
Design of Experiments for Model Discrimination Hybridising Analytical and Data-Driven Approaches

Simon Olofsson¹ Marc Peter Deisenroth^{1,2} Ruth Misener¹

Abstract

Healthcare companies must submit pharmaceutical drugs or medical devices to regulatory bodies before marketing new technology. Regulatory bodies frequently require transparent and interpretable computational modelling to justify a new healthcare technology, but researchers may have several competing models for a biological system and too little data to discriminate between the models. In design of experiments for model discrimination, the goal is to design maximally informative physical experiments in order to discriminate between rival predictive models. Prior work has focused either on analytical approaches, which cannot manage all functions, or on data-driven approaches, which may have computational difficulties or lack interpretable marginal predictive distributions. We develop a methodology introducing Gaussian process surrogates in lieu of the original mechanistic models. We thereby extend existing design and model discrimination methods developed for analytical models to cases of non-analytical models in a computationally efficient manner.

1. Introduction

A patient's health and safety is an important concern, and healthcare is a heavily regulated industry. In the USA and the EU, private and public regulatory bodies exist on federal/union and state levels (Field, 2008; Hervey, 2010). Healthcare companies applying to market a new drug or medical device must submit extensive technical information to the regulatory bodies. In the USA and the EU, the Food & Drug Administration (FDA) and European Medicines

Agency (EMA) handle these applications, respectively. The FDA require that the technical information contains e.g. the chemical composition of the drug, how the drug affects the human body (pharmacodynamics), how the body affects the drug (pharmacokinetics), and methods for drug manufacturing, packaging and quality control (U.S. Food & Drug Administration, 2018). Likewise, applying to market a new medical device requires submitting extensive technical information. A medical device can be any appliance, software or material used for non-pharmacological diagnosis, monitoring or treatment of a disease or condition (European Union, 2017). An important part of proving the efficacy and safety of a new drug or medical device is showing that its effects on the patient can be predicted and interpreted, e.g. via mathematical and computational models.

These transparency and interpretability requirements, combined with limited amounts of available experimental data, make models learnt solely from observed data unsuitable for proving the efficacy and safety of a new drug or medical device to the regulatory bodies. Hence researchers use explicit parametric models. For drugs, the pharmacokinetics (how the human body affects the drug) is often modelled using systems of ordinary differential equations. These models elucidate how the drug is absorbed and distributed through the body (e.g. brain, kidneys and liver) under different dosing profiles and methods of drug administration (e.g. orally or intravenously) (see Figure 1). The EMA pharmacokinetic model guidelines state that regulatory submissions should include detailed descriptions of the models, i.e. justification of model assumptions, parameters and their biochemical plausibility, as well as parameter sensitivity analyses (European Medicines Agency, 2011, p. 16; 2016, p. 3).

A key problem is to identify a mathematical model that usefully explains and predicts the behaviour of the pharmacokinetic process (Heller et al., 2018). Researchers will often have several models (i.e. hypotheses) about an underlying mechanism of the process, but insufficient experimental data to discriminate between the models (Heller et al., 2018). Additional experiments are generally expensive and time-consuming. Pharmacokinetic experiments typically take 8–24 hours for mice and rats, and 12–100 hours for simians and humans (Ogungbenro & Aarons, 2008; Tuntland et al.,

¹Dept. of Computing, Imperial College London, United Kingdom. ²PROWLER.io, United Kingdom. Correspondence to: Simon Olofsson <simon.olofsson15@imperial.ac.uk>, Marc Peter Deisenroth <m.deisenroth@imperial.ac.uk>, Ruth Misener <r.misener@imperial.ac.uk>.




Figure 1. Different aspects of pharmacokinetic models. (Top left) simple pharmacokinetic model with four compartments: digestive system, liver, bloodstream and other tissue. The drug can enter as tablets into the digestive system or as an injection into the bloodstream, and leave the system by being excreted as waste. (Bottom left) the effect on drug concentration in a patient from two different dosing profiles: (i) half the dose twice per day, or (ii) the full dose once per day. (Centre) two different methods of drug administration: orally in the form of tablets, or intravenously as an injection, and (Right) their different effects: an injection has a quicker effect than tablets.

2014). It is key to design experiments yielding maximally informative outcomes with respect to model discrimination.

Contribution We tackle limitations in existing analytical and data-driven approaches to design of experiments for model discrimination by bridging the gap between the classical analytical methods and the Monte Carlo-based approaches. These limitations include the classical analytical methods’ difficulty to manage non-analytical model functions, and the computational cost for the data-driven approaches. To this end, we develop a novel method of replacing the original parametric models with probabilistic, non-parametric Gaussian process (GP) surrogates learnt from model evaluations. The GP surrogates are flexible regression tools that allow us to extend classical analytical results to non-analytical models in a computationally efficient manner, while providing us with confidence bounds on the model predictions.

The method described in this paper has been implemented in a python software package, *GPdoemd* (2018), publicly available via GitHub under an MIT License.

2. Model Discrimination

Model discrimination aims to discard inaccurate models, i.e., hypotheses that are not supported by the data. Assume $M > 1$ models \mathcal{M}_i are given, with corresponding model

functions $f_i(x, \theta_i)$. The model is a hypothesis and collection of assumptions about a biological system, e.g. the pharmacokinetic model in the top left of Figure 1, and the function is the model’s mathematical formulation that can be used for predictions, e.g. the drug concentrations on the right in Figure 1. Each model function f_i takes as inputs the design variables $x \in \mathcal{X} \subset \mathbb{R}^d$ and model parameters $\theta_i \in \Theta_i \subset \mathbb{R}^{D_i}$. The design variables x , e.g. drug administration and dosing profile, specify a physical experiment that can be run on the system described by the models, with \mathcal{X} defining the operating system boundaries. The classical setting allows tuning of the model parameters θ_i to fit the model to data. The system has $E \geq 1$ target dimensions, i.e. $f_i: \mathbb{R}^{d+D_i} \rightarrow \mathbb{R}^E$. We denote the function for each target dimension e as $f_{i,(e)}$. This notation distinguishes between the e^{th} target dimension $f_{(e)}$ of a model function f , and model \mathcal{M}_i ’s function f_i . Table A in the supplementary material summarises the notation used.

We start with an initial set of N_0 experimental measurements $y_n = f_{\text{true}}(x_n) + \epsilon_n$, $n = 1, \dots, N_0$. A common assumption in pharmacokinetics is that the measurement noise term ϵ_n is i.i.d. zero-mean Gaussian distributed with covariance Σ_{exp} (Steimer et al., 1984; Ette & Williams, 2004; Tsimring, 2014). Skew in the data distribution can be handled, e.g. through log-transformation. The experimental data set is denoted $\mathcal{D}_{\text{exp}} = \{y_j, x_j\}$. The initial N_0 experimental data points are insufficient to discriminate

between the models, i.e. $\forall i : p(\mathcal{M}_i | \mathcal{D}_{\text{exp}}) \approx 1/M$. We are agnostic to the available experimental data, and wish to incorporate all of it in the model discrimination.

