate Bayesian computation (for recent reviews see e.g. Lintusaari et al., 2017; Sisson et al., 2018).

While considerable research effort has focused on developing efficient likelihood-free inference methods (e.g. Papamakarios et al., 2019; Chen & Gutmann, 2019; Gutmann & Corander, 2016; Papamakarios & Murray, 2016; Ong et al., 2018), the quality of the estimated parameters $\boldsymbol{\theta}$ ultimately depends on the quality of the data $\mathbf{y}$ that are available for inference in the first place. We here consider the scenario where we have control over experimental designs $\mathbf{d}$ that affect the data collection process. For example, these might be the spatial location or time at which we take measurements, or they might be the stimulus that is used to perturb the natural system. Our overall goal is then to find experimental designs $\mathbf{d}$ that yield the most information about the model parameters.

In Bayesian experimental design (BED), we construct and optimise a utility function $U(\mathbf{d})$ that indicates the value of a design $\mathbf{d}$. A choice for the utility function that is firmly rooted in information theory is the mutual information (MI) $I(\boldsymbol{\theta}, \mathbf{y}; \mathbf{d})$ between model parameters and data,

$$I(\boldsymbol{\theta}, \mathbf{y}; \mathbf{d}) = \mathbb{E}_{p(\boldsymbol{\theta}, \mathbf{y} \mid \mathbf{d})} \left[\log \frac{p(\boldsymbol{\theta} \mid \mathbf{y}, \mathbf{d})}{p(\boldsymbol{\theta})} \right], \quad (1)$$

where $p(\boldsymbol{\theta} \mid \mathbf{y}, \mathbf{d})$ is the posterior distribution of $\boldsymbol{\theta}$ given data $\mathbf{y}$ that were obtained with design $\mathbf{d}$, and $p(\boldsymbol{\theta})$ is the prior belief of $\boldsymbol{\theta}$.[1] MI describes the expected uncertainty reduction in the model parameters $\boldsymbol{\theta}$ when collecting data with experimental design $\mathbf{d}$, which makes it an effective utility function for BED. Since MI depends on the full posterior $p(\boldsymbol{\theta} \mid \mathbf{y}, \mathbf{d})$, it is sensitive to deviations from the prior, e.g. nonlinear correlations or multi-modality, that other utility functions do not capture (Ryan et al., 2016).

*Equal contribution [1]School of Informatics, University of Edinburgh, Edinburgh, UK. Correspondence to: Steven Kleinegesse <steven.kleinegesse@ed.ac.uk>.

---

[1]It is typically assumed that $p(\boldsymbol{\theta})$ is not affected by $\mathbf{d}$.

Although a quantity with highly desirable properties, MI is notoriously difficult to compute and maximise. The key problems in the context of implicit models are 1) that the posteriors are hard to estimate and that obtaining samples from them is expensive as well; 2) that the functional relationship between the designs and the MI is unknown and that (approximate) gradients are generally not available either.

The first problem made the community turn to approximate Bayesian computation and less powerful utility functions, e.g. based on the posterior variance (Drovandi & Pettitt, 2013), or approximations with kernel density estimation (Price et al., 2016). More recently, Kleinegesse & Gutmann (2019) have shown that likelihood-free inference by density ratio estimation (Thomas et al., 2016) can be used to estimate the posterior and the mutual information at the same time, which alleviates the first problem to some extent. The second problem has been addressed by using gradient-free optimisation techniques such as grid-search, sampling (Müller, 1999), evolutionary algorithms (Price et al., 2018), Gaussian-process surrogate modelling (Overstall & McGree, 2018) and Bayesian optimisation (Kleinegesse & Gutmann, 2019). However, these approaches generally do not scale well with the dimensionality of the design variables $\mathbf{d}$ (e.g. Spall, 2005).

In BED, our primary goal is to find the designs $\mathbf{d}$ that maximise the MI and not estimating the MI to a high accuracy. Rather than spending resources on estimating MI accurately for non-optimal designs, a potentially more cost-effective approach is thus to relax the problem and to determine the designs that maximise a lower bound of the MI instead, while tightening the bound at the same time. In our paper, we take this approach and show that for a large class of implicit models, the lower bound can be tightened and maximised by gradient ascent, which addresses the aforementioned scalability issues. Our approach leverages the Mutual Information Neural Estimation (MINE) method of Belghazi et al. (2018) to perform BED — we thus call it MINEBED.[2] We show later on that in addition to gradient-based experimental design, MINEBED also provides us with an approximation of the posterior $p(\boldsymbol{\theta} \mid \mathbf{y}, \mathbf{d})$ so that no separate and oftentimes expensive likelihood-free inference step is needed.

Related Work Foster et al. (2019a) have recently considered the use of MI lower bounds for experimental design. Their approach is based on variational approximations to the posterior and likelihood. While the authors note that this introduces a bias, in practice, this may actually be acceptable. However, the optimal designs were determined by Bayesian optimisation and, like in the aforementioned previous work,

---

[2]The research code for this paper is available at: https://github.com/stevenkleinegesse/minebed

this approach may thus also suffer from scalability problems. The authors rectify this in a follow-up paper (Foster et al., 2019b) but their experiments focus on the explicit setting where likelihood functions are tractable or approximations are available, unlike our paper. In theory, their method is also applicable in the likelihood-free setting.

2. The MINEBED Method

We here show how to perform experimental design for implicit models by maximising a lower bound of the mutual information (MI) between model parameters $\boldsymbol{\theta}$ and data $\mathbf{y}$. The main properties of our approach are that for a large and interesting class of implicit models, the designs can be found by gradient ascent and that the method simultaneously yields an estimate of the posterior.

We start with the MINE-f MI lower bound of Belghazi et al. (2018), which we shall call $\widehat{I}(\mathbf{d}, \boldsymbol{\psi})$,

$$
\begin{aligned}
\widehat{I}(\mathbf{d}, \boldsymbol{\psi}) = \mathbb{E}_{p(\boldsymbol{\theta}, \mathbf{y} \mid \mathbf{d})} \left[T_{\boldsymbol{\psi}}(\boldsymbol{\theta}, \mathbf{y}) \right] \\
- e^{-1} \mathbb{E}_{p(\boldsymbol{\theta}) p(\mathbf{y} \mid \mathbf{d})} \left[e^{T_{\boldsymbol{\psi}}(\boldsymbol{\theta}, \mathbf{y})} \right],
\end{aligned}
\tag{2}
$$

where $T_{\boldsymbol{\psi}}(\boldsymbol{\theta}, \mathbf{y})$ is a neural network parametrised by $\boldsymbol{\psi}$, with $\boldsymbol{\theta}$ and $\mathbf{y}$ as input. This is also the lower bound of Nguyen et al. (2010) and the f-GAN KL of Nowozin et al. (2016). We note that while we focus on the above lower bound, our approach could also be applied to other MI lower bounds of e.g. Poole et al. (2019), Foster et al. (2019a) and Foster et al. (2019b). Moreover, comparing the performance of different MI lower bounds is not the aim of this paper and we do not claim that the particular bound in (2) is superior to others.

