Optimal experiment design with applications to Pharmacokinetic modeling*

Murat K. Erdal[1], Kevin W. Plaxco[1], Julian Gerson[1], Tod E. Kippin[1], João P. Hespanha [1]

***Abstract*— We study the problem of designing an input to a dynamical system that is optimal at estimating unknown parameters in the system's model. We take the A and D optimality criteria on the Fisher Information Matrix associated with the estimation problem as our optimization objective. Our main motivation is the estimation of the physiological parameters that appear in pharmacokinetic dynamics using a relatively short set of measurements. In this context, model inputs correspond to the intravenous injection of drugs and input selection needs to consider safety constraints that include max-min instantaneous injection rates and total dosage amount. We divide the time interval available for the experiment into learning and optimization stages. We use the initial learning stage to obtain a preliminary estimate for the system's model. Then we find an optimal input for the optimization stage so that we can improve upon this initial estimate.**

I. INTRODUCTION

Learning system dynamics from measured data can be a significant challenge when long time series of measurements are difficult to obtain. This challenge is especially important in biology-related fields such as pharmacokinetics, where subject-to-subject variability is large and time-series of measurements can be difficult to collect. In such systems, it is crucial to maximize the information content of the limited data that can be collected.

Pharmacokinetics is the study of how drugs diffuse within and are removed from the body [1] and pharmacokinetics provide a model for this process. Even for the same drug at the same dose, these dynamics typically exhibit significant variability across individuals. Because of this, the development of new drugs [2] and/or the design of individualized treatment regimens [3], [4] require understanding the individualized pharmacokinetics based on a single or a small number of identification experiments. This has historically proven difficult, as the time resolution of the traditional methods of measuring drug concentration is quite poor. The recently developed Electrochemical-aptamer-based (E-AB) sensor platform, however, supports seconds-resolution in these measurements, in turn providing an unprecedentedly high-precision window into pharmacokinetics [4].

Motivated by these recent developments, we are interested in scenarios where a single experiment is performed to estimate drug-specific pharmacokinetic parameters of an animal, and relatively little information is known about those parameters before the experiment starts. Such experiments typically consist of administering a drug and measuring the time evolution of the drug concentration in one or several compartments in the body. The experiment design problem consists of selecting a time-profile for the drug administration to obtain the best possible estimates for the pharmacokinetic parameters while respecting key safety constraints that are typical drug-injection protocols.

We use the Fisher Information Matrix (FIM) to assess the information content that can be extracted from the measurements collected in an experiment. We recall that the inverse of the FIM is equal to the error covariance of an optimal estimator that is able to achieve the Cramér–Rao bound and therefore a FIM with a "small" inverse is associated with an estimation problem that can be solved with small mean-square errors. The FIM can thus be viewed as a measure of information content [5]. We consider both the A and D scalar criteria to assess optimality on FIM, which are common in the experiment design literature [6].

It is well known that the exact value of the FIM depends on the true system. This issue has been even called the "Achilles' heel of optimal input design" in [7], as it does not really make sense to do system identification if we know the system in the first place. This problem usually resolved with some prior information taken from previous experiments or literature. However, especially for drugs with large subject-to-subject variability, little prior information may be available on the pharmacokinetic parameter values. For such problems, we propose a two-stage approach for optimal experiment design that uses an initial "learning" stage to build coarse parameter estimates, followed by an "optimization" stage that injects the drug according to a time profile that has been optimized for the specific individual.

Optimal experiment design problem enjoys a rich history, and it can be posed for different purposes such as system identification or adaptive control. These differing objectives change the criterion by which we call an input profile optimal. We focus on experiment design for system identification but we refer to [8], [9] for a detailed survey of differing objectives.

The complexity of finding solutions to experiment design problems is highly dependent on the specifics of the problem in the hand. For example, it can be posed as a convex program by using linear matrix inequalities (LMI) when the input is optimized over a spectrum [7], [10]. Unfortunately, optimizing over the spectrum prevents us from imposing most time-domain constraints, which are required for safety in pharmacokinetic studies.

It was shown that solving experiment design problems

*This work was supported by National Institutes of Health under Grant R01 A1145206 and National Science Foundation under Grant ECCS 2029985

[1]University of California, Santa Barbara, Santa Barbara, CA 93106, The USA. Correspondence: m_erdal@ucsb.edu

with optimality criterion based on the FIM with time-domain amplitude constraints is generally NP-hard [11]. Convex semidefinite relaxations have been used to deal with this complexity issue. However, the solution to the relaxed problem is usually not compatible with the original feasible space. This requires a projection back to the original feasible space from the relaxed solutions. Examples of such projections are discussed in [11], [12]. Similar relaxation techniques were also used to solve the optimal experiment design problem, together with a constraint on the sum of the input in [13].

We propose to solve the original experiment design problem using a software tool, CasADI [14], that employs automatic differentiation together with an nonlinear program solver, IPOPT [15], without having to relax the problem or projecting back to the feasible space. This approach allows us to execute the optimization procedure fast and reliably, and thereby similar techniques recently attracted attention [16].

We formally introduce the optimal experiment design problem in Section II for a generic discrete-time linear time invariant (LTI) system. The specific pharmacokinetic model and the relevant safety constraints are introduced in Section III. We introduce our two-stage approach that involves the initial identification and the subsequent optimization stages in Section IV. We, then, validate our approach on a hypothetical drug that follows a three compartment model, and discuss the numerical results in Section V.

II. Optimal Experiment Design

We consider a parametrized discrete linear system model of the form

$$x(t+1) = A(\theta)x(t) + B(\theta)u(t) \tag{1a}$$
$$y(t) = C(\theta)x(t) + D(\theta)u(t) + \eta(t) \tag{1b}$$

$\forall t \in \{1, \dots, N\}$, where $\theta \in \mathbb{R}^{n_\theta}$ is a vector of unknown parameters, $x(t) \in \mathbb{R}^{n_x}$ is the state of the system, $u(t) \in \mathbb{R}^{n_u}$ a known input, $y(t) \in \mathbb{R}^{n_y}$ the measured output, and $\eta(t) \in \mathbb{R}^{n_y}$ stochastic measurement noise. When the initial condition $x(1)$ is unknown, we include it in the vector θ.