2.1. Classical Analytical Approach

Methods for tackling the design of experiments for discriminating simple, analytical models have been around for over 50 years. [Box & Hill \(1967\)](#) study the expected change $\mathbb{E}[\Delta S] = \mathbb{E}[S_{N+1}] - S_N$ in Shannon entropy $S_N = -\sum_i \pi_{i,N} \log \pi_{i,N}$ from making an additional experimental observation, where the posterior model probability $\pi_{i,N+1} = p(\mathbf{y} | \mathcal{M}_i) \pi_{i,N} / p(\mathbf{y})$ and $p(\mathbf{y}) = \sum_i p(\mathbf{y} | \mathcal{M}_i) \pi_{i,N}$. [Box & Hill \(1967\)](#) develop a design criterion D_{BH} by maximising the upper bound on $\mathbb{E}[\Delta S]$. The expression for D_{BH} can be found in the supplementary material. [MacKay \(1992a;b\)](#) derives a similar expression. [Box & Hill \(1967\)](#) choose the next experiment by finding $\mathbf{x}^* = \arg \max_{\mathbf{x}} D_{\text{BH}}$. Design of experiments continues until there exists a model probability $\pi_{N,i} \approx 1$.

[Buzzi-Ferraris et al. \(1990\)](#) propose a new design criterion D_{BF} (see the supplementary material) from a frequentist point-of-view. They suggest using a χ^2 test with $NE - D_i$ degrees of freedom for the model discrimination. The null hypothesis for each model is that the model has generated the experimental data. Design of experiments continues until only one model passes the χ^2 test. Models are not ranked against each other since [Buzzi-Ferraris et al. \(1990\)](#) argue this simply leads to the least inaccurate model being selected. The χ^2 procedure is more robust against—but not immune to—favouring the least inaccurate model.

[Michalik et al. \(2010\)](#) proceed from the Akaike information criterion (AIC) as the model discrimination criterion to derive a design criterion D_{AW} from the Akaike weights w_i . The expression for D_{AW} can be found in the supplementary material. Design of experiments continues until there exists an Akaike weight $w_i \approx 1$. [Box & Hill \(1967\)](#) and [Michalik et al. \(2010\)](#) implicitly favour the least inaccurate model.

In order to account for the uncertainty in the model parameters θ , the classical analytical approach ([Prasad & Someswara Rao, 1977](#); [Buzzi-Ferraris et al., 1984](#)) is to approximate the model parameters as being Gaussian distributed $\mathcal{N}(\hat{\theta}, \Sigma_{\theta})$ around the best-fit parameter values $\hat{\theta}$. The covariance Σ_{θ} is given by a first-order approximation

$$\Sigma_{\theta}^{-1} = \sum_{n=1}^N \nabla_{\theta} \mathbf{f}_i^n \top \Sigma_{\text{exp}}^{-1} \nabla_{\theta} \mathbf{f}_i^n, \quad (1)$$

where $\mathbf{f}_i^n = \mathbf{f}_i(\mathbf{x}_n, \theta)$, $\mathbf{x}_n \in \mathcal{D}_{\text{exp}}$, and the gradient $\nabla_{\theta} \mathbf{f}_i = \partial \mathbf{f}_i(\mathbf{x}, \theta_i) / \partial \theta_i |_{\theta_i = \hat{\theta}_i}$. The predictive distribution with θ approximately marginalised out becomes $p(\mathbf{f}_i | \mathbf{x}, \mathcal{D}_{\text{exp}}) = \mathcal{N}(\mathbf{f}_i(\mathbf{x}, \hat{\theta}), \hat{\Sigma}_i(\mathbf{x}))$, where the covariance is $\hat{\Sigma}_i(\mathbf{x}) = \nabla_{\theta} \mathbf{f}_i \Sigma_{\theta} \nabla_{\theta} \mathbf{f}_i^{\top}$.

Limitations Most model functions for healthcare applications such as pharmacokinetics or medical devices are not analytical. They are often, from a practical point-of-view, complex black boxes of legacy code, e.g. large systems of ordinary or partial differential equations. We can evaluate the model function \mathbf{f}_i , but the gradients $\nabla_{\theta} \mathbf{f}_i$ are not readily available. Specialist software can retrieve the gradients $\nabla_{\theta} \mathbf{f}_i$ from systems of ordinary or partial differential equations, but this requires implementing and maintaining the models in said specialist software packages. Other types of mathematical models may require solving an optimisation problem ([Boukouvala et al., 2016](#)). For model discrimination, we wish to be agnostic with regards to the software implementation or model type, since this flexibly (i) allows faster model prototyping and development, and (ii) satisfies the personal preferences of researchers and engineers.

2.2. Bayesian Design of Experiments

Methods for accommodating non-analytical model functions have developed in parallel with increasing computer speed. These methods are typically closer to fully Bayesian than the classical analytical methods. [Vanlier et al. \(2014\)](#) approximate the marginal predictive distributions of M models and their Jensen-Shannon divergence using Markov chain Monte Carlo (MCMC) sampling and a k -nearest neighbours density estimate. On the downside, the density estimates become less accurate as the number of experimental observations increases ([Vanlier et al., 2014](#)) and the method is computationally intensive ([Ryan et al., 2016](#)).

Instead of studying design of experiments for model discrimination, statisticians have focused on design of experiments and model discrimination separately ([Chaloner & Verdinelli, 1995](#)). Design of experiments can also be used to aid model parameter estimation. They solve problems of the type

$$\mathbf{x}^* = \arg \max_{\mathbf{x}} \mathbb{E}_{\theta, \mathbf{y}_{N+1}} [U(\mathbf{y}_{N+1}, \mathbf{x}, \theta)] \quad (2)$$

where the utility function $U(\cdot, \cdot, \cdot)$ serves to either discriminate models or estimate parameters. Criteria for model discrimination are handled separately, usually under the name of model selection or hypothesis testing.

[Ryan et al. \(2015\)](#) use a Laplace approximation of the posterior distribution $p(\theta | \mathcal{D}_{\text{exp}})$ combined with importance sampling. [Drovandi et al. \(2014\)](#) develop a method based on sequential Monte Carlo (SMC) that is faster than using MCMC. [Woods et al. \(2017\)](#) use a Monte Carlo approximation $\Phi(\mathbf{x}) = \frac{1}{B} \sum_b^B U(\mathbf{y}_b, \mathbf{x}, \theta_b)$ with $(\mathbf{y}_b, \theta_b) \sim p(\mathbf{y}_b, \theta_b | \mathbf{x})$ on which they place a Gaussian process prior.