Importantly, Belghazi et al. (2018) showed that the bound in (2) can be tightened by maximising it with respect to the neural network parameters $\boldsymbol{\psi}$ by gradient ascent and that for flexible enough neural networks we obtain $I(\boldsymbol{\theta}, \mathbf{y}; \mathbf{d}) = \max_{\boldsymbol{\psi}} \widehat{I}(\mathbf{d}, \boldsymbol{\psi})$. Our experimental design problem can thus be formulated as

$$
\mathbf{d}^* = \arg \max_{\mathbf{d}} \max_{\boldsymbol{\psi}} \left\{ \widehat{I}(\mathbf{d}, \boldsymbol{\psi}) \right\}.
\tag{3}
$$

The main difficulty is that for implicit models, the expectations in the definition of $\widehat{I}(\mathbf{d}, \boldsymbol{\psi})$ generally depend in some complex manner on the design variables $\mathbf{d}$, complicating gradient-based optimisation.

2.1. Experimental Design with Gradients

We are interested in computing the derivative $\nabla_{\mathbf{d}} \widehat{I}(\mathbf{d}, \boldsymbol{\psi})$,

$$
\begin{aligned}
\nabla_{\mathbf{d}} \widehat{I}(\mathbf{d}, \boldsymbol{\psi}) = \nabla_{\mathbf{d}} \mathbb{E}_{p(\boldsymbol{\theta}, \mathbf{y} \mid \mathbf{d})} \left[T_{\boldsymbol{\psi}}(\boldsymbol{\theta}, \mathbf{y}) \right] \\
- \nabla_{\mathbf{d}} e^{-1} \mathbb{E}_{p(\boldsymbol{\theta}) p(\mathbf{y} \mid \mathbf{d})} \left[e^{T_{\boldsymbol{\psi}}(\boldsymbol{\theta}, \mathbf{y})} \right].
\end{aligned}
\tag{4}
$$

For implicit models we do not have access to the gradients of the joint distribution $p(\boldsymbol{\theta}, \mathbf{y} \mid \mathbf{d})$ and marginal distribution

$p(\mathbf{y}|\mathbf{d})$. Thus, we cannot simply pull $\nabla_\mathbf{d}$ into the integrals defining the expectations and compute the derivatives. For the same reason we cannot use score-function estimators (Mohamed et al., 2019), as they rely on the analytic derivative of the log densities.

The alternative pathwise gradient estimators (see Mohamed et al., 2019, for a review), however, are well suited for implicit models. They require that we can sample from the data-generating distribution $p(\mathbf{y}|\boldsymbol{\theta},\mathbf{d})$ by sampling from a base distribution $p(\boldsymbol{\epsilon})$ and then transforming the samples through a (nonlinear) deterministic function $\mathbf{h}(\boldsymbol{\epsilon};\boldsymbol{\theta},\mathbf{d})$ called the sampling path (Mohamed et al., 2019), i.e.

$$\mathbf{y} \sim p(\mathbf{y}|\boldsymbol{\theta},\mathbf{d}) \iff \mathbf{y} = \mathbf{h}(\boldsymbol{\epsilon};\boldsymbol{\theta},\mathbf{d}), \quad \boldsymbol{\epsilon} \sim p(\boldsymbol{\epsilon}). \quad (5)$$

For implicit models, this requirement is automatically satisfied because the models are specified by such transformations in the first place.

We first factorise the joint $p(\boldsymbol{\theta},\mathbf{y}|\mathbf{d}) = p(\mathbf{y}|\boldsymbol{\theta},\mathbf{d})p(\boldsymbol{\theta})$, where we again assume that $p(\boldsymbol{\theta})$ is not affected by $\mathbf{d}$. Specifying $p(\mathbf{y}|\boldsymbol{\theta},\mathbf{d})$ implicitly in terms of the sampling path $\mathbf{h}(\boldsymbol{\epsilon};\boldsymbol{\theta},\mathbf{d})$ then allows us to invoke the law of the unconscious statistician (e.g. Grimmett & Stirzaker, 2001) and, for instance, rewrite the first expectation in (2) as

$$\mathbb{E}_{p(\boldsymbol{\theta},\mathbf{y}|\mathbf{d})}\left[T_{\boldsymbol{\psi}}(\boldsymbol{\theta},\mathbf{y})\right] = \mathbb{E}_{p(\boldsymbol{\theta})p(\boldsymbol{\epsilon})}\left[T_{\boldsymbol{\psi}}(\boldsymbol{\theta},\mathbf{h}(\boldsymbol{\epsilon};\boldsymbol{\theta},\mathbf{d}))\right], \quad (6)$$

Since neither $p(\boldsymbol{\epsilon})$ nor $p(\boldsymbol{\theta})$ directly depend on $\mathbf{d}$, taking the derivative of the above expectation, which proved problematic before, yields

$$\nabla_\mathbf{d}\mathbb{E}_{p(\boldsymbol{\theta},\mathbf{y}|\mathbf{d})}\left[T_{\boldsymbol{\psi}}(\boldsymbol{\theta},\mathbf{y})\right] = \mathbb{E}_{p(\boldsymbol{\theta})p(\boldsymbol{\epsilon})}\left[\nabla_\mathbf{d}T_{\boldsymbol{\psi}}(\boldsymbol{\theta},\mathbf{h}(\boldsymbol{\epsilon};\boldsymbol{\theta},\mathbf{d}))\right]. \quad (7)$$

We can compute the derivative of the neural network output with respect to designs via the chain rule, i.e.

$$\nabla_\mathbf{d}T_{\boldsymbol{\psi}}(\boldsymbol{\theta},\mathbf{h}(\boldsymbol{\epsilon};\boldsymbol{\theta},\mathbf{d})) = \boldsymbol{J}_{\boldsymbol{y}}^\top\nabla_\mathbf{y}T_{\boldsymbol{\psi}}(\boldsymbol{\theta},\mathbf{y})|_{\mathbf{y}=\mathbf{h}(\boldsymbol{\epsilon};\boldsymbol{\theta},\mathbf{d})}, \quad (8)$$

where $\boldsymbol{J}_{\boldsymbol{y}}$ is the Jacobian matrix containing the derivatives of the vector-valued $\mathbf{y} = \mathbf{h}(\boldsymbol{\epsilon};\boldsymbol{\theta},\mathbf{d})$ with respect to $\mathbf{d}$[3] and $\nabla_\mathbf{y}T_{\boldsymbol{\psi}}(\boldsymbol{\theta},\mathbf{y})$ is the gradient of the neural network $T_{\boldsymbol{\psi}}(\boldsymbol{\theta},\mathbf{y})$ with respect to its inputs $\mathbf{y}$, which is provided by most machine-learning libraries via auto-differentiation. The computation of the Jacobian $\boldsymbol{J}_{\boldsymbol{y}}$, or the Jacobian-gradient product, is the main technical obstacle, and below we discuss in which cases we can compute it and what to do when we cannot.

For the second expectation in (4), which is defined over the marginal distribution, we have to apply pathwise gradient estimators in a similar manner to above. However, we only know the sampling path of the data-generating distribution, and not the sampling path of the marginal random variable

---

[3]We use the definition that $\boldsymbol{J}_{\boldsymbol{y}\,ij} = \partial y_i/\partial d_j$.

