Optimal Bayesian Experimental Design for Priors of Compact Support with Application to Shock-Tube Experiments for Combustion Kinetics

Fabrizio Bisetti^{1,2}, Daesang Kim^{1,2}, Omar Knio¹, Quan Long^{1,*} and Raul Tempone¹

¹*SRI Center for Uncertainty Quantification, King Abdullah University of Science and Technology,
23955 Thuwal, Saudi Arabia*

²*Clean Combustion Research Center, King Abdullah University of Science and Technology,
23955 Thuwal, Saudi Arabia*

SUMMARY

The analysis of reactive systems in combustion science and technology relies on detailed models comprising many chemical reactions that describe the conversion of fuel and oxidizer into products and the formation of pollutants. Shock-tube experiments are a convenient setting for measuring the rate parameters of individual reactions. The temperature, pressure, and concentration of reactants are chosen to maximize the sensitivity of the measured quantities to the rate parameter of the target reaction. In this study, we optimize the experimental setup computationally by optimal experimental design (OED) in a Bayesian framework. We approximate the posterior probability density functions (pdf) using truncated Gaussian distributions in order to account for the bounded domain of the uniform prior pdf of the parameters. The underlying Gaussian distribution is obtained in the spirit of the Laplace method, more precisely, the mode is chosen as the maximum a posteriori (MAP) estimate, and the covariance is chosen as the negative inverse of the Hessian of the misfit function at the MAP estimate. The model related entities are obtained from a polynomial surrogate. The optimality, quantified by the information gain measures, can be estimated efficiently by a rejection sampling algorithm against the underlying Gaussian probability distribution, rather than against the true posterior. This approach offers a significant error reduction when the magnitude of the invariants of the posterior covariance are comparable to the size of the bounded domain of the prior. We demonstrate the accuracy and superior computational efficiency of our method for shock-tube experiments aiming to measure the model parameters of a key reaction which is part of the complex kinetic network describing the hydrocarbon oxidation. In the experiments, the initial temperature and fuel concentration are optimized with respect to the expected information gain in the estimation of the parameters of the target reaction rate. We show that the expected information gain surface can change its “shape” dramatically according to the level of noise introduced into the synthetic data. The information that can be extracted from the data saturates as a logarithmic function of the number of experiments, and few experiments are needed when they are conducted at the optimal experimental design conditions. Furthermore, inversion of the legacy data indicates the validity and robustness of our designs. Copyright © 2016 John Wiley & Sons, Ltd.

Received ...

KEY WORDS: Bayesian optimal experimental design; Laplace approximation; Truncated Gaussian approximation; Polynomial surrogates; Shock-tube; Combustion kinetics

*Correspondence to: Quan Long, SRI Center for Uncertainty Quantification, King Abdullah University of Science and Technology, 23955 Thuwal, Saudi Arabia. E-mail: quan.long@kaust.edu.