Limitations These methods agnostically accommodate non-analytical models using a Monte Carlo-approach but require exhaustive model sampling in the model parameter space. On a case study with four models, each with ten

model parameters, and two design variables, the Vanlier et al. (2014) MCMC method requires six days of wall-clock time to compute two sets of four marginal predictive distributions. This cost is impractical for pharmacokinetic applications, where a physical experiment takes $\approx 1-4$ days. SMC methods can converge faster than MCMC methods (Woods et al., 2017). But SMC methods can suffer from sample degeneracy, where only a few particles receive the vast majority of the probability weight (Li et al., 2014). Also, convergence analysis for MCMC and SMC methods is difficult.

Furthermore, these methods are only for design of experiments. Once an experiment is executed, the model discrimination issue remains. In this case the marginal predictive distributions $p(\mathbf{f}_i | \mathbf{x}, \mathcal{D}_{\text{exp}})$ would enable calculating the model likelihoods.

3. Method

We wish to exploit results from the classical analytical approach to design of experiments for model discrimination, while considering that we may deal with non-analytical model functions. The key idea is replacing the original (non-analytical) model functions with analytical and differentiable surrogate models. The surrogate models can be learnt from model evaluations. We choose to use GPs to construct the surrogate models because they are flexible regression tools that allow us to extract gradient information.

A GP is a collection of random variables, any finite subset of which is jointly Gaussian distributed (Rasmussen & Williams, 2006, p. 13) and fully specified by a mean function $m(\cdot)$ and a kernel function $k(\cdot, \cdot)$. Prior knowledge about the model functions \mathbf{f}_i can be incorporated into the choice of m and k . To construct the GP surrogates, target data is needed. The purpose of the surrogate model is to predict the outcome of a model evaluation $f_i(\mathbf{x}, \boldsymbol{\theta}_i)$, rather than the outcome of a physical experiment $f_{\text{true}}(\mathbf{x})$. Hence, for each model \mathcal{M}_i we create a data set $\mathcal{D}_{\text{sim},i} = \{\mathbf{f}_\ell, \tilde{\mathbf{x}}_\ell, \tilde{\boldsymbol{\theta}}_\ell; \ell = 1, \dots, N_{\text{sim},i}\}$ of simulated observations, i.e. model evaluations $\tilde{\mathbf{f}}_\ell = \mathbf{f}_i(\tilde{\mathbf{x}}_\ell, \tilde{\boldsymbol{\theta}}_\ell)$, where $\tilde{\mathbf{x}}_\ell$ and $\tilde{\boldsymbol{\theta}}_\ell$ are sampled from \mathcal{X} and Θ_i , respectively. Note the difference between the experimental data set \mathcal{D}_{exp} and the simulated data sets $\mathcal{D}_{\text{sim},i}$: whereas the data set \mathcal{D}_{exp} has few (noisy, biological) observations, $\mathcal{D}_{\text{sim},i}$ has many (noise-free, artificial) observations.

We place independent GP priors $f_{i,(e)} \sim \mathcal{GP}(m_{i,(e)}, k_{i,(e)})$ on each target dimension $e = 1, \dots, E$ of the model function \mathbf{f}_i , $i = 1, \dots, M$. For notational simplicity, let $f = f_{i,(e)}$, $m = m_{i,(e)}$, $k = k_{i,(e)}$ and $\mathbf{z}^\top = [\mathbf{x}^\top, \boldsymbol{\theta}_i^\top]$. Given $\mathcal{D}_{\text{sim},i}$ with model evaluations $[\tilde{\mathbf{f}}]_\ell = \tilde{f}_\ell$ at locations $[\tilde{\mathbf{Z}}]_\ell = \tilde{\mathbf{z}}_\ell$ the predictive distribution $p(f | \mathbf{z}_*, \tilde{\mathbf{f}}, \tilde{\mathbf{Z}}) = \mathcal{N}(\mu(\mathbf{z}_*), \sigma^2(\mathbf{z}_*))$ at a test point \mathbf{z}_* is given by

$$\mu(\mathbf{z}_*) = m(\mathbf{z}_*) + \mathbf{k}_*^\top \mathbf{K}^{-1}(\tilde{\mathbf{f}} - \mathbf{m}), \quad (3)$$

$$\sigma^2(\mathbf{z}_*) = k(\mathbf{z}_*, \mathbf{z}_*) - \mathbf{k}_*^\top \mathbf{K}^{-1} \mathbf{k}_*, \quad (4)$$

where the covariance vector is $[\mathbf{k}_*]_\ell = k(\tilde{\mathbf{z}}_\ell, \mathbf{z}_*)$, the kernel matrix is $[\mathbf{K}]_{\ell_1, \ell_2} = k(\tilde{\mathbf{z}}_{\ell_1}, \tilde{\mathbf{z}}_{\ell_2}) + \sigma_n^2 \delta_{\ell_1, \ell_2}$, and the target mean is $[\mathbf{m}]_\ell = m(\tilde{\mathbf{z}}_\ell)$, with a small artificial noise variance $\sigma_n^2 \ll 1$ added for numerical stability. Note the simplified notation $\mu = \mu_{i,(e)}$ and $\sigma^2 = \sigma_{i,(e)}^2$.

Pharmacokinetic models are often formulated as systems of ordinary differential equations, e.g. Han et al. (2006), where the solutions are smooth functions. If we wish to model these functions with a GP, it makes sense to use a covariance function that encodes the smoothness of the underlying function, e.g. the radial basis function (RBF) covariance function $k_{\text{RBF}}(\mathbf{z}, \mathbf{z}') = \rho \exp(-\frac{1}{2}(\mathbf{z} - \mathbf{z}')^\top \boldsymbol{\Lambda}^{-1}(\mathbf{z} - \mathbf{z}'))$, since it corresponds to Bayesian linear regression with an infinite number of basis functions (Rasmussen & Williams, 2006). Our method extends to other choices of covariance function.

Each function sample f drawn from a zero-mean GP is a linear combination $f(\mathbf{x}) = \sum_j \omega_j k(\mathbf{x}, \mathbf{x}_j)$ of covariance function terms (Rasmussen & Williams, 2006, p. 17). Different model characteristics, e.g. smoothness or periodicity, make alternate covariance functions appropriate. The RBF covariance function is common, but its strong smoothness prior may be too unrealistic for some systems. Other choices include: Matérn, polynomial, periodic, and non-stationary covariance functions. Sums or products of valid covariance functions are still valid covariance functions. Covariance functions can also model different behaviour over different input regions (Ambrogioni & Maris, 2016). Choosing the covariance function requires some prior knowledge about the modelled function characteristics, e.g. a periodic covariance function may be appropriate to model periodic drug dosing.

In k_{RBF} , and other similar covariance functions, ρ^2 is the signal variance and $\boldsymbol{\Lambda} = \text{diag}(\lambda_{1:D}^2)$ is a diagonal matrix of squared length scales, with $\mathbf{z} \in \mathbb{R}^{D'}$. Together, ρ , $\boldsymbol{\Lambda}$ and the artificial noise variance σ_n^2 (in \mathbf{K}) are the GP hyperparameters. Because the data set $\mathcal{D}_{\text{sim},i}$ is arbitrarily large, we use point-estimate hyperparameters learnt through standard evidence maximisation.