$\mathbf{y} \sim p(\mathbf{y} \mid \mathbf{d})$. Thus, we need to re-write the marginal distribution as $p(\mathbf{y} \mid \mathbf{d}) = \mathbb{E}_{p(\boldsymbol{\theta}')}\left[p(\mathbf{y} \mid \boldsymbol{\theta}',\mathbf{d})\right]$, where $\boldsymbol{\theta}'$ follows the same distribution as $\boldsymbol{\theta}$, which then allows us to follow the procedure as described above.

This gives us the following pathwise gradient estimator of the MI lower bound

$$\nabla_\mathbf{d}\widehat{I}(\mathbf{d},\boldsymbol{\psi}) = \mathbb{E}_{p(\boldsymbol{\theta})p(\boldsymbol{\epsilon})}\left[\nabla_\mathbf{d}T_{\boldsymbol{\psi}}(\boldsymbol{\theta},\mathbf{y})\right]$$
$$- e^{-1}\mathbb{E}_{p(\boldsymbol{\theta})p(\boldsymbol{\epsilon})p(\boldsymbol{\theta}')}\left[e^{T_{\boldsymbol{\psi}}(\boldsymbol{\theta},\mathbf{y}')}\nabla_\mathbf{d}T_{\boldsymbol{\psi}}(\boldsymbol{\theta}',\mathbf{y}')\right], \quad (9)$$

where $\mathbf{y} = \mathbf{h}(\boldsymbol{\epsilon};\boldsymbol{\theta},\mathbf{d})$, $\mathbf{y}' = \mathbf{h}(\boldsymbol{\epsilon};\boldsymbol{\theta}',\mathbf{d})$ and the gradients of $T_{\boldsymbol{\psi}}(\boldsymbol{\theta},\mathbf{y})$ with respect to $\mathbf{d}$ are given in (8). Throughout this work, we approximate all expectations, i.e. in the MI lower bound in (2) and in the corresponding gradients in (9), via a sample-average.

With the gradient in (9) in hand, we can maximise the MI lower bound in (2) with respect to both the neural network parameters $\boldsymbol{\psi}$ and the experimental designs $\mathbf{d}$ jointly by gradient ascent. Gradient-based optimisation is key to scaling experimental design to higher dimensions and the joint maximisation tightens the lower bound automatically as the (locally) optimal designs are determined more accurately.

The above development assumes that we can compute the Jacobian $\boldsymbol{J}_{\boldsymbol{y}}$ or the Jacobian-gradient product in (8). While generally not guaranteed to be computable for implicit models, we do have access to the exact Jacobian in a number of important cases: Interestingly, for several implicit models considered in the literature, the data generating distribution $p(\mathbf{y}|\boldsymbol{\theta},\mathbf{d})$ might be intractable, but $\mathbf{h}(\boldsymbol{\epsilon};\boldsymbol{\theta},\mathbf{d})$ can be expressed analytically in closed form and derivatives with respect to $\mathbf{d}$ are readily available. Examples include the noisy linear model and the pharmacokinetic model considered in this work.

Another broad class of implicit models where the Jacobian is exactly computable is the case where the generative process separates into a black-box latent process that does not depend on $\mathbf{d}$ and a differentiable observation process that is governed or affected by $\mathbf{d}$. In many cases, the latent process is described by stochastic differential equations and the design variables $\mathbf{d}$ are the measurement times. Whenever we have access to the differential equations, which we most often have when solving them numerically, the Jacobian can be computed exactly too. The pharmacokinetic model can be considered an example of this category as well.

Finally, if the implicit model is written in a differential programming framework, the required Jacobian-gradient product in (8) can be computed via their built-in automatic differentiation functionality. Given the usefulness of automatic differentiation packages, we expect their usage only to increase and hence also the range of models that our method can be applied to.

2.2. Experimental Design without Gradients

We have seen that, for a large class of implicit models, the optimisation problem in (3) can be solved by gradient ascent on the design variables **d**. However, there might be cases of implict models and design problems for which the required gradient can neither be computed exactly nor be reasonably approximated. For example, when the implicit model is described by a nonlinear stochastic differential equation and the experimental design is concerned with interventions and not measurements as assumed above, the gradient cannot be computed exactly.

However, we can fall back to gradient-free methods to solve the optimisation problem in (3), at the expense of reduced scalability. We next describe how Bayesian optimisation (BO) (see Shahriari et al., 2016, for a review) can be used to learn the optimal design **d*** when gradients are not available. We decided to use BO as it smoothes out Monte-Carlo error and has been shown to perform well in practice.

In brief, BO works by defining a probabilistic surrogate model of the function to be optimised. For the surrogate model we will be using a Gaussian process (GP) (Rasmussen & Williams, 2005) with a Matérn-5/2 kernel. Based on this surrogate model we have to compute and optimise an acquisition function over the domain, which tells us where to evaluate the function next. We will be using expected improvement (EI), as this has been shown to generally be a cost-effective choice (Shahriari et al., 2016).

Using an initial design $\mathbf{d}_0$, we train the neural network $T_{\boldsymbol{\psi}}(\boldsymbol{\theta}, \mathbf{y})$ by gradient ascent on $\boldsymbol{\psi}$. We then use the obtained parameters $\boldsymbol{\psi}^*$ to compute the lower bound $\widehat{I}(\mathbf{d}_0, \boldsymbol{\psi}^*)$, update our probabilistic surrogate model given the computed value, and use the acquisition function to decide for which design $\mathbf{d}_1$ to evaluate the lower bound next. For this evaluation, we here opt to re-train the neural network. This is because BO often involves significant exploration and thus may lead to design updates that are vastly different from previous designs. This will change the data distributions significantly enough that the neural network has trouble re-training its parameters, a common issue in transfer learning (Pan & Yang, 2010). This procedure is repeated until we have converged to the optimal design **d***.

We note that BO has been used in BED for implicit models before (e.g. Foster et al., 2019a; Kleinegesse & Gutmann, 2019; Kleinegesse et al., 2020), as it is a popular black-box optimisation tool. Furthermore, Foster et al. (2019a) have used BO to optimise a parametrised MI lower bound, based on the Donsker-Varadhan representation of mutual information (Donsker & Varadhan, 1983). We would like to emphasise that we only present the use of BO for MINEBED as a fall-back method in cases where the required Jacobian cannot be computed.

2.3. Estimating the Posterior Distribution

It has been shown that the learned neural network $T_{\boldsymbol{\psi}^*}(\boldsymbol{\theta}, \mathbf{y})$ is related to the posterior and prior distribution when the bound is tight (Nguyen et al., 2010; Belghazi et al., 2018; Poole et al., 2019),

$$p(\boldsymbol{\theta} \mid \mathbf{y}, \mathbf{d}) = e^{T_{\boldsymbol{\psi}^*}(\boldsymbol{\theta}, \mathbf{y}) - 1} p(\boldsymbol{\theta}). \qquad (10)$$

This relationship allows us obtain an estimate of the posterior $p(\boldsymbol{\theta} \mid \mathbf{y}, \mathbf{d}^*)$ without any additional work after we have determined the optimal design $\mathbf{d}^*$ by maximisation of the MI lower bound. Since sampling from the prior $p(\boldsymbol{\theta})$ is often easily possible, the above expression further allows us to cheaply obtain posterior samples $\boldsymbol{\theta}^{(i)} \sim p(\boldsymbol{\theta} \mid \mathbf{y}, \mathbf{d}^*)$ in an amortised way for any data $\mathbf{y}$ by re-weighing samples from the prior and then using categorical sampling. In practice, any MCMC sampling scheme could be used in place as well, since we are able to quickly evaluate the posterior density via (10).