A. Fisher Information Matrix

The Fisher Information Matrix (FIM) associated with the estimation of the parameter θ is denoted by $I_\theta \in \mathbb{R}^{n_\theta \times n_\theta}$ and is defined as a positive definite matrix whose $(i,j)^{th}$ entry is equal to

$$I_{\theta_{ij}}(\bar{u}) = \mathbb{E}\left[\frac{\partial \log p(\bar{y};\bar{u},\theta)}{\partial \theta_i} \frac{\partial \log p(\bar{y};\bar{u},\theta)}{\partial \theta_j} \middle| \theta \right], \tag{2}$$

where $\bar{y} \in \mathbb{R}^{Nn_y}$ and $\bar{u} \in \mathbb{R}^{Nn_u}$ are the observations and the inputs vector, defined as

$$\bar{y} = \begin{bmatrix} y(1)^T & y(2)^T & \dots y(N)^T \end{bmatrix}^T$$
$$\bar{u} = \begin{bmatrix} u(1)^T & u(2)^T \dots & u(N)^T \end{bmatrix}^T,$$

and $p(\bar{y};\bar{u},\theta)$ is the probability density function (pdf) of the measurements in $\bar{y}$ given the input sequence $\bar{u}$ and the parameter vector θ.

The following result provides an explicit formula for the FIM associated with the parameter θ in (1), based on the matrix

$$M(\theta) = \begin{bmatrix} Cx(1) & D & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ CA^{N-1}x(1) & CA^{N-2}B & \dots & D \end{bmatrix}, \tag{3}$$

where we omitted from the right-hand side the dependence of A, B, C, D, and $x(1)$ on the parameter vector θ.

Theorem 1: Assume that the measurement noise vectors $\eta(t)$ in the LTI system (1) are identically independently distributed zero-mean Gaussian random vectors with covariance matrix equal to Σ_n. Then the entries of the FIM are quadratic in the input $\bar{u}$ with

$$I_{\theta_{ij}}(\bar{u}) = \bar{u}_{agg}^T M(\theta,i,j)\bar{u}_{agg}, \tag{4}$$

where $\bar{u}_{agg} = \begin{bmatrix} 1 & \bar{u}^T \end{bmatrix}^T$ and

$$M(\theta,i,j) = \frac{\partial M(\theta)}{\partial \theta_i}^T \left(I \otimes \Sigma_n^{-1}\right) \frac{\partial M(\theta)}{\partial \theta_j},$$

Proof: Because the vectors $\eta(t)$ are identically independently distributed zero-mean Gaussian random vectors with a covariance matrix Σ_n, each measurement $y(t)$ is also a Gaussian random vector with a covariance of Σ_n. Moreover, the expected value of the output of an LTI system (1) at the time t is given by the variation of constants formula [17].

$$\mathbb{E}\left[y(t);\bar{u},\theta\right] = CA^t x(1) + C\sum_{\ell=1}^{t-1} A^{t-\ell-1}Bu(\ell) + Du(t),$$

where the matrices A, B, C, D are all functions of parameter vector θ but we drop the θ notation in these matrices for brevity. We can stack each of these expected outputs and express this computation in vector form

$$\mu_{\bar{y}} := \mathbb{E}\left[\bar{y};\bar{u},\theta\right] = M(\theta) \underbrace{\begin{bmatrix} 1 \\ \bar{u} \end{bmatrix}}_{\bar{u}_{agg}}. \tag{5}$$

Now that we also have an explicit formula for the expected value of $\bar{y}$, we can express the log likelihood of $\bar{y}$ as,

$$\begin{aligned} \log p\left(\bar{y};\bar{u},\theta\right) = &-\frac{n_y N}{2}\log(2\pi) + \frac{1}{2}\log\det\left(\bar{\Sigma}^{-1}\right) \\ &- \frac{1}{2}\left(\bar{y}-\mu_{\bar{y}}\right)^T \bar{\Sigma}^{-1}\left(\bar{y}-\mu_{\bar{y}}\right), \end{aligned} \tag{6}$$

where $\bar{\Sigma} = (I_N \otimes \Sigma_n)$ We take the derivative with respect to the i^{th} entry of θ,

$$\begin{aligned} \frac{\partial \log p\left(\bar{y};\bar{u},\theta\right)}{\partial \theta_i} &= \frac{\partial \mu_{\bar{y}}}{\partial \theta_i}^T \bar{\Sigma}^{-1}\left(\bar{y}-\mu_{\bar{y}}\right) \\ &= \bar{u}_{agg}^T \frac{\partial M(\theta)}{\partial \theta_i}^T \bar{\Sigma}^{-1}\left(\bar{y}-\mu_{\bar{y}}\right). \end{aligned}$$

Then we take the expected value as in (2), which leads to,

$$\begin{aligned} I_{\theta_{ij}}(\bar{u}) &= \mathbb{E}\left[\frac{\partial \log p(\bar{y};\bar{u},\theta)}{\partial \theta_i} \frac{\partial \log p(\bar{y};\bar{u},\theta)}{\partial \theta_j} \middle| \theta \right] \\ &= \frac{\partial \mu_{\bar{y}}}{\partial \theta_i}^T \bar{\Sigma}^{-1} \frac{\partial \mu_{\bar{y}}}{\partial \theta_j} \\ &= \bar{u}_{agg}^T M(\theta,i,j)\bar{u}_{agg}, \end{aligned} \tag{7}$$

where $M(\theta,i,j) = \frac{\partial M(\theta)}{\partial \theta_i}^T \left(I \otimes \Sigma_n^{-1}\right) \frac{\partial M(\theta)}{\partial \theta_j}$. ■

B. Optimization criteria

The A and D optimality criteria are commonly used to construct a scalar metric that captures the information content of a set of measurements and are formally defined by

$$f_A(I_\theta) = \mathrm{Tr}\left(I_\theta^{-1}\right), \qquad f_D(I_\theta) = -\log\det(I_\theta), \tag{8}$$

respectively [6]. For Gaussian measurement vectors, the A optimality criterion is proportional to the mean square estimation error (or the sum of the square of the estimation errors associated with each entry of θ) of an optimal estimator, whereas the D optimality criteria is proportional to the volume of the confidence region around the estimates defined by level sets of the posterior pdf of the parameters given the measurements.