We study a single model function $\mathbf{f} = \mathbf{f}_i$ and its GP surrogates. Using the GP surrogates we can predict the outcome of a model evaluation $\mathbf{f}(\mathbf{x}, \boldsymbol{\theta}) \sim p(\mathbf{f} | \mathbf{z}, \mathcal{D}_{\text{sim},i}) = \mathcal{N}(\boldsymbol{\mu}(\mathbf{z}), \boldsymbol{\Sigma}(\mathbf{z}))$, where $\boldsymbol{\mu} = [\mu_{(1:E)}]^\top$ and $\boldsymbol{\Sigma} = \text{diag}(\sigma_{(1:E)}^2)$. We approximate the (marginal) predictive distribution $p(\mathbf{f} | \mathbf{x}, \mathcal{D}_{\text{sim},i}, \mathcal{D}_{\text{exp}})$ to account for the model parameter uncertainty $\boldsymbol{\theta} \sim p(\boldsymbol{\theta} | \mathcal{D}_{\text{exp}})$ due to parameter estimation from the noisy experimental data \mathcal{D}_{exp} .

3.1. Model Parameter Marginalisation

We assume that the experimental observations $\mathbf{y}_{1:N} \in \mathcal{D}_{\text{exp}}$ have i.i.d. zero-mean Gaussian distributed measurement

noise with known covariance Σ_{exp} . Let $\nabla_{\theta}g$ denote the gradient $\partial g(\mathbf{x}, \boldsymbol{\theta})/\partial \boldsymbol{\theta}|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}}$ of a function g with respect to the model parameter $\boldsymbol{\theta}$, evaluated at $\boldsymbol{\theta} = \hat{\boldsymbol{\theta}}$, and $\nabla_{\boldsymbol{\theta}}^2 g$ denote the corresponding Hessian.

We approximate the model parameter posterior distribution $p(\boldsymbol{\theta} | \mathcal{D}_{\text{exp}}) \approx \mathcal{N}(\hat{\boldsymbol{\theta}}, \Sigma_{\theta})$ as a Gaussian distribution around the best-fit model parameter values $\hat{\boldsymbol{\theta}}$. Using the same first-order approximation as the classical analytical method in (1), with $\mathbb{E}_{\boldsymbol{f}}[\boldsymbol{f} | \mathbf{x}, \boldsymbol{\theta}] = \boldsymbol{\mu}(\mathbf{x}, \boldsymbol{\theta})$, yields

$$\Sigma_{\theta}^{-1} = \sum_{n=1}^N \nabla_{\theta} \boldsymbol{\mu}^n \top \Sigma_{\text{exp}}^{-1} \nabla_{\theta} \boldsymbol{\mu}^n,$$

where $\boldsymbol{\mu}^n = \boldsymbol{\mu}(\mathbf{x}_n, \boldsymbol{\theta})$, with $\mathbf{x}_n \in \mathcal{D}_{\text{exp}}$. We approximate the marginal predictive distribution $p(\boldsymbol{f} | \mathbf{x}, \mathcal{D}_{\text{sim}}, \mathcal{D}_{\text{exp}}) \approx \mathcal{N}(\check{\boldsymbol{\mu}}(\mathbf{x}), \check{\Sigma}(\mathbf{x}))$ using a first- or second-order Taylor expansion around $\boldsymbol{\theta} = \hat{\boldsymbol{\theta}}$, where

$$\check{\boldsymbol{\mu}}(\mathbf{x}) = \mathbb{E}_{\theta}[\boldsymbol{\mu} | \mathbf{x}, \boldsymbol{\theta}] \quad (5)$$

$$\check{\Sigma}(\mathbf{x}) = \mathbb{E}_{\theta}[\Sigma | \mathbf{x}, \boldsymbol{\theta}] + \mathbb{V}_{\theta}[\boldsymbol{\mu} | \mathbf{x}, \boldsymbol{\theta}], \quad (6)$$

which we will detail in the following.

First-order Taylor approximation assumes that the model function \boldsymbol{f} is (approximately) linear in the model parameters in some neighbourhood around $\boldsymbol{\theta} = \hat{\boldsymbol{\theta}}$

$$\boldsymbol{\mu}_{(e)}(\mathbf{x}, \boldsymbol{\theta}) \approx \boldsymbol{\mu}_{(e)}(\mathbf{x}, \hat{\boldsymbol{\theta}}) + \nabla_{\theta} \boldsymbol{\mu}_{(e)}(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}), \quad (7)$$

$$\sigma_{(e)}^2(\mathbf{x}, \boldsymbol{\theta}) \approx \sigma_{(e)}^2(\mathbf{x}, \hat{\boldsymbol{\theta}}) + \nabla_{\theta} \sigma_{(e)}^2(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}). \quad (8)$$

Inserting (7) into (5) yields the approximate mean of the marginal predictive distribution

$$\check{\boldsymbol{\mu}}(\mathbf{x}) \approx \boldsymbol{\mu}(\mathbf{x}, \hat{\boldsymbol{\theta}}). \quad (9)$$

Similarly, inserting (8) into (6), the approximate expectation of the variance becomes

$$\mathbb{E}_{\theta}[\Sigma | \mathbf{x}, \boldsymbol{\theta}] \approx \Sigma(\mathbf{x}, \hat{\boldsymbol{\theta}}), \quad (10)$$

and the variance of the mean becomes

$$\mathbb{V}_{\theta}[\boldsymbol{\mu} | \mathbf{x}, \boldsymbol{\theta}] \approx \nabla_{\theta} \boldsymbol{\mu} \Sigma_{\theta} \nabla_{\theta} \boldsymbol{\mu}^{\top}. \quad (11)$$

(9)–(11) yield a tractable first-order Taylor approximation to $p(\boldsymbol{f} | \mathbf{x}, \mathcal{D}_{\text{sim}}, \mathcal{D}_{\text{exp}}) = \mathcal{N}(\check{\boldsymbol{\mu}}(\mathbf{x}), \check{\Sigma}(\mathbf{x}))$.