3. Experiments

We here showcase our approach to BED for implicit models. First, to demonstrate scalability, we use gradient ascent to find optimal designs for a linear model with several noise sources and a pharmacokinetic model (Ryan et al., 2014). Then we design experiments to locate a gas leak in a fluid simulation experiment (Asenov et al., 2019) where we do not have sampling path gradients. An important property of mutual information is that it is sensitive to multi-modal posteriors, unlike previous popular approaches. We verify this in the supplementary material using a multi-modal, oscillatory toy model.

3.1. Noisy Linear Model

We first consider the classical model where it is assumed that a response variable y has a linear relationship with an experimental design d, governed by some model parameters $\boldsymbol{\theta} = [\theta_0, \theta_1]^\top$, the offset and the gradient, that we wish to estimate. If we then wish to make D measurements in order to estimate $\boldsymbol{\theta}$, we construct a design vector $\mathbf{d} = [d_1, \ldots, d_D]^\top$ that consists of individual experimental designs. For each d_i we have an independent measurement y_i, which gives us the data vector $\mathbf{y} = [y_1, \ldots, y_D]^\top$. There may also be several noise sources affecting the outcome of measurements, complicating the parameter estimation of such a model. If all of these noise sources happen to be Gaussian, we can write down an analytic likelihood and compute the posterior distribution, as well as MI, exactly.

For this toy model, we choose to include two different noise sources, a Gaussian noise source $\mathcal{N}(\epsilon; 0, 1)$ and a Gamma noise source $\Gamma(\nu; 2, 2)$. The sampling path is given by

$$\mathbf{y} = \theta_0 \mathbf{1} + \theta_1 \mathbf{d} + \boldsymbol{\epsilon} + \boldsymbol{\nu}, \qquad (11)$$
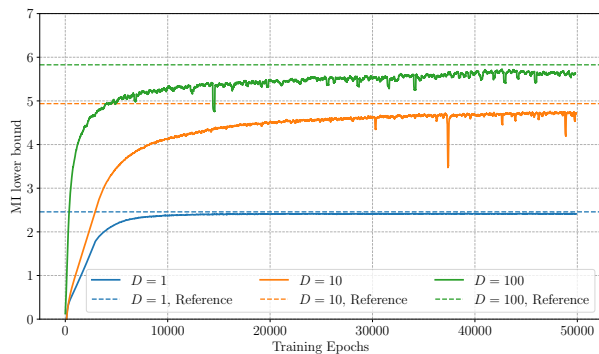
Figure 1. MI lower bound as a function of neural network training epochs for the $D = 1$ dimensional (blue), $D = 10$ dimensional (orange) and $D = 100$ dimensional (green) noisy linear model. Shown are the moving averages with a window size of 100 and the dotted lines are numerical reference MI values computed at d^*.

where $\boldsymbol{\epsilon} = [\epsilon_1, \ldots, \epsilon_D]^\top$ and $\boldsymbol{\nu} = [\nu_1, \ldots, \nu_D]^\top$ consist of i.i.d. samples; $\mathbf{1}$ is a D-dimensional vector of ones. While numerical integration could be used to approximate the posterior and MI, we here consider it an implicit model to test our approach. The Jacobian matrix $\boldsymbol{J_y}$, needed for the computation of the gradient in (8), is

$$\boldsymbol{J_y} = \theta_1 \mathbb{1}, \tag{12}$$

where $\mathbb{1}$ is the identity matrix. Knowing the sampling path and its gradient with respect to designs then allows us to compute an estimate of the gradients $\nabla_{\mathbf{d}} \widehat{I}(\mathbf{d}, \boldsymbol{\psi})$ in (9) and thus find the optimal designs $\mathbf{d}^* = [d_1^*, \ldots, d_D^*]^\top$.

At first, we wish to take 1 measurement only, i.e. we set $D = 1$. We randomly initialise a design $d \in [-10, 10]$ and sample 30,000 parameters from the prior distribution $p(\boldsymbol{\theta}) = \mathcal{N}(\boldsymbol{\theta}; \mathbf{0}, 3^2 \mathbb{I})$. For every training epoch we generate 30,000 new data samples because we continuously update the design variable and hence the data-generating distribution changes. For the neural network $T_{\boldsymbol{\psi}}(\boldsymbol{\theta}, \mathbf{y})$ we use one layer of 100 hidden units, with a ReLU activation layer after the input layer. We update the neural network parameters and designs using the Adam optimiser, with learning rates of $l_{\boldsymbol{\psi}} = 10^{-4}$ and $l_{\mathbf{d}} = 10^{-2}$, respectively. Because the design domain is bounded, we ignore any design updates that would go beyond the specified domain of $d \in [-10, 10]$.

The blue curve in Figure 1 shows the MI lower bound as a function of training epochs for one-dimensional designs. It can be seen that the lower bound $\widehat{I}(\mathbf{d}, \boldsymbol{\psi})$ converges after roughly 15,000 epochs. The final MI lower bound is around 2.5, close to a reference MI computed numerically.[4] The optimal design found is $d^* = -10$, i.e. at the boundary; we

[4]See the supplementary material for how this is computed.

show the convergence of the design as a function of training epochs in the supplementary material. This optimal design value is intuitive, as there is a larger signal-to-noise ratio at large designs for the linear model.

Having found an optimal design d^*, we can perform a real-world experiment to obtain data y^* and use these to estimate the model parameters $\boldsymbol{\theta}$. We here use $\boldsymbol{\theta}_{\text{true}} = [2, 5]^\top$ to generate y^*. Using (10) we then compute the posterior density, which we show to the left in Figure 2. The posterior density is relatively broad for both parameters, with the mode being quite far from the ground truth. By means of (10) and categorical sampling we obtain posterior samples that we use to estimate the parameters, resulting in $\widehat{\theta}_0 = -0.736 \pm 3.074$ and $\widehat{\theta}_1 = 5.049 \pm 0.401$, where the error indicates a 68% confidence interval. Because the model parameters are two-dimensional, accurately estimating both with one observation is naturally impossible.

We now assume that we have enough budget to take $D = 10$ measurements of the response variable. It is generally worthwhile to do hyper-parameter optimisation to choose an appropriate neural network architecture for $T_{\boldsymbol{\psi}}(\boldsymbol{\theta}, \mathbf{y})$. By means of grid search (see the supplementary material), we found a neural network with one hidden layer and 150 hidden units to be working best. As before, we used the Adam optimiser with learning rates of $l_{\boldsymbol{\psi}} = 10^{-4}$ and $l_{\mathbf{d}} = 10^{-2}$. While finding an optimal architecture is usually difficult, this problem is mitigated by the fact that our method only requires synthetic data that we can generate ourselves. The resulting MI lower bound for $D = 10$ is shown in Figure 1 (orange curve), including a reference MI computed at $\mathbf{d}^*$. The final MI lower bound is around 4.5 with a bias to the reference MI value that is slightly larger than for $D = 1$. We show the optimal designs $\mathbf{d}^* = [d_1^*, \ldots, d_{10}^*]^\top$ in the bottom of Figure 2. The design dimensions cluster in three distinct regions: Two clusters are at the two boundaries which is ideal to learn the slope θ_1, and one cluster is in the middle where the response is most strongly affected by the offset θ_0. The supplementary material shows the convergence of each design dimension as a function of training epochs. As done for the one-dimensional model, we would then perform a real-world experiment at $\mathbf{d}^*$ to observe the response $\mathbf{y}^*$, which we generate using $\boldsymbol{\theta}_{\text{true}} = [2, 5]^\top$, and use this to estimate the parameters $\boldsymbol{\theta}$. We show the resulting posterior density, again computed by using (10), to the right in Figure 2. Due to having more measurements, the posterior distribution is narrower and more accurate for $D = 10$ than for $D = 1$. By using categorical sampling we obtain the estimates $\widehat{\theta}_0 = 1.441 \pm 0.826$ and $\widehat{\theta}_1 = 4.949 \pm 0.108$.