Since the criteria in (8) generally depend on the value of the unknown parameter θ, we need to minimize the expected value of one of these criteria with respect the prior distribution of θ, leading to the following optimization

$$\bar{u}^* = \underset{\bar{u}\in\mathcal{U}}{\arg\min}\, \mathbb{E}_\theta\Big[f\big(I_\theta(\bar{u})\big)\Big], \tag{9}$$

where $f(\cdot)$ is either $f_A(\cdot)$ or $f_D(\cdot)$ in (8), $\mathbb{E}_\theta$ refers to the expectation taken over the prior distribution of θ, and $\mathcal{U}$ is a set of admissible control inputs.

In general, the expected value in the right-hand side of (9) cannot be computed in closed form and needs to be replaced by an empirical average, leading to the following sample average approximation (SAA) [18]:

$$\bar{u}^* = \underset{\bar{u}\in\mathcal{U}}{\arg\min}\, \frac{1}{K}\sum_{k=1}^{K} f\big(I_{\theta_k}(\bar{u})\big), \tag{10}$$

where the θ_k are independent random samples extracted from the prior distribution for θ. Since minimizing the criteria in (10) is computationally difficult for a large number of samples K, we take advantage of the fact that both optimality criteria functions in (8) are convex with respect to the cone of positive definite matrices [19] and use Jensen's inequality to conclude that

$$\frac{1}{K}\sum_{k=1}^{K} f\big(I_{\theta_k}(\bar{u})\big) \geqslant f\Big(\frac{1}{K}\sum_{k=1}^{K} I_{\theta_k}(\bar{u})\Big). \tag{11}$$

Rather than minimizing the empirical average in the left-hand side of this inequality, we minimize instead its lower bound in the right-hand side.

$$\bar{u}^* = \underset{\bar{u}\in\mathcal{U}}{\arg\min}\, f\left(\hat{I}(\bar{u})\right), \quad \hat{I}(\bar{u}) \coloneqq \frac{1}{K}\sum_{k=1}^{K} I_{\theta_k}(\bar{u}). \tag{12}$$

This optimization is much more attractive, because we can conclude from (4) in Theorem 1 that the $(i,j)^{th}$ entry of $\bar{I}$ is given by

$$\hat{I}_{ij}(\bar{u}) = \bar{u}_{agg}^T\Big(\frac{1}{K}\sum_{k=1}^{K} M(\theta_k,i,j)\Big)\bar{u}_{agg}, \tag{13}$$

which shows that the average $\frac{1}{K}\sum_{k=1}^{K} M(\theta_k,i,j)$ does not depend on the control input $\bar{u}$ and can be pre-computed without adding complexity to the optimization in (12).

We acknowledge that minimizing a lower-bound is, in general, not ideal since it may not lead to a smaller real objective. However, our results, presented in Table II, for a generic numerical example indicate that the lower bound in (11) indeed leads to more accurate estimates.

III. Pharmacokinetic Model

We use the discretized version of a classical compartmental model [20] with n_x compartments and n_u inputs:

$$\begin{aligned} x_i(t+1) = {} & (1-k_{E_i})\,x_i(t) \\ & + \sum_{\substack{j=1\\ j\neq i}}^{n_x} k_{j_i}\left(x_j(t)-x_i(t)\right) + \sum_{\ell=1}^{n_u} k_{U_{i\ell}}u_\ell(t), \end{aligned} \tag{14}$$

where x_i is the drug concentration in compartment i, k_{E_i} is the elimination rate in compartment i, k_{j_i} is the rate constant associated with the transfer from the j^{th} into i^{th} compartments, $u_\ell(t)$ is the ℓ^{th} exogenous injection at time t, and $k_{U_{i\ell}}$ is the transfer rate for u_ℓ into the i^{th} compartment.

In this model, we have $k_{j_i} = 0$ when the concentration of compartment j does not affect the concentration in compartment i, $k_{E_i} = 0$ when the drug is not being metabolized nor excreted from compartment i, and $k_{U_{i\ell}} = 0$ when there is no drug infusion into compartment i from ℓ^{th} infusion channel.

We also assume that not all these compartments are available to be observed. Instead, we assume that there are n_y sensors each of which measuring a different compartment with some noise in the measured values.

We can express the dynamics (14) in terms of the parametrized discrete LTI system in (1) with $D = \bar{0} \in \mathbb{R}^{n_y\times n_u}$ and appropriate choices for the matrices A, B, C: The non-diagonal entries of $A \in \mathbb{R}^{n_x\times n_x}$ are given by $A_{ij} = k_{j_i}$. The diagonal entries of A are given by

$$A_{ii} = -k_{E_i} - \sum_{\substack{j=1\\ j\neq i}}^{n} k_{j_i}.$$

The entries of the matrix $B \in \mathbb{R}^{n_x\times n_u}$ are given by $B_{i\ell} = k_{U_{i\ell}}$. By assuming that specific entries of the matrices A and B are zero, we are implicitly introducing some knowledge regarding the underlying network of compartments for the particular drug.

Finally, each row i of the matrix $C \in \mathbb{R}^{n_y\times n_x}$ has a single entry at the entry (i,j) where j corresponds to the compartment whose concentration is measured by the ith sensor.

The vector $\theta \in \mathbb{R}^{n_\theta}$ of unknown parameters contains all the pharmacokinetic parameters, including all the nonzero entries of the matrix B, the nonzero non-diagonal entries of A, and the nonzero elimination rates in diagonal entries of A. When the initial condition of the system is not known, it needs to be added to the parameters vector.. We define the feasible set $\mathcal{U}$ in (12) next.

Pharmacokinetic Constraints

When conducting experiments to collect pharmacokinetic measurements from living animals one needs to follow pre-approved safety procedures that specify maximum and minimum instantaneous injection rates and maximum dosage over a certain period of time. Therefore, the set $\mathcal{U}$ of admissible inputs in the problem (12) needs to be restricted to the input profiles that satisfy

$$u_{\min} \leqslant u(t) \leqslant u_{\max}, \ \forall t \in \{1, \dots, N-1\} \tag{15a}$$

$$\sum_{t=1}^{N-1} u(t) \leqslant d_{\max}, \tag{15b}$$

where the vector inequalities are element-wise, $u_{\min}$ and $u_{\max}$ are the minimum and maximum infusion rates, and $d_{\max}$ is the maximum total dose. We consider input profiles that have $N-1$ time instances. This is because, the D matrix in the pharmacokinetic model is assumed to be a zero matrix. Therefore, with N measurements, we cannot observe the impact of $u(N)$ on the system.