Second-order Taylor expansion assumes that the model function \boldsymbol{f} is (approximately) quadratic in the model parameters in some neighbourhood around $\boldsymbol{\theta} = \hat{\boldsymbol{\theta}}$

$$\begin{aligned} \boldsymbol{\mu}_{(e)}(\mathbf{x}, \boldsymbol{\theta}) &\approx \boldsymbol{\mu}_{(e)}(\mathbf{x}, \hat{\boldsymbol{\theta}}) + \nabla_{\theta} \boldsymbol{\mu}_{(e)}(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}) \\ &\quad + \frac{1}{2}(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}})^{\top} \nabla_{\theta}^2 \boldsymbol{\mu}_{(e)}(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}), \end{aligned} \quad (12)$$

$$\begin{aligned} \sigma_{(e)}^2(\mathbf{x}, \boldsymbol{\theta}) &\approx \sigma_{(e)}^2(\mathbf{x}, \hat{\boldsymbol{\theta}}) + \nabla_{\theta} \sigma_{(e)}^2(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}) \\ &\quad + \frac{1}{2}(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}})^{\top} \nabla_{\theta}^2 \sigma_{(e)}^2(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}). \end{aligned} \quad (13)$$

Inserting (12) into (5) yields the approximate mean of the marginal predictive distribution

$$\check{\boldsymbol{\mu}}_{(e)}(\mathbf{x}) \approx \boldsymbol{\mu}_{(e)}(\mathbf{x}, \hat{\boldsymbol{\theta}}) + \frac{1}{2} \text{tr}(\nabla_{\theta}^2 \boldsymbol{\mu}_{(e)} \Sigma_{\theta}). \quad (14)$$

Similarly, inserting (13) into (6) yields the diagonal approximate expectation of the variance with elements

$$\mathbb{E}_{\theta}[\sigma_{(e)}^2 | \mathbf{x}, \boldsymbol{\theta}] \approx \sigma_{(e)}^2(\mathbf{x}, \hat{\boldsymbol{\theta}}) + \frac{1}{2} \text{tr}(\nabla_{\theta}^2 \sigma_{(e)}^2 \Sigma_{\theta}). \quad (15)$$

The variance of the mean is a dense matrix with elements

$$\begin{aligned} \mathbb{V}_{\theta}[\boldsymbol{\mu}_{(e_1)}, \boldsymbol{\mu}_{(e_2)} | \mathbf{x}, \boldsymbol{\theta}] &\approx \nabla_{\theta} \boldsymbol{\mu}_{(e_1)} \Sigma_{\theta} \nabla_{\theta} \boldsymbol{\mu}_{(e_2)}^{\top} \\ &\quad + \frac{1}{2} \text{tr}(\nabla_{\theta}^2 \boldsymbol{\mu}_{(e_1)} \Sigma_{\theta} \nabla_{\theta}^2 \boldsymbol{\mu}_{(e_2)} \Sigma_{\theta}). \end{aligned} \quad (16)$$

(14)–(16) yield a tractable second-order Taylor approximation to $p(\boldsymbol{f} | \mathbf{x}, \mathcal{D}_{\text{sim}}, \mathcal{D}_{\text{exp}}) = \mathcal{N}(\check{\boldsymbol{\mu}}(\mathbf{x}), \check{\Sigma}(\mathbf{x}))$.

Figure 2 illustrates the process of constructing the GP surrogate models and the approximations of the marginal predictive distributions. Firstly, the model parameters are estimated from \mathcal{D}_{exp} . Secondly, \mathcal{D}_{sim} is created from model evaluations and used (i) to learn the GP surrogate’s hyperparameters, and (ii) as target data for GP predictions. Thirdly, the approximation $\mathcal{N}(\check{\boldsymbol{\mu}}(\mathbf{x}), \check{\Sigma}(\mathbf{x}))$ of the marginal predictive distribution is computed. Fourthly and lastly, the marginal predictive distribution is used to design the next experiment and approximate the model likelihood.

4. Results

The following demonstrates that our approach for model discrimination with GP surrogate models exhibits a similar performance to the classical analytical method, but can be used to design experiments to discriminate between non-analytical model functions. The GP surrogate method performance will be studied using two different case studies.

Case study 1 has analytical model functions, so we can compare how well our new method performs compared to the classical analytical methods in Section 2.1. Case study 2 has pharmacokinetic-like models consisting of systems of ordinary differential equations.

In both case studies we compute the following metrics:

- (A) the average number of additional experiments $N - N_0$ required for all incorrect models to be discarded.
- (SE) the standard error of the average (A).
- (S) success rate; the proportion of tests in which all incorrect models were discarded.
- (F) failure rate; the proportion of tests in which the correct (data-generating) model was discarded.

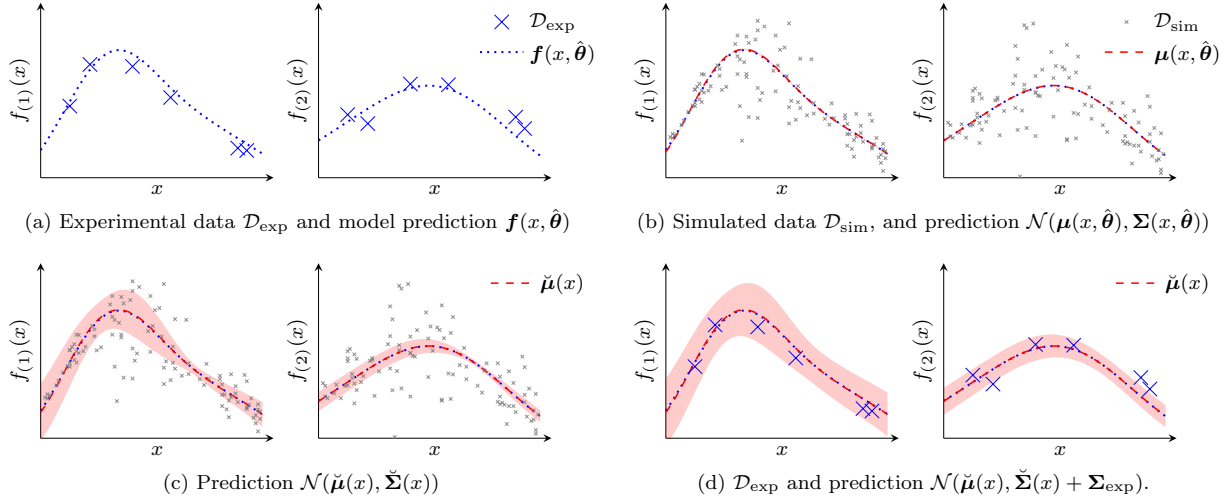


Figure 2. Process of constructing the approximation of the marginal predictive distribution to design a new experiment and perform model discrimination. (a) Step 1, use experimental data \mathcal{D}_{exp} to perform model parameter estimation and find $\hat{\theta}$. (b) Step 2, generate simulated data \mathcal{D}_{sim} by sampling from the model f . Use \mathcal{D}_{sim} to learn hyperparameters of the GP surrogate model. (c) Step 3, use Taylor approximation to compute $\check{\mu}$ and $\check{\Sigma}$. Use approximation of marginal predictive distribution to design a new experiment. (d) Step 4, use the approximation of the predictive distribution to compute the model likelihood and perform model discrimination.

(I) the proportion of inconclusive tests (all models were deemed inaccurate or the maximum number of additional experiments $N_{\text{max}} - N_0$ was reached).

We compare the design criteria (DC) D_{BH} , D_{BF} and D_{AW} described in Section 2.1 as well as random uniform sampling (denoted Uni.). We also compare the three different criteria for model discrimination (MD) described in Section 2.1: updated posterior likelihood $\pi_{N,i}$, the χ^2 adequacy test, and the Akaike information criterion (AIC).