In order to test the scalability and robustness of our method, we further consider the noisy linear model for $D = 100$ dimensions. We show the corresponding mutual information lower bound in Figure 1 as the green curve. The final
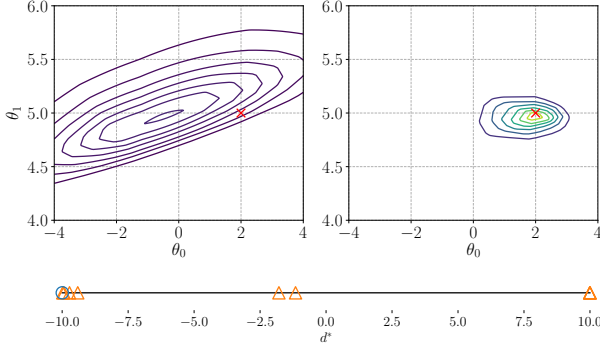
Figure 2. Posterior density for the one-dimensional (left) and 10-dimensional (right) noisy linear model. Real-world observations $\mathbf{y}^*$ are generated at $\mathbf{d}^*$ using $\boldsymbol{\theta}_{\text{true}} = [2, 5]^\top$ (shown as a red cross). The bottom line shows the optimal designs for $D = 1$ (blue circle) and $D = 10$ (orange triangles).

mutual information lower bound for $D = 100$ is higher than for the $D = 10$ or $D = 1$ models, which should be expected due to having more measurements. Even in such high-dimensions, we find that we can converge to a lower bound that only has a small bias with respect to a reference mutual information value (shown as a dashed, green line). Interestingly, we find that extrapolating the design strategy from $D = 10$ to $D = 100$ yields sub-optimal results. Instead, one should focus measurements around zero and avoid the domain boundaries. See the supplementary materials for the results, the description of the setup, and further analysis and figures.

3.2. Pharmacokinetic Model

In this experiment we are concerned with finding the optimal blood sampling time during a pharmacokinetic (PK) study. These types of studies involve the administration of a drug to a patient or group of patients and then measuring the drug concentration at certain times to study the underlying kinetics. We shall use the compartmental model introduced by Ryan et al. (2014) to simulate the drug concentration at a particular time. Their model is governed by three parameters that we wish to estimate: the absorption rate k_a, the elimination rate k_e, and the volume of distribution V which indicates how the drug will distribute in the body. We will assume that we have a group of D patients but for each of the patients, we can take only one blood sample at time t.

The model consist of a latent variable $z(t)$ for the drug concentration at design time t,

$$z(t) = \frac{D_V}{V} \frac{k_a}{k_a - k_e} \left[e^{-k_e t} - e^{-k_a t} \right] (1 + \epsilon), \quad (13)$$

where $D_V = 400$ is a constant and $\epsilon \sim \mathcal{N}(0, 0.01)$ is mul-

tiplicative noise modelling heteroscedasticity that is often observed in drug concentration data. We further have an additive observation model that yields the measured drug concentration y at design time t, i.e.

$$y(t) = z(t) + \nu, \quad (14)$$

where $\nu \sim \mathcal{N}(0, 0.1)$. We note that (13) is a simplified model in that it assumes that the parameters $\boldsymbol{\theta} = [k_a, k_e, V]^\top$ are the same for the group of patients considered, and that it does not model the time-course of the concentration for each patient individually but only the marginal distribution at design time t. If we wanted to take several measurements per patient instead of one, we had to model the concentration correlation across time and e.g. use models described by stochastic differential equations (see e.g. Donnet & Samson, 2013).

We follow Ryan et al. (2014) and assume the following prior for $\boldsymbol{\theta}$

$$\log \boldsymbol{\theta} \sim \mathcal{N} \left(\begin{bmatrix} \log 1 \\ \log 0.1 \\ \log 20 \end{bmatrix}, \begin{bmatrix} 0.05 & 0 & 0 \\ 0 & 0.05 & 0 \\ 0 & 0 & 0.05 \end{bmatrix} \right), \quad (15)$$

with the additional constraint that $k_a > k_e$, enforced via rejection sampling. The PK model in (14) can be written in terms of the sampling path $y = h(\epsilon, \nu \mid \boldsymbol{\theta}, t)$. The corresponding sampling path gradients are

$$\frac{\partial y(t)}{\partial t} = \frac{D_V}{V} \frac{k_a}{k_a - k_e} \left[-k_e e^{-k_e t} + k_a e^{-k_a t} \right] (1 + \epsilon) \quad (16)$$

Knowing the sampling path derivative in (16) allows us to find optimal blood sampling times by gradient ascent on the MI lower bound. Similar to the linear model, we here assume that we wish to take D measurements. In practice, this means that we have D patients and take one blood sample per patient. Our design vector then becomes $\mathbf{d} = [t_1, \ldots, t_D]^\top$, with corresponding observations $\mathbf{y} = [y_1, \ldots, y_D]^\top$.

We start with the simpler setting of only performing one experiment, i.e. we set $D = 1$. At first, we randomly initialise a design time in $t \in [0, 24]$ hours, the design domain for this experiment. We then draw 30,000 model parameter samples $\boldsymbol{\theta}^{(i)}$ from the prior distribution in (15). At every epoch we simulate corresponding data samples $y^{(i)}$ using the current design t and the sampling path in (14). For $T_\psi(\boldsymbol{\theta}, \mathbf{y})$ we use a neural network with one hidden layer of 100 hidden units, as well as a ReLU activation function after the input layer. For the simpler setting of one-dimensional designs, we do not perform hyper-parameter optimisation of the network architecture. We then follow the procedure outlined in Section 2 to train the neural network in order to obtain an optimal design time. We use the Adam optimiser to learn
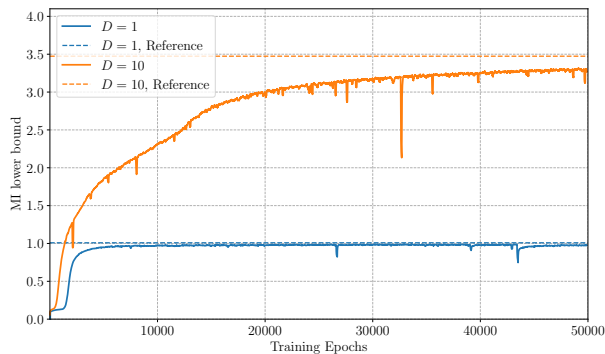
Figure 3. MI lower bound as a function of neural network training epochs for the one-dimensional (blue) and 10-dimensional (orange) PK model. Shown are the moving averages with a window size of 100 and the dotted lines are numerical reference MI values computed at t^*.
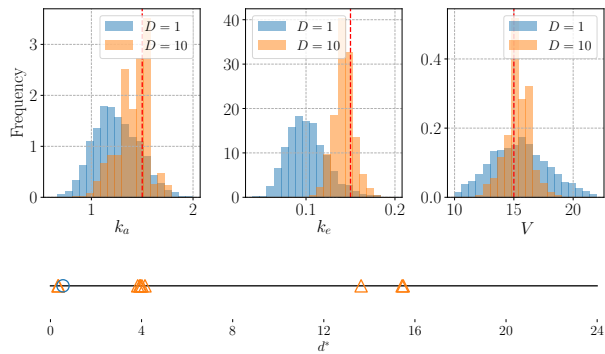
Figure 4. Posterior marginal distributions of the model parameters for the one-dimensional (blue) and 10-dimensional (orange) PK model, with red-dotted lines showing the true model parameters. The bottom line shows the optimal designs for $D = 1$ (blue circle) and $D = 10$ (orange triangles).