IV. Two-Stage Approach

To address the lack of prior knowledge on the values of the pharmacokinetic parameters, we divide the experiment into two stages: An initial "learning" stage in which a fixed time profile of drug infusion is used, followed by an "optimization" stage where we use an infusion profile that has been optimized for the specific individual. This two-stage approach enable us to gather data during the learning stage to construct a coarse estimate for the parameters, which is used to compute an optimal infusion profile for the optimization stage. The overall procedure consists of the following steps:

1) Apply a fixed drug infusion profile from times $t = 1$ through $t = N_0$ and collect drug concentration measurements.
2) Compute a "coarse" posterior distribution for the unknown parameter vector θ, given the data collected for $t \in \{1, \dots, N_0\}$.
3) Compute the optimal infusion profile for times $t = N_0+1$ through $t = N-1$ by solving an optimization of the form (12), using samples drawn from distributions computed in step 2.
4) Apply the optimized drug infusion profile from times $t = N_0 + 1$ through the final time $t = N - 1$ and collect drug concentration measurements.
5) Compute a final estimate for the unknown parameter vector θ given the whole data collected for $t \in \{1, \dots, N\}$.

In the remainder of this section, we first discuss the approach used for system identification in steps 2 and 5, followed by the optimization performed in step 3.

A. System Identification

The approach used for system identification relies on the following property of multi-variable Gaussian distributions, which enables the computation of the posterior mean and covariance of the parameters θ given the measurement $\bar{y}$ from the joint distribution of θ and $\bar{y}$ through an optimization, rather than through integration of the joint density.

Lemma 1: For a given input $\bar{u}$, if the conditional distribution of θ given the measurements in $\bar{y}$ is a multivariate Gaussian, then this distribution has the following mean and covariance matrix:

$$\mathbb{E}\left[\theta|\bar{y};\bar{u}\right] = \underset{\theta\in\mathbb{R}^{n_\theta}}{\arg\max} \log p_{\Theta,\bar{Y}}(\bar{\theta},\bar{y};\bar{u}) \tag{16a}$$

$$\mathrm{Cov}\left[\theta|\bar{y};\bar{u}\right] = -\left(\frac{\partial \log p_{\Theta,\bar{Y}}(\theta,\bar{y};\bar{u})}{\partial\theta^2}\right)^{-1}, \tag{16b}$$

where $p_{\bar{Y},\Theta}(\theta,\bar{y};\bar{u})$ is the joint distribution of θ and $\bar{y}$, for the given input $\bar{u}$.

Proof: We can use the Bayes' rule to expand the logarithm of the conditional pdf of θ given the measurements $\bar{y}$ as follows

$$\log p_{\Theta|\bar{Y}}(\theta|\bar{y};\bar{u}) + \log p_{\bar{Y}}(\bar{y};\bar{u}) = \log p_{\Theta,\bar{Y}}(\theta,\bar{y};\bar{u}),$$

where $\log p_{\bar{Y}}(\bar{y};\bar{u})$ denotes the marginal distribution of $\bar{y}$. Since the conditional distribution of θ given $\bar{y}$ is assumed to be Gaussian, the logarithm of its pdf is of the form

$$\begin{aligned}\log p_{\bar{\theta}|\bar{Y}}(\theta|\bar{y};\bar{u}) = &-\frac{n_\theta}{2}\log(2\pi) + \frac{n_\theta}{2}\log\det(\Sigma^{-1}_{\theta|\bar{y}}) \\ &- \frac{1}{2}\left(\theta-\mu_{\theta|\bar{y}}\right)^T \Sigma^{-1}_{\theta|\bar{y}}\left(\theta-\mu_{\theta|\bar{y}}\right),\end{aligned} \tag{17}$$

where $\mu_{\theta|\bar{y}}$ and $\Sigma_{\theta|\bar{y}}$ denote the mean and covariance matrix of the conditional distribution. From this formula, we conclude that

$$\mathbb{E}\left[\theta|\bar{y};\bar{u}\right] = \mu_{\theta|\bar{y}} = \underset{\theta\in\mathbb{R}^{n_\theta}}{\arg\max} \log p_{\bar{Y},\Theta}(\bar{\theta},\bar{y};\bar{u})$$

$$\mathrm{Cov}\left[\theta|\bar{y};\bar{u}\right] = \Sigma_{\theta|\bar{y}} = -\left(\frac{\partial \log p_{\Theta,\bar{Y}}(\theta,\bar{y};\bar{u})}{\partial\theta^2}\right)^{-1}.$$

■

The usefulness of Lemma 1 stems from the observation that the joint pdf of θ and $\bar{y}$ that appears in the formulas in the right-hand side of (16) can be easily computed in closed form for the system (1) and is given by

$$p_{\Theta,\bar{Y}}\left(\theta,\bar{y};\bar{u}\right) = p_{\bar{Y}|\Theta}\left(\bar{y}|\theta;\bar{u}\right) p_\Theta(\theta), \tag{18}$$

where $p_{\bar{Y}|\Theta}\left(\bar{y}|\theta;\bar{u}\right)$ is defined as in (6). The result then shows that we can optimize/differentiate (18) to obtain the mean and variance of the posterior distribution using (16).

In general, the posterior conditional distribution in (17) is not Gaussian, but we can still regard the right-hand side of (17) as a second order Taylor series approximation to the exact distribution. In this case, the formulas in (16) will have some error that should be small as long as the second order approximation is reasonable.

Furthermore, if we assume that the prior distribution on θ in (18) is uninformative in the sense that it has a large covariance, optimizing/differentiating (18) will be approximately the same as optimizing/differentiating (6). Therefore, we solve the maximum likelihood estimation (MLE) problem for system identification

$$\hat{\theta} = \underset{\theta\in\mathbb{R}^{n_\theta}}{\arg\max} \log p\left(\bar{y}|\theta;\bar{u}\right), \tag{19}$$

and take the solution to be approximately (16a). We also note that inverse of the hessian of the MLE problem at the solution will give us the approximate value of (16b).