4.1. Case Study 1: Comparison to Analytical Method

Case study 1 is from the seminal paper by Buzzi-Ferraris et al. (1984). With two measured states y_1, y_2 and two design variables $x_1, x_2 \in [5, 55]$, we discriminate between four chemical kinetic models \mathcal{M}_i , each with four model parameters $\theta_{i,j}$ that we fix to $\theta_{i,j} \in [0, 1]$:

$$\begin{aligned} \mathcal{M}_1: & f_{(1)} = \theta_{1,1}x_1x_2/g_1, & f_{(2)} = \theta_{1,2}x_1x_2/g_1, \\ \mathcal{M}_2: & f_{(1)} = \theta_{2,1}x_1x_2/g_2^2, & f_{(2)} = \theta_{2,2}x_1x_2/h_{2,1}^2, \\ \mathcal{M}_3: & f_{(1)} = \theta_{3,1}x_1x_2/h_{3,1}^2, & f_{(2)} = \theta_{3,2}x_1x_2/h_{3,2}^2, \\ \mathcal{M}_4: & f_{(1)} = \theta_{4,1}x_1x_2/g_4, & f_{(2)} = \theta_{4,2}x_1x_2/h_{4,1}, \end{aligned}$$

where $g_i = 1 + \theta_{i,3}x_1 + \theta_{i,4}x_2$ and $h_{i,j} = 1 + \theta_{i,2+j}x_j$. Buzzi-Ferraris et al. (1984) generate the experimental data \mathcal{D}_{exp} from model \mathcal{M}_1 using $\theta_{1,1} = \theta_{1,3} = 0.1$ and $\theta_{1,2} = \theta_{1,4} = 0.01$ and Gaussian noise covariance $\Sigma_{\text{exp}} = \text{diag}(0.35, 2.3e-3)$. We start each test with $N_0 = 5$ randomly sampled experimental observations, and set a maximum budget of $N_{\text{max}} - N_0 = 40$ additional experiments.

Tables 1–3 compare the performance in 500 completed test instances of the classical analytical approach to the GP surrogate method using first- or second-order Taylor approximations. For this case study, the GP surrogate method performs similarly to the classical analytical approach. We see that the average number of required additional experiments (A) increases using π for model discrimination, but that the model discrimination success rate (S) also increases. The model discrimination failure rate drops for almost all MD/DC combinations using the GP surrogate method. We also see that even for this simple case study, the difference in performance between the design criteria and random sampling is significant. The low average (A) for $\pi/\text{Uni.}$ in Table 1 is likely an artefact due to the large failure rate (F).

4.2. Case Study 2: Non-Analytical Models

Case study 2 is from Vanlier et al. (2014). There are four different models, described by ordinary differential equations. The supplementary material contains the full model descriptions. The models describe a biochemical network, and are similar to typical pharmacokinetic models in (i) the use of reversible reactions, and (ii) the number of parameters and equations (see e.g. Han et al. (2006)).

We assume two measured states ($E = 2$), with model parameters $\theta_i \in [0, 1]^{10}$, initial concentrations $\mathbf{c} \in [0, 1]^E$, measurement time point $t \in [0, 20]$, and design variables $\mathbf{x}^\top = [t, \mathbf{c}^\top]$. We follow Vanlier et al. (2014) and let \mathcal{M}_1 generate the observed experimental data, with random uniformly sampled true model parameters and experimental




Figure 3. Results from case study 2 with 4 models using the D_{AW} design criterion and AIC for model discrimination. (a)–(d) show the evolution of the Akaike weights w_i for model $\mathcal{M}_1, \dots, \mathcal{M}_4$, respectively, for 63 completed tests. (e) shows the average Akaike weight evolutions (with one standard deviation) in (a)–(d). The results indicate that models \mathcal{M}_1 and \mathcal{M}_2 are indiscriminable.

measurement noise covariance $\Sigma_{\text{exp}} = 9.0 \times 10^{-4} \mathbf{I}$. Each test is initiated with $N_0 = 20$ observations at random design locations $\mathbf{x}_1, \dots, \mathbf{x}_{N_0}$. We set a maximum budget of $N_{\text{max}} - N_0 = 100$ additional experiments.

Case study 1 shows no obvious benefit in using the second-order over a first-order Taylor approximation. Hence, in case study 2 we only performance test the GP surrogate method with the first-order Taylor approximation.

Table 1. Results from case study 1 using the analytical approach.

MD	π	χ^2	AIC	π	χ^2
DC	D_{BH}	D_{BF}	D_{AW}	Uni.	Uni.
A	2.60	2.87	2.08	3.43	10.39
SE	0.04	0.12	0.04	0.11	0.45
S [%]	86.4	64.2	62.4	59.0	85.8
F [%]	13.6	5.0	37.6	40.6	4.4
I [%]	0.0	30.8	0.0	0.4	9.8

Table 2. Results from case study 1 using the GP surrogate method with first-order Taylor approximation.

MD	π	χ^2	AIC	π	χ^2
DC	D_{BH}	D_{BF}	D_{AW}	Uni.	Uni.
A	4.31	2.23	2.72	10.99	10.02
SE	0.09	0.06	0.08	0.34	0.45
S [%]	95.6	47.4	88.6	69.6	83.6
F [%]	4.4	4.8	11.4	30.4	4.8
I [%]	0.0	47.8	0.0	0.0	11.6

Table 3. Results from case study 1 using the GP surrogate method with second-order Taylor approximation.

MD	π	χ^2	AIC	π	χ^2
DC	D_{BH}	D_{BF}	D_{AW}	Uni.	Uni.
A	4.14	2.29	2.64	9.53	10.11
SE	0.09	0.07	0.06	0.30	0.47
S [%]	96.9	46.6	90.1	64.3	80.2
F [%]	3.1	9.9	9.9	35.7	7.0
I [%]	0.0	43.5	0.0	0.0	12.8

The results in Table 4 for 4 models show that the rates of inconclusive results are higher for case study 2 than for case study 1. Furthermore, the ratio of successful model discriminations to failed model discriminations are lower for π/D_{BH} and AIC/D_{AW} than they are in case study 1. This is because for case study 2 the experimental noise is often larger than the difference between the model predictions. Figure 3 shows the evolutions of the Akaike weights w_i (Michalik et al., 2010) for all 63 completed tests. The Akaike weight evolutions show that especially models \mathcal{M}_1 and \mathcal{M}_2 are difficult to discriminate between, given that we only consider two measured states. Corresponding graphs for π and χ^2 are similar-looking. One can also see from the model formulations (see supplementary material) that all models are similar.