ψ and t with learning rates of $l_\psi = 10^{-3}$ and $l_t = 10^{-2}$, respectively.

In Figure 3 we show the MI lower bound as a function of training epochs for the simple one-dimensional PK model (blue curve). The MI lower bound converges to a value of about 1 and has a negligibly small bias compared to a reference MI value at the optimal design $t^* = 0.551$. In the supplementary material, we give details on the computation of the reference value and illustrate the convergence of the design as a function of training epochs.

Using the optimal design t^* we would then do a real-world experiment to obtain data y^*. Here, we generate y^* using $\theta_{\text{true}} = [1.5, 0.15, 15]^\top$. Using the trained neural network $T_{\psi^*}(\boldsymbol{\theta}, \mathbf{y})$, as well as (10), we then compute the posterior density $p(\boldsymbol{\theta} \mid t^*, y^*)$ and use categorical sampling to obtain posterior samples. The marginal posterior distribution of each model parameter is shown in Figure 4. The posterior mean is close to its true value for V, less close for k_a, but relatively far for k_e. This is because with only one measurement at low design times we can only quite accurately determine the ratio of k_a/V, which can be seen from their joint distribution shown in the supplementary material.

We now consider a more complex setting of performing $D = 10$ measurements. As before, appropriate choices of neural network architectures are not obvious. Thus, we test one-layer neural networks with different number of hidden units, as well as different learning rate schedules. We found that the best-performing neural network had a hidden layer of size 300 and a multiplicative learning rate schedule with a multiplying factor of 0.8. In practice, this means that we multiply the initial learning rates $l_\psi = 10^{-3}$ and $l_{\mathbf{d}} = 10^{-2}$ by 0.8 every 5,000 epochs. We again note that further

hyper-parameter search would be possible. Similar to before we use 30,000 prior samples, and corresponding new data generated at each epoch, to train the neural network. We show the resulting lower bound $\widehat{I}(\mathbf{d}, \psi)$ as a function of training epochs in Figure 3 (orange curve). The neural network for $D = 10$ takes longer to converge than for $D = 1$ and ends up at a higher MI lower bound; this is intuitive as more data naturally yields more information about the model parameters. The bias of the final MI lower bound to a reference MI value is larger for $D = 10$ than for $D = 1$ due to the increased dimensions but it is still reasonably small. The optimal designs, shown in the bottom of Figure 4, form three clusters similar to what has been observed by Ryan et al. (2014). The early cluster allows us to determine the ratio k_a/V as before, the late cluster reduces the effect of k_a, while the middle cluster, which is needed to identify the third parameter, corresponds to where the mean response tends to be highest (Ryan et al., 2014). The convergence of the designs are shown in the supplementary material. Similarly to before, we then take a real-world measurement $\mathbf{y}^*$ at $\mathbf{d}^*$, simulated using $\theta_{\text{true}} = [1.5, 0.15, 15]^\top$. Using (10) and categorical sampling we obtain samples from the posterior distribution, with resulting marginal samples for each model parameter shown in Figure 4.[5] As opposed to $D = 1$, where we could only estimate the ratio k_a/V well, for $D = 10$ we can estimate all parameters more precisely and accurately.

3.3. Locating a Gas Leak

In this experiment we aim to locate a gas leak in a two-dimensional space. In order to find the source $[\theta_x, \theta_y]^\top$ of the leak, we need to decide at which optimal position $\mathbf{d} = [d_x, d_y]^\top$ to measure the gas concentration g. We here

---

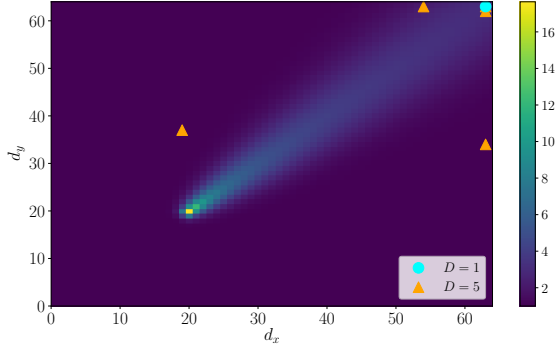[5]See the supplementary material for the joint distributions.

Figure 5. Marginal gas concentration as a function of $\mathbf{d}$ with the gas leak at the true source location. Also shown are the optimal designs for $D = 1$ measurement (cyan circles) and $D = 5$ measurements (orange triangles).
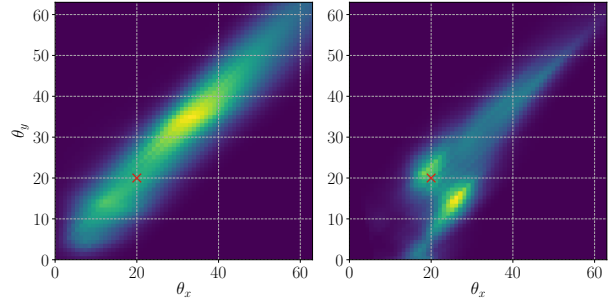

Figure 6. Marginal Posterior density for the source locations of the gas leak with one measurement (left) and five measurements (right). Shown as the red cross is the true gas leak location.

use the fluid simulator of Asenov et al. (2019) to forward-simulate the movement of gas particles. This fluid simulator solves the Navier-Stokes equations to yield gas concentrations g over a two-dimensional grid as a function of time t. Throughout this experiment we shall be assuming a steady-state solution, which is achieved by ignoring early time-steps. By using the translational properties of their simulator we can compute g at any position $\mathbf{d}$ for any source location $[\theta_x, \theta_y]^\top$ using only one forward simulation (see Asenov et al., 2019, for more information), greatly speeding up computation times. We assume that our measurements are subject to additive Gamma observational noise $\eta \sim \Gamma(2, 0.2)$. In the fluid simulator, we choose to set the gas viscosity to 0, the gas diffusion to 10^{-4}, and set the wind speed to 0.1. Furthermore, the discrete grid size is set to 64×64; a finer grid would yield more accurate localisations, at the cost of higher computation times.