B. Experiment Design

We use the approximate Gaussian posterior distribution (16) obtained in the learning stage as a prior for the experiment design problem (12) under the constraints introduced in Section III.

$$\min_{\bar{u}\in\mathbb{R}^{N-1}} f\left(\frac{1}{K}\sum_{k=1}^{K} I_{\theta_k}(\bar{u})\right) \tag{20a}$$

$$\begin{aligned} \text{s.t.} \sum_{i=N_0+1}^{N_0+N-1} \bar{u}(i) &\leqslant d_{\max} \\ \bar{u}(i) &\geqslant u_{\min},\ \forall i \in \{N_0+1,\dots,N-1\} \\ \bar{u}(i) &\leqslant u_{\max},\ \forall i \in \{N_0+1,\dots,N-1\} \\ \bar{u}(i) &= \bar{u}_0(i),\ \forall i \in \{1,2,\dots,N_0\} \end{aligned} \tag{20b}$$

where the optimality criteria f is as in (8), $\hat{I}_\theta$ is defined as in (13), N_0 is the initial input horizon, $\bar{u}_0$ is the initial input profile applied during this time, and N is the number of samples for the optimal experiment design. We highlight that FIM, $\hat{I}_\theta$, depends on a large $\bar{u}$ with $N-1$ entries. However, first N_0 entries of it are fixed to $\bar{u}_0(i)$ in order to account for the initial data collection period.

We then use a nonlinear programming solver to solve (20) but it is typically important to pick good initializations for the solver due to the lack of convexity. We choose a finite family of "Bang-bang" injection protocols with control inputs whose values switch between rates $u_{\min}$ and $u_{\max}$. This type of input profiles is commonly used in practice and also found to be performing well in previous studies [13]. Motivated by this, we propose to initialize the nonlinear programming solver with the best such input. Over a finite discrete-time horizon of length N, there is only a finite number of bang-bang inputs but comparing all of them is generally not practical because the number of possible combinations grows super-exponentially with the horizon length N even for a single input. The following process of choosing initializations is defined to be a scalar input but for a multi-channel input case, one can use this procedure for each channel independently, which would take the complexity to the n_uth power.

We look for the combinations in a coarse manner in order to limit the number of possible combinations. We divide the overall design interval of length $N-N_0-1$ into m segments of length $N-N_0-1/m$ and keep the input constant over each segment. More precisely we choose ℓ segments to infuse at the maximum allowed rate and at the minimum rate for the remaining $m-\ell$ segments. Therefore, we use sets of length ℓ to define each input we consider.

$$u_{\bar{c}}(t) = \begin{cases} u_{\max} & \left\lceil \frac{t}{m} \right\rceil \in \bar{c} \\ u_{\min} & \text{otherwise} \end{cases} \tag{21}$$

$\forall t \in \{1,2,\dots,N-N_0-1\}$, where $\lceil x \rceil$ stands for the smallest integer larger than $x \in \mathbb{Q}$, and $\bar{c}$ is any subset of $\{1,2,\dots,m\}$ with ℓ elements. We note that integers, $m, \ell \in \mathbb{Z}$, must satisfy

$$\frac{\ell}{m} \leqslant \frac{d_{\max} - u_{\min}(N-N_0-1)T}{(N-N_0-1)(u_{\max}-u_{\min})T}$$

in order to satisfy the total dosage constraint.

Now we choose m, ℓ such that the total dosage requirement is satisfied and the total number of possible combinations is not very high. Then, we define the set of all possible combinations, $C(m,l)$, and find the optimal combination,

$$\bar{c}^* = \{c_1^*, c_2^*, \dots, c_l^*\} = \underset{\bar{c}\in C(s,l)}{\arg\min}\, f\Big(\frac{1}{K}\sum_{k=1}^{K} I_{\theta_k}(u_{\bar{c}})\Big)$$

where f is one of the criteria in (8), and the coarse input $u_{\bar{c}}$ is defined as (21).

Once we have the initialization, we run a primal-dual interior point solver [15] to solve the problem (12).

V. Results and Discussion

We use a hypothetical drug that does not appear in the body naturally and follows the three compartment model depicted in Fig. 1.

The drug is injected into the blood stream (which corresponds to the "vein" compartment) and diffuses both to the brain (which is represented by the "brain" compartment) and to the rest of body that we represent by a single "distribution compartment." We can measure drug concentration in the first two compartments, which is in line with the current capabilities of the electrochemical sensing [4], [21], [22]

The corresponding pharmacokinetic model follows from (14) and is of the form

$$x(t+1) = \begin{bmatrix} -k_{V_0} & 0 & k_{D_V} \\ k_{V_B} & -k_{B_0} & 0 \\ k_{V_D} & 0 & -k_{D_0} \end{bmatrix} x(t) + \begin{bmatrix} k_{U_V} \\ 0 \\ 0 \end{bmatrix} u(t) \tag{22a}$$

$$y(t) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} x(t) + \eta_t, \tag{22b}$$

where $k_{V_0} = k_{E_V} + k_{D_V} - 1,\ k_{B_0} = k_{E_B} + k_{V_B} - 1,\ k_{D_0} = k_{V_D} - 1$, η_t are independent and identically distributed zero-mean Gaussian noise and we use the notation described

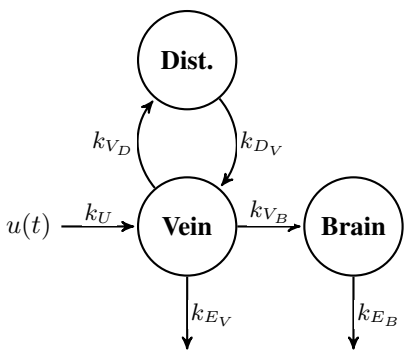

Fig. 1. Network view of three compartment model.