To verify that the Table 4 results for 4 models are indeed due to the indiscriminability of model \mathcal{M}_1 and \mathcal{M}_2 , we carry out another set of tests where we only consider models \mathcal{M}_1 , \mathcal{M}_3 and \mathcal{M}_4 . The results in Table 4 for 3 models show that removing model \mathcal{M}_2 enables the GP surrogate method to correctly discriminate between the remaining models more successfully. In practice, researchers faced with apparent model indiscriminability have to rethink their experimental set-up to lower measurement noise or add measured states.

Table 4. Results from case study 2.

MD	4 models			3 models		
	π	χ^2	AIC	π	χ^2	AIC
DC	D_{BH}	D_{BF}	D_{AW}	D_{BH}	D_{BF}	D_{AW}
A	20.10	39.83	29.62	15.80	21.91	9.74
SE	3.72	12.09	7.72	2.05	2.52	1.70
S [%]	15.9	9.5	33.3	89.5	77.2	95.6
F [%]	7.9	0.0	7.9	6.1	0.9	1.8
I [%]	76.2	90.5	58.7	4.4	21.9	2.6

5. Discussion

The healthcare industry already uses machine learning to help discover new drugs (Pfizer, 2016). But passing the pre-clinical stage is expensive. DiMasi et al. (2016) estimate the average pre-approval R&D cost to \$2.56B (in 2013 U.S. dollars) with a yearly increase of 8.5 %. For the pre-clinical stage (before human testing starts), DiMasi et al. (2016) estimate the R&D cost to \$1.1B.

Model discrimination, which takes place in the pre-clinical stage, may lower drug development cost by identifying a model structure quickly. Heller et al. (2018) explain that a major hurdle for *in silico* pharmacokinetics is the difficulty of model discrimination. Scannell & Bosley (2016) and Plenge (2016) also attribute a significant part of the total R&D cost for new drugs to inaccurate models passing the pre-clinical stage only to be discarded in later clinical stages. Inaccurate models need to be discarded sooner rather than later to avoid expensive failures.

Design of experiments for model discrimination has typically focussed either on analytical approaches, which cannot manage all functions, or on data-driven approaches, which may have computational difficulties or lack interpretable marginal predictive distributions. We leverage advantages of both approaches by hybridising the analytical and data-driven approaches, i.e. we replace the analytical functions from the classical methods with GP surrogates.

The novel GP surrogate method is highly effective. Case study 1 directly compares it with the analytical approach on a classical test instance. The similar performance indicates successful hybridisation between analytical and data-driven approaches. Case study 2 considers an industrially relevant problem suffering from model indiscriminability. The GP surrogate method successfully avoids introducing overconfidence into the model discrimination despite the assumptions used to derive the marginal predictive distribution.

The parameter estimation step is limited to few and noisy data points, e.g. from time-consuming pharmacokinetic experiments. If the approximation $p(\boldsymbol{\theta} | \mathcal{D}_{\text{exp}}) \approx \mathcal{N}(\hat{\boldsymbol{\theta}}, \boldsymbol{\Sigma}_{\theta})$ is not accurate, the Section 3.1 Taylor approximations of $\check{\boldsymbol{\mu}}$ and $\check{\boldsymbol{\Sigma}}$ can be extended to a mixture of Gaussian distributions $p(\boldsymbol{\theta} | \mathcal{D}_{\text{exp}}) = \sum_j \omega_j \mathcal{N}(\hat{\boldsymbol{\theta}}_j, \boldsymbol{\Sigma}_{\theta,j})$, with $\sum_j \omega_j = 1$, if $\{\hat{\boldsymbol{\theta}}_j, \boldsymbol{\Sigma}_{\theta,j}\}$ are given. The analytical approximations sacrifice some accuracy for computational tractability: Our computational time for design of experiments in case study 2 was on the order of minutes, whereas the method of Vanlier et al. (2014) requires days. Our results indicate that the approximations work well in practice. GP surrogate methods could sharply reduce the cost of drug development’s pre-clinical stage, e.g. by shrinking the time needed to decide on a new experiment or by finding a good model structure in a reduced number of pharmacokinetic experiments.

Our new method directly applies to neighbouring fields. For example, dual control theory deals with control of unknown systems, where the goal is to balance controlling the system optimally vs. learning more about the system. When rival models predict the unknown system behaviour, learning about the system involves designing new experiments, i.e. control inputs to the system (Cheong & Manchester, 2014). This design of experiments problem, e.g. in Larsson et al. (2013), would be likely to have very strict time requirements and benefit from confidence bounds on the model predictions. Our method also applies to neighbouring application areas, e.g. distinguishing between microalgae metabolism models (Ulmasov et al., 2016) or models of bone neotissue growth (Mehrian et al., 2018).

A very hot topic in machine learning research is learning in implicit generative models (Mohamed & Lakshminarayanan, 2017), e.g. generative adversarial networks (GANs) (Goodfellow et al., 2014), which employ a form of model discrimination. Here the likelihood function $p(\mathbf{y} | \mathcal{M}_i)$ is not given explicitly, so the goal is to learn an approximation $q(\mathbf{y} | \mathcal{M}_i)$ of the data distribution $p(\mathbf{y})$. The model discrimination tries to identify from which distribution $q(\mathbf{y} | \mathcal{M}_i)$ or $p(\mathbf{y})$ a random sample has been drawn. For the case of design of experiments for model discrimination, we have a prescribed probabilistic model, since we explicitly specify an approximation of the likelihood $p(\mathbf{y} | \mathcal{M}_i, \mathbf{x}, \boldsymbol{\theta}_i)$. But we have (i) the physical experiments, the true but noisy data generator that is expensive to sample, and (ii) the rival models, deterministic and relatively cheap (but probably inaccurate!) data generators. Design of experiments for model discrimination also ties in with several other well-known problems in machine learning, e.g. density estimation, distribution divergence estimation, and model likelihood estimation.

6. Conclusions

We have presented a novel method of performing design of experiments for model discrimination. The method hybridises analytical and data-driven methods in order to exploit classical results while remaining agnostic to the model contents. Our method of approximating the marginal predictive distribution performs similarly to classical methods in a case study with simple analytical models. For an industrially-relevant system of ordinary differential equations, we show that our method both successfully discriminates models and avoids overconfidence in the face of model indiscriminability. The trade-off between computational complexity and accuracy that we introduce is useful for design of experiments for discrimination e.g. between pharmacokinetic models, where the duration of the experiments is shorter than or one the same time-scale as the computational time for Monte Carlo-based methods.

Acknowledgements

This work has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement no.675251, and an EPSRC Research Fellowship to R.M. (EP/P016871/1).