While we keep most parameters of the fluid simulator constant, we assume uncertainty over the wind direction W_d, which introduces additional stochasticity. To do so, we put a (discretised) Gaussian prior over the wind direction W_d, centered at $45°$ from east with a standard deviation of $5°$. In practice, this means that we define the wind direction over a grid $W_d \in \{30.0, 32.5, \ldots, 60.0\}$ with a corresponding set of prior probabilities given by $p(W_d) \propto \mathcal{N}(W_d; 45, 5^2)$. Including this nuisance parameter, the model parameters for this problem are now given by $\boldsymbol{\theta} = [\theta_x, \theta_y, W_d]^\top$.

For this particular fluid simulator, we do not have gradients of the sampling path. Thus, we have to revert to using Bayesian optimisation to optimise the design variable $\mathbf{d}$, as explained in Section 2.2. In order to train $T_{\boldsymbol{\psi}}(\boldsymbol{\theta}, \mathbf{y})$ at a particular design $\mathbf{d}$ we use a one-layer neural network with 100 hidden units and a ReLU activation function after the input layer. The neural network parameters $\boldsymbol{\psi}$ are still

optimised by means of the Adam optimiser with a learning rate of $l_{\boldsymbol{\psi}} = 10^{-3}$. Before training, we simulate 10,000 uniform prior samples $\boldsymbol{\theta}^{(i)}$. During every re-training of $T_{\boldsymbol{\psi}}(\boldsymbol{\theta}, \mathbf{y})$ at a different design $\mathbf{d}$ we then simulate 10,000 corresponding samples of the gas concentration g. This procedure is repeated until we converge to the optimal design of $\mathbf{d}^* = [64, 64]^\top$. We visualise the optimal design location in Figure 5 and the supplementary material shows a plot of the GP posterior mean of the MI lower bound after training. The optimal design matches the fact that the wind is pointing north east, hence making upstream designs in that direction more informative. To illustrate this, in Figure 5 we also show the gas concentration on a grid of $\mathbf{d}$, marginalised over the wind speed, using $[\theta_x, \theta_y]^\top_{\text{true}} = [20, 20]^\top$ as the true source location.

Similar to before, given an optimal design $\mathbf{d}^*$ we can make a real-world measurement of the gas concentration g^*, generated using $[\theta_x, \theta_y]^\top_{\text{true}} = [20, 20]^\top$ and a true wind direction of $45°$. We then use the trained network $T_{\boldsymbol{\psi}^*}(\boldsymbol{\theta}, g)$ at the optimal design $\mathbf{d}^*$ to obtain $p(\boldsymbol{\theta} \mid g^*, \mathbf{d}^*)$ without any additional computation using (10). We evaluate this density over the 64×64 grid of possible gas leak locations, as shown in the left in Figure 6, where we have marginalised out the wind direction because it is a nuisance variable. While broad, the posterior accurately captures the direction of the gas leak relative to the measurement location.

As can be seen in the left in Figure 6, accurately locating a gas leak with only one measurement is extremely difficult. We wish to rectify this by increasing the number of measurements to $D = 5$. As a consequence, our design variable $\mathbf{d} = [d_{1,x}, d_{1,y}, \ldots, d_{5,x}, d_{5,y}]^\top$ is now 10-dimensional, with corresponding gas concentration measurements of $\mathbf{g} = [g_1, \ldots, g_5]^\top$. As before, we test different architectures and find out that a one-layer neural network with 200 hidden units has the best performance; all other

settings are kept as before. Repeating the BO optimisation procedure we obtain an optimal design $\mathbf{d}^*$ that we visualise in Figure 5. The optimal designs are spread out at the edges of the grid where the wind is pointing to. This is again intuitive, as we want to be measuring upstream to ensure that we do not get concentration measurements of zero. Furthermore, we have one optimal design away from the edges that helps us localise the gas leak better. Depending on the starting position of the experimenter, one could then determine the optimal flight path to take all of these measurements.

Taking real-world observations at $\mathbf{d}^*$ yields gas concentration measurements $\mathbf{g}^*$. Using these and the trained neural network we compute the posterior density by means of (10), which we show in the right in Figure 6. The posterior distribution is more concentrated for 5 measurements than for 1 measurement. While it is bimodal, the true source location is covered by one of the modes. This bimodality arises because of the uncertainty in the wind direction; if we know the wind direction, the posterior is unimodal (see the supplementary material).

4. Conclusions

In this work we have presented a novel approach to Bayesian experimental design for implicit models, where the data-generation distribution is intractable but sampling is possible. Our method uses a neural network to maximise a lower bound on the mutual information between data and parameters. This allows us to find both the optimal experimental designs and the corresponding posterior distribution by training a single neural network.

Whenever we have access to the gradients of the data sampling path of the implicit model, we can update neural network parameters and designs at the same time using gradient ascent, greatly enhancing scalability. While this is possible for a large and interesting class of implicit models, in situations where gradients with respect to the designs are not available, we can still update the neural network parameters using gradient ascent and the designs by means of e.g. Bayesian Optimisation.

In a variety of simulation studies, we found that our method provides suitable and intuitive experimental designs, as well as reasonable posterior distributions, even in higher design dimensions.

It would be interesting to apply our method to sequential experimental design where we can update the posterior of the model parameters after each real-world experiment performed. While there is some work on extending sequential experimental design to the significantly more difficult area of implicit models (e.g. Kleinegesse et al., 2020; Hainy et al., 2016), we believe that there is still much to explore, especially in the context of our method. Additionally, it might

be worthwhile to investigate the extension of our method to the setting of experimental design for model discrimination or the prediction of future observations.

References

Agostinelli, S., Allison, J., Amako, K., Apostolakis, J., M Araujo, H., Arce, P., Asai, M., A Axen, D., Banerjee, S., Barrand, G., Behner, F., Bellagamba, L., Boudreau, J., Broglia, L., Brunengo, A., Chauvie, S., Chuma, J., Chytracek, R., Cooperman, G., and Zschiesche, D. Geant4-a simulation toolkit. *Nuclear Instruments and Methods in Physics Research*, 506:250, 07 2003.

Asenov, M., Rutkauskas, M., Reid, D., Subr, K., and Ramamoorthy, S. Active localization of gas leaks using fluid simulation. *IEEE Robotics and Automation Letters*, PP, 01 2019.

Belghazi, M. I., Baratin, A., Rajeshwar, S., Ozair, S., Bengio, Y., Courville, A., and Hjelm, D. Mutual information neural estimation. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 531–540, Stockholmsmässan, Stockholm Sweden, 10–15 Jul 2018. PMLR.

Chen, Y. and Gutmann, M. U. Adaptive Gaussian copula ABC. In Chaudhuri, K. and Sugiyama, M. (eds.), *Proceedings of Machine Learning Research*, volume 89 of *Proceedings of Machine Learning Research*, pp. 1584–1592. PMLR, 16–18 Apr 2019.

Corander, J., Fraser, C., Gutmann, M., Arnold, B., Hanage, W., Bentley, S., Lipsitch, M., and Croucher, N. Frequency-dependent selection in vaccine-associated pneumococcal population dynamics. *Nature Ecology & Evolution*, 1: 1950–1960, 2017. ISSN 2397-334X.

Donnet, S. and Samson, A. A review on estimation of stochastic differential equations for pharmacokinetic/pharmacodynamic models. *Advanced Drug Delivery Reviews*, 65(7):929 – 939, 2013. ISSN 0169-409X. doi: https://doi.org/10.1016/j.addr.2013.03.005. Mathematical modeling of systems pharmacogenomics towards personalized drug delivery.