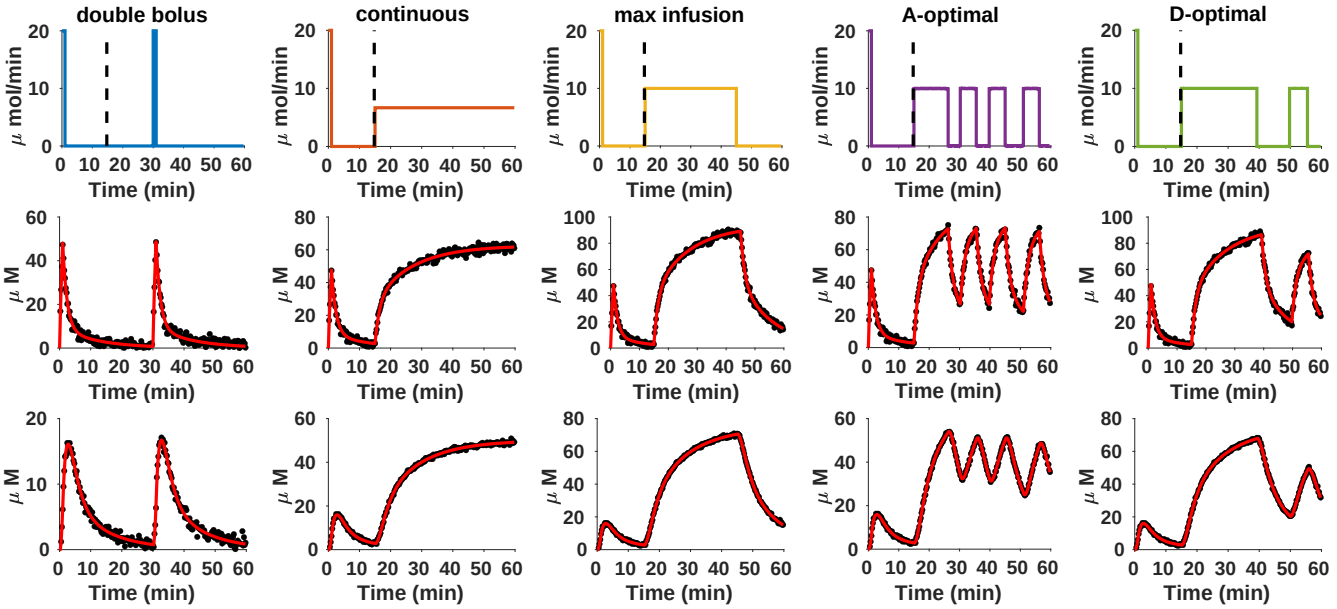

Fig. 2. Comparison of different drug administration profiles. The drug administration profiles are shown in the top row, where a dashed black line separates the initial data collection period from the experiment design period. Each input profile is shown with a color consistent with Fig. 4. The second and the third rows show the corresponding simulated vein and brain responses, respectively, with noisy measurements (black dots), and the estimated system response (red line). Each column of plots corresponds to specific drug administration profile, from left to right: injection of a second bolus injection at the maximum level allowed, continuous administration of the drug at the maximum constant level allowed by the total dose constraint, administration of a maximum magnitude with duration limited by the total dose constraint, the A-optimal input, and the D-optimal input.

in Section III, with the understanding that the subscripts V, B and D stand for the vein, brain and the distribution compartments. We assume that the additive noise in-vein and in-brain measurements are not correlated and they have variances σ_V^2 and σ_B^2,respectively. We use the values given in Table I to simulate the system and find an optimal input profile. The values we use are consistent with what we observe during pharmacokinetic experiments carried out in our laboratory [4], [23].

We follow a scenario where we have a design horizon of 60 minutes, i.e. $N = 241$ (including the last sample), and we are able to sample the system every 15 seconds [3]. We divide this horizon into two stages: a 15 minutes, i.e. $N_0 = 60$, long learning stage and 45 minutes long optimization stage. We assume that we have two allowed drug administration protocols;

1) Bolus injection at 20μmol/min over a minute and can be repeated every 30 minutes.
2) Continuous infusion at maximum rate of 20μmol/min for a total of maximum 300μmol of drug.

We start by applying the bolus injection protocol. We then collect data during the learning stage, and use this data to obtain a distribution over parameter values θ, using the methodology in Section IV-A.

Once we have this distribution, we sample from it to compute the average matrices appearing in the quadratic terms of FIM as in (13). We use $K = 1000$ samples for this purpose. Then we design the optimal input profile in the optimization stage, following the continuous infusion protocol with $u_{\min} = 0$. This is done by using the method in Section IV with $N_0 = 60$ and $N = 160$ corresponding to a 15 and 45 minutes stages, respectively.

We use the initialization approach described in Section IV-B with $m = 15, \ell = 10$, which corresponds to 3003 possible combinations. We then use a nonlinear solver to find local solutions to the optimization problem (20). We particularly use IPOPT [15] together with linear algebra solver MA57 from HSL [24], and software tool that uses automatic differentiation (AD) to compile our problem, CasADI [14]. An important advantage of using an AD-based solver is that the solver, once created (ahead of the experiment), does not need to numerically approximate to a large hessian and gradient at every iteration. This allows fast enough computation times for real-time experiments.

We present our findings in Fig. 2, where we compare drug injection profiles commonly used in pharmacological studies with the ones that result from the approach proposed here, based on optimizing the A and D criteria.

All the drug administration profiles considered in Fig. 2 start with the same initial bolus that is used to obtain a coarse

Parameter	Value
k_U	0.75 min/L
k_{D_V}	0.06
k_{V_B}	0.08
k_{V_D}	0.04
k_{E_V}	0.08
k_{E_B}	0.02
σ_V^2	2.5 μM^2
σ_B^2	0.25 μM^2