References

- Ambrogioni, L. and Maris, E. Analysis of nonstationary time series using locally coupled gaussian processes. *ArXiv e-prints*, 2016.
- Boukouvala, F., Misener, R., and Floudas, C. A. Global optimization advances in mixed-integer nonlinear programming, MINLP, and constrained derivative-free optimization, CDO. *Eur J Oper Res*, 252(3):701 – 727, 2016.
- Box, G. E. P. and Hill, W. J. Discrimination among mechanistic models. *Technometrics*, 9(1):57–71, 1967.
- Buzzi-Ferraris, G., Forzatti, P., Emig, G., and Hofmann, H. Sequential experimental design for model discrimination in the case of multiple responses. *Chem Eng Sci*, 39(1): 81–85, 1984.
- Buzzi-Ferraris, G., Forzatti, P., and Canu, P. An improved version of a sequential design criterion for discriminating among rival multiresponse models. *Chem Eng Sci*, 45(2): 477–481, 1990.
- Chaloner, K. and Verdinelli, I. Bayesian experimental design: A review. *Statist Sci*, 10(3):273–304, 1995.
- Cheong, S. and Manchester, I. R. Model predictive control combined with model discrimination and fault detection. In *IFAC 19*, 2014.
- DiMasi, J. A., Grabowski, H. G., and Hansen, R. W. Innovation in the pharmaceutical industry: New estimates of R&D costs. *J Health Econ*, 47:20–33, 2016.
- Drovandi, C. C., McGree, J. M., and Pettitt, A. N. A sequential Monte Carlo algorithm to incorporate model uncertainty in Bayesian sequential design. *J Comput Graph Stat*, 23(1):3–24, 2014.
- Ette, E. I. and Williams, P. J. Population pharmacokinetics i: Background, concepts, and models. *Ann Pharmacother*, 38:1702–1706, 2004.
- European Medicines Agency. Guideline on the use of pharmacogenetic methodologies in the pharmacokinetic evaluation of medicinal products, December 2011.
- European Medicines Agency. Guideline on the qualification and reporting of physiologically based pharmacokinetic (PBPK) modelling and simulation, July 2016.
- European Union. Regulation (EU) 2017/745, April 2017.
- Field, R. I. Why is health care regulation so complex? *Pharm Ther*, 33(10):607–608, 2008.
- Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *NIPS 27*, 2014.
- GPdoemd. Github repository. <https://github.com/cog-imperial/GPdoemd>, 2018.
- Han, J.-J., Plawecki, M. H., Doerschuk, P. C., Ramchandani, V. A., and O'Connor, S. Ordinary differential equation models for ethanol pharmacokinetic based on anatomy and physiology. In *EMBS 28*, 2006.
- Heller, A. A., Lockwood, S. Y., Janes, T. M., and Spence, D. M. Technologies for measuring pharmacokinetic profiles. *Annu Rev Anal Chem*, 11(1), 2018.
- Hervey, T. The impacts of European Union law on the health care sector: Institutional overview. *Eurohealth*, 16(4):5–7, 2010.
- Larsson, C. A., Annergren, M., Hjalmarsson, H., Rojas, C. R., Bombois, X., Mesbah, A., and Modén, P. E. Model predictive control with integrated experiment design for output error systems. In *ECC 2013*, pp. 3790–3795, 2013.
- Li, T., Sun, S., Sattar, T. P., and Corchado, J. M. Fight sample degeneracy and impoverishment in particle filters: A review of intelligent approaches. *Expert Syst Appl*, 41(8):3944–3954, 2014.
- MacKay, D. J. C. Bayesian interpolation. *Neural Comput*, 4:415–447, 1992a.
- MacKay, D. J. C. Information-based objective functions for active data selection. *Neural Comput*, 4:590–604, 1992b.
- Mehrian, M., Guyot, Y., Papantoniou, I., Olofsson, S., Sonnaert, M., Misener, R., and Geris, L. Maximizing neotissue growth kinetics in a perfusion bioreactor: An in silico strategy using model reduction and Bayesian optimization. *Biotechnol Bioeng*, 115(3):617–629, 2018.
- Michalik, C., Stuckert, M., and Marquardt, W. Optimal experimental design for discriminating numerous model candidates: The AWDC criterion. *Ind Eng Chem Res*, 49: 913–919, 2010.
- Mohamed, S. and Lakshminarayanan, B. Learning in implicit generative models. *ArXiv e-prints*, 2017.
- Ogungbenro, K. and Aarons, L. How many subjects are necessary for population pharmacokinetic experiments? Confidence interval approach. *Eur J Clin Pharmacol*, 64(7):705, 2008.

- Pfizer. IBM and Pfizer to accelerate immuno-oncology research with Watson for drug discovery, 1 December 2016. Press release.
- Plenge, R. M. Disciplined approach to drug discovery and early development. *Sci Transl Med*, 8(349):349ps15, 2016.
- Prasad, K. B. S. and Someswara Rao, M. Use of expected likelihood in sequential model discrimination in multiresponse systems. *Chem Eng Sci*, 32:1411–1418, 1977.
- Rasmussen, C. E. and Williams, C. K. I. *Gaussian processes for machine learning*. MIT Press, 2006.
- Ryan, E. G., Drovandi, C. C., and Pettitt, A. N. Fully Bayesian experimental design for pharmacokinetic studies. *Entropy*, 17:1063–1089, 2015.
- Ryan, E. G., Drovandi, C. C., McGree, J. M., and Pettitt, A. N. A review of modern computational algorithms for Bayesian optimal design. *Int Stat Rev*, 84:128–154, 2016.
- Scannell, J. W. and Bosley, J. When quality beats quantity: Decision theory, drug discovery, and the reproducibility crisis. *PLOS ONE*, 11(2):1–21, 2016.
- Steimer, J.-L., Mallet, A., Golmard, J.-L., and Boisvieux, J.-F. Alternative approaches to estimation of population pharmacokinetic parameters: Comparison with the nonlinear mixed-effect model. *Drug Metab Rev*, 15(1–2): 265–292, 1984.
- Tsimring, L. S. Noise in biology. *Rep Prog Phys*, 77(2): 026601, 2014.
- Tuntland, T., Ethell, B., Kosaka, T., Blasco, F., Zang, R. X., Jain, M., Gould, T., and Hoffmaster, K. Implementation of pharmacokinetic and pharmacodynamic strategies in early research phases of drug discovery and development at Novartis Institute of Biomedical Research. *Front Pharmacol*, 5:174, 2014.
- Ulmasov, D., Baroukh, C., Chachuat, B., Deisenroth, M. P., and Misener, R. Bayesian optimization with dimension scheduling: Application to biological systems. In Kravanja, Z. and Bogataj, M. (eds.), *26th European Symposium on Computer Aided Process Engineering*, volume 38 of *Computer Aided Chemical Engineering*, pp. 1051 – 1056. 2016.
- U.S. Food & Drug Administration. Food and drugs, 21 C.F.R. § 314, 25 January 2018.
- Vanlier, J., Tiemann, C. A., Hilbers, P. A. J., and van Riel, N. A. W. Optimal experiment design for model selection in biochemical networks. *BMC Syst Biol*, 8(20), 2014.
- Woods, D. C., Overstall, A. M., Adamou, M., and Waite, T. W. Bayesian design of experiments for generalized linear models and dimensional analysis with industrial and scientific application. *Qual Eng*, 29(1):91–103, 2017.