Donsker, M. D. and Varadhan, S. R. S. Asymptotic evaluation of certain markov process expectations for large time.

iv. *Communications on Pure and Applied Mathematics*, 36(2):183–212, 1983.

Drovandi, C. C. and Pettitt, A. N. Bayesian experimental design for models with intractable likelihoods. *Biometrics*, 69(4):937–948, 2013. doi: 10.1111/biom.12081.

Foster, A., Jankowiak, M., Bingham, E., Horsfall, P., Teh, Y. W., Rainforth, T., and Goodman, N. Variational Bayesian optimal experimental design. In *Advances in Neural Information Processing Systems 32*, pp. 14036–14047. Curran Associates, Inc., 2019a.

Foster, A., Jankowiak, M., O'Meara, M., Whye Teh, Y., and Rainforth, T. A Unified Stochastic Gradient Approach to Designing Bayesian-Optimal Experiments. *arXiv e-prints*, art. arXiv:1911.00294, Nov 2019b.

Grimmett, G. and Stirzaker, D. *Probability and random processes*, volume 80. Oxford university press, 2001.

Gutmann, M. U. and Corander, J. Bayesian optimization for likelihood-free inference of simulator-based statistical models. *Journal of Machine Learning Research*, 17(125): 1–47, 2016.

Hainy, M., Drovandi, C. C., and McGree, J. Likelihood-free extensions for Bayesian sequentially designed experiments. In Kunert, J., Muller, C. H., and Atkinson, A. C. (eds.), *11th International Workshop in Model-Oriented Design and Analysis (mODa 2016)*, pp. 153–161, Hamminkeln, Germany, 2016. Springer. doi: 10.1007/978-3-319-31266-8_18.

Kleinegesse, S. and Gutmann, M. U. Efficient Bayesian experimental design for implicit models. In Chaudhuri, K. and Sugiyama, M. (eds.), *Proceedings of Machine Learning Research*, volume 89 of *Proceedings of Machine Learning Research*, pp. 476–485. PMLR, Apr 2019.

Kleinegesse, S., Drovandi, C., and Gutmann, M. U. Sequential Bayesian Experimental Design for Implicit Models via Mutual Information. *arXiv e-prints [accepted at Bayesian Analysis]*, art. arXiv:2003.09379, 2020.

Lintusaari, J., Gutmann, M., Dutta, R., Kaski, S., and Corander, J. Fundamentals and recent developments in approximate Bayesian computation. *Systematic Biology*, 66(1):e66–e82, January 2017. ISSN 1063-5157. doi: 10.1093/sysbio/syw077.

M. Schafer, C. and Freeman, P. Likelihood-free inference in cosmology: Potential for the estimation of luminosity functions. 209:3–19, 01 2012.

Mohamed, S., Rosca, M., Figurnov, M., and Mnih, A. Monte Carlo Gradient Estimation in Machine Learning. *arXiv e-prints*, art. arXiv:1906.10652, Jun 2019.

Müller, P. Simulation-based optimal design. *Bayesian Statistics*, 6:459–474, 1999.

Nguyen, X., Wainwright, M. J., and Jordan, M. I. Estimating divergence functionals and the likelihood ratio by convex risk minimization. *IEEE Transactions on Information Theory*, 56(11):5847–5861, Nov 2010.

Nowozin, S., Cseke, B., and Tomioka, R. F-GAN: Training generative neural samplers using variational divergence minimization. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, NIPS'16, pp. 271–279, Red Hook, NY, USA, 2016. Curran Associates Inc. ISBN 9781510838819.

Ong, V., Nott, D., Tran, M.-N., Sisson, S., and Drovandi, C. Likelihood-free inference in high dimensions with synthetic likelihood. *Computational Statistics & Data Analysis*, 128:271–291, Dec 2018.

Overstall, A. M. and McGree, J. M. Bayesian design of experiments for intractable likelihood models using coupled auxiliary models and multivariate emulation. *ArXiv e-prints: 1803.07018*, March 2018.

Pan, S. J. and Yang, Q. A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22 (10):1345–1359, Oct 2010.

Papamakarios, G. and Murray, I. Fast epsilon-free inference of simulation models with Bayesian conditional density estimation. In Lee, D. D., Sugiyama, M., Luxburg, U. V., Guyon, I., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 29*, pp. 1028–1036. Curran Associates, Inc., 2016.

Papamakarios, G., Sterratt, D., and Murray, I. Sequential neural likelihood: Fast likelihood-free inference with autoregressive flows. In Chaudhuri, K. and Sugiyama, M. (eds.), *Proceedings of Machine Learning Research*, volume 89, pp. 837–848, Proceedings of Machine Learning Research, 2019. PMLR.

Poole, B., Ozair, S., Van Den Oord, A., Alemi, A., and Tucker, G. On variational bounds of mutual information. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 5171–5180, Long Beach, California, USA, 09–15 Jun 2019. PMLR.

Price, D. J., Bean, N. G., Ross, J. V., and Tuke, J. On the efficient determination of optimal Bayesian experimental designs using ABC: A case study in optimal observation of epidemics. *Journal of Statistical Planning and Inference*, 172:1 – 15, 2016. ISSN 0378-3758. doi: https://doi.org/10.1016/j.jspi.2015.12.008.

Price, D. J., Bean, N. G., Ross, J. V., and Tuke, J. An induced natural selection heuristic for finding optimal Bayesian experimental designs. *Computational Statistics and Data Analysis*, 126:112 – 124, 2018. ISSN 0167-9473. doi: https://doi.org/10.1016/j.csda.2018.04.011.

Rasmussen, C. E. and Williams, C. K. I. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2005. ISBN 026218253X.

Ryan, E., Drovandi, C. C., Thompson, H., and Pettitt, A. N. Towards Bayesian experimental design for nonlinear models that require a large number of sampling times. *Computational Statistics and Data Analysis*, 70:45–60, February 2014.

Ryan, E. G., Drovandi, C. C., McGree, J. M., and Pettitt, A. N. A review of modern computational algorithms for Bayesian optimal design. *International Statistical Review*, 84(1):128–154, 2016.

Shahriari, B., Swersky, K., Wang, Z., Adams, R. P., and de Freitas, N. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104 (1):148–175, Jan 2016.

Sisson, S., Fan, Y., and Beaumont, M. *Handbook of Approximate Bayesian Computation.*, chapter Overview of Approximate Bayesian Computation. Chapman and Hall/CRC Press, 2018.

Spall, J. *Introduction to Stochastic Search and Optimization: Estimation, Simulation, and Control*. Wiley Series in Discrete Mathematics and Optimization. Wiley, 2005.

Thomas, O., Dutta, R., Corander, J., Kaski, S., and Gutmann, M. U. Likelihood-free inference by ratio estimation. *ArXiv e-prints: 1611.10242*, art. arXiv:1611.10242, Nov 2016.

Vo, B. N., Drovandi, C. C., Pettitt, A. N., and Simpson, M. J. Quantifying uncertainty in parameter estimates for stochastic models of collective cell spreading using approximate Bayesian computation. *Mathematical Biosciences*, 263:133 – 142, 2015.