TABLE I

NUMERICAL VALUES FOR OUR HYPOTHETICAL EXPERIMENT

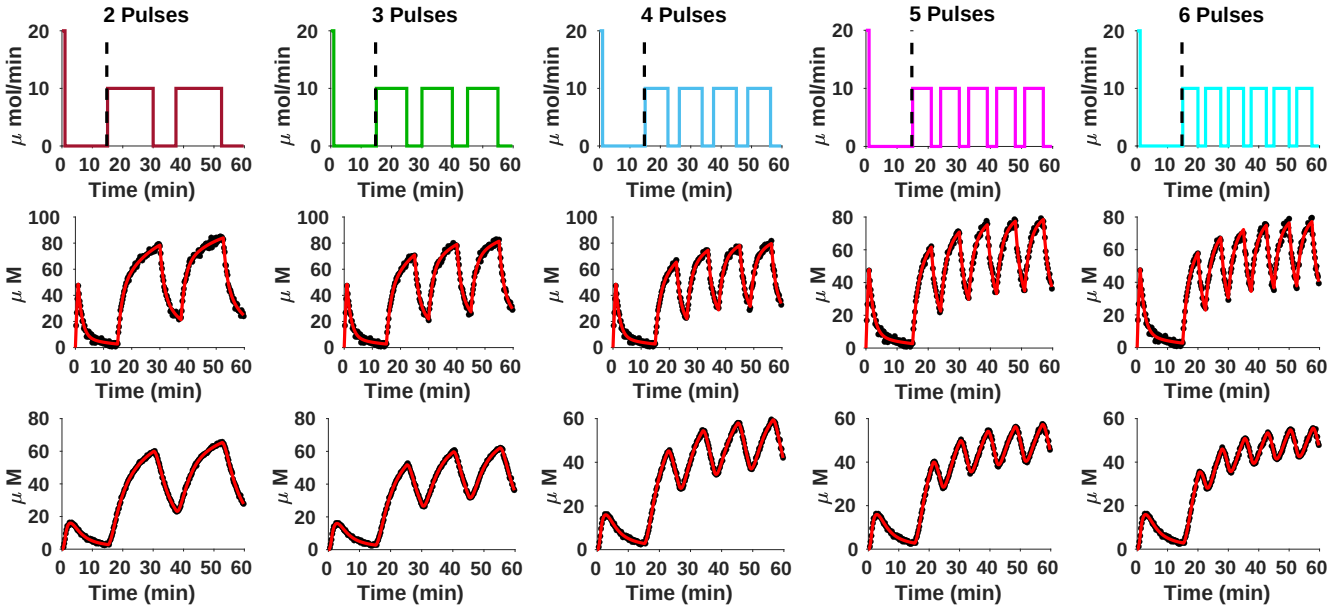

Fig. 3. Additional drug administration profiles, defined by different number of equal length pulses. The plots follow the same formating as Fig. 2.

estimate for the unknown parameters. Two of them (on the right-hand side of the figure) then proceed to apply optimal injection profiles, whereas the three other follow injection profiles commonly used by practitioners: The first of these corresponds to the injection of a second bolus injection at the maximum level allowed, the second corresponds to the continuous administration of the drug at the maximum constant level allowed by the total dose constraint, and the third to a pulse with magnitude $u_{\max}$ and duration limited by the total dose constraint. In contrast, both optimal profiles consist of multiple pulses. This led us to question how some obvious pre-determined pulsed profiles compare with our proposed optimal input profiles. We consider 5 more profiles with two to six pulses as shown in Fig. 3.

We present all our results in Table II. A and D optimal infusion profiles outperform other candidates in their respective optimality cost measures, as expected. We see that the number of pulses in the infusion profile make a difference as the infusion profiles with 4 and 2 pulses perform very well in terms of A and D optimality costs after their optimal counterparts. However, the number of pulses is not the only factor and the differences in individual pulse lengths also matter as can be seen in the varying widths of pulses in optimal profiles.

We also looked at how the performance of an infusion profile evolves over time. This gives us an idea of how the pulses and their widths affect the identification performance. We choose five input profiles to do this analysis: the continuous infusion, the input profiles with 2 and 4 pulses, and the optimal input profiles. We chose the inputs with 2 and 4 pulses since these correspond to the number of pulses in the optimal profiles, and the continuous infusion is for the case with no pulse.

We estimate the parameters every minute after the initial learning phase which all input profiles in Fig. 2 and Fig. 3 share. We use (16b) to approximate the posterior covariance of the parameter estimates. This allows the computation of 95% marginal confidence intervals (CIs) on each estimate. We finally report the evolution of the sum of all 95% CIs in Fig. 4, which should parallel the A-optimal costs and more meaningful for experimenters.

We observe that the continuous infusion steadily decreases the sum of CIs. The decrease slows down after about minute 40 which corresponds to the time in-vein and in-brain measurements starts leveling off.

Input	Identification		Sampled FIM	
	A-opt	D-opt	A-opt	D-opt
double bolus	1.52×10^{-4}	-76.88	1.51×10^{-4}	-76.69
continuous	1.84×10^{-4}	-85.82	1.25×10^{-4}	-86.37
max infusion	9.26×10^{-5}	-90.87	6.70×10^{-5}	-91.33
2 Pulses	6.12×10^{-5}	-91.23	5.29×10^{-5}	-91.17
3 Pulses	4.77×10^{-5}	-90.87	4.38×10^{-5}	-90.74
4 Pulses	4.58×10^{-5}	-90.20	4.08×10^{-5}	-90.18
5 Pulses	4.40×10^{-5}	-89.74	4.13×10^{-5}	-89.61
6 Pulses	4.95×10^{-5}	-88.95	4.33×10^{-5}	-89.10
A-optimal	4.36×10^{-5}	-90.76	$4.00 \times 10^{-5}*$	-90.47
D-optimal	5.90×10^{-5}	-92.22	4.79×10^{-5}	-92.01*

TABLE II

VALUES FOR THE OPTIMALITY CRITERIA OBTAINED WITH THE FIVE DIFFERENT INPUT PROFILES SHOWN IN FIG. 2 AND IN FIG. 3. IN THE COLUMNS LABELED AS "IDENTIFICATION", WE PROVIDE THE VALUE OF THE A AND D OPTIMALITY CRITERIA COMPUTED FOR THE ERROR COVARIANCE MATRICES OBTAINED FROM THE SYSTEM IDENTIFICATION APPROACH OUTLINED IN SECTION IV-A, WHEREAS IN THE ROWS LABELS AS "SAMPLED FIM" WE PROVIDE THE CORRESPONDING VALUE OF THE OPTIMIZATION CRITERIA IN (12) USED FOR EXPERIMENT DESIGN. THE VALUES HIGHLIGHTED WITH $*$ ARE THUS NECESSARILY THE SMALLEST WITHIN THEIR COLUMNS.

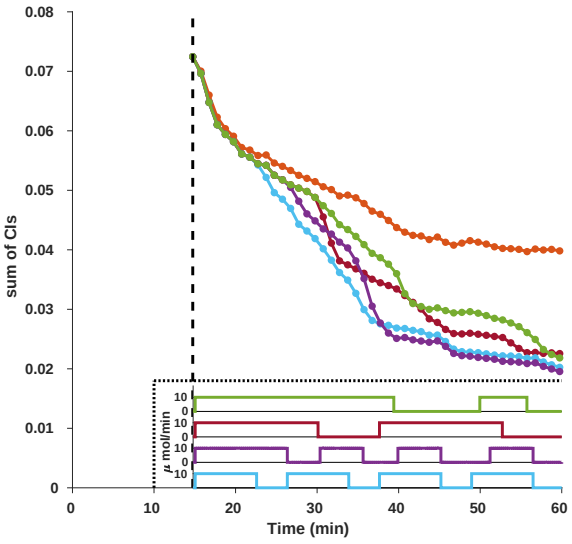

Fig. 4. Change in the sum of confidence intervals as more data is given to the estimator for the continuous (orange), the 2 pulses (dark red), the 4 pulses (turquoise), D-optimal (light green), and A-optimal (puple) infusion profiles. We also provide the infusion profiles (except the continuous infusion) with a common x-axis in a small box at the bottom to make it easier to see how the switching of the input value changes the CIs.

Some more interesting insights can be gained when comparing the other input profiles. The inputs with both 4 pulses and 2 pulses enables faster decreases in the beginning, compared to A and D optimal inputs, respectively. This can be attributed to the impact of earlier switching off of the input. However, in the long run the optimal inputs outperform their competitors. This is expected as optimal inputs are optimized for the entire horizon length not for shorter horizons.

VI. Conclusions and Future Work

We have shown how experiment design can be used to improve the accuracy of parameter estimation when there is very little prior information about the parameter values.

In the context of realistic pharmacokinetic parameters and models, we see that optimal infusion profiles that are designed based on an initial coarse system identification procedure can decrease the uncertainty around the identified parameters significantly, leading to 30-50% reductions in sum of the confidence intervals. Because the method does not depend on values obtained from literature or from prior experiments, this work is especially important for systems with dynamics that are highly subject-dependent.

In a future work, we will consider additional safety requirements that are common in animal experiments, such as a maximum allowed concentration in a body part. We are also excited about the real world application of this technique and what it will reveal about the pharmacokinetics of new and existing drugs.

References

[1] T. Loftsson, *Essential pharmacokinetics: A primer for pharmaceutical scientists*. Academic Press, 2015.

[2] N. Arroyo-Currás, G. Ortega, D. A. Copp, K. L. Ploense, Z. A. Plaxco, T. E. Kippin, J. a. P. Hespanha, and K. W. Plaxco, "High-precision control of plasma drug levels using feedback-controlled dosing," *ACS pharmacology & translational science*, vol. 1, no. 2, pp. 110–118, 2018.

[3] N. Arroyo-Currás, J. Somerson, P. A. Vieira, K. L. Ploense, T. E. Kippin, and K. W. Plaxco, "Real-time measurement of small molecules directly in awake, ambulatory animals," *Proceedings of the National Academy of Sciences*, vol. 114, no. 4, pp. 645–650, 2017.

[4] P. Dauphin-Ducharme, K. Yang, N. Arroyo-Currás, K. L. Ploense, Y. Zhang, J. Gerson, M. Kurnik, T. E. Kippin, M. N. Stojanovic, and K. W. Plaxco, "Electrochemical aptamer-based sensors for improved therapeutic drug monitoring and high-precision, feedback-controlled drug delivery," *ACS sensors*, vol. 4, no. 10, pp. 2832–2837, 2019.

[5] L. Ljung, *System Identification: Theory for the User*. Prentice Hall information and system sciences series, Prentice Hall PTR, 1999.

[6] A. Pázman, *Foundations of optimum experimental design*, vol. 14. Springer, 1986.

[7] H. Hjalmarsson, "System identification of complex and structured systems," *European journal of control*, vol. 15, no. 3-4, pp. 275–310, 2009.

[8] L. Pronzato, "Optimal experimental design and some related control problems," *Automatica*, vol. 44, no. 2, pp. 303–325, 2008.

[9] M. Gevers, "Identification for control: From the early achievements to the revival of experiment design," *European journal of control*, vol. 11, no. 4-5, pp. 335–352, 2005.

[10] H. Jansson and H. Hjalmarsson, "Input design via lmis admitting frequency-wise model specifications in confidence regions," *IEEE transactions on Automatic Control*, vol. 50, no. 10, pp. 1534–1549, 2005.

[11] I. R. Manchester, "Input design for system identification via convex relaxation," in *49th IEEE Conference on Decision and Control (CDC)*, pp. 2041–2046, IEEE, 2010.

[12] I. R. Manchester, "Amplitude-constrained input design: Convex relaxation and application to clinical neurology," *IFAC Proceedings Volumes*, vol. 45, no. 16, pp. 1617–1622, 2012.

[13] J. Maidens and M. Arcak, "Semidefinite relaxations in optimal experiment design with application to substrate injection for hyperpolarized mri," in *2016 American Control Conference (ACC)*, pp. 2023–2028, IEEE, 2016.

[14] J. A. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, "Casadi: a software framework for nonlinear optimization and optimal control," *Mathematical Programming Computation*, vol. 11, no. 1, pp. 1–36, 2019.

[15] A. Wächter and L. T. Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," *Mathematical programming*, vol. 106, no. 1, pp. 25–57, 2006.

[16] B. P. Duarte, A. C. Atkinson, and N. M. Oliveira, "Optimal experimental design for linear time invariant state–space models," *Statistics and Computing*, vol. 31, no. 4, pp. 1–20, 2021.

[17] J. P. Hespanha, *Linear systems theory*. Princeton university press, 2018.

[18] A. J. Kleywegt, A. Shapiro, and T. Homem-de Mello, "The sample average approximation method for stochastic discrete optimization," *SIAM Journal on Optimization*, vol. 12, no. 2, pp. 479–502, 2002.

[19] S. Boyd, S. P. Boyd, and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.

[20] P. Macheras and A. Iliadis, *Modeling in biopharmaceutics, pharmacokinetics and pharmacodynamics: homogeneous and heterogeneous approaches*, vol. 30, ch. 8. Springer, 2016.

[21] J. G. Roberts and L. A. Sombers, "Fast scan cyclic voltammetry: Chemical sensing in the brain and beyond," *Analytical chemistry*, vol. 90, no. 1, p. 490, 2018.

[22] M. Sarter and Y. Kim, "Interpreting chemical neurotransmission in vivo: techniques, time scales, and theories," *ACS chemical neuroscience*, vol. 6, no. 1, pp. 8–10, 2015.

[23] J. Gerson, M. Erdal, K. Plaxco, T. Kippin, and J. Hespanha, "Vein-to-brain: Simultaneous, seconds-resolved measurements of intracranial and intravenous drug levels provide a highly time-resolved picture of drug transport," *The FASEB Journal*, vol. 35, 2021.

[24] A. HSL, "collection of fortran codes for large-scale scientific computation," 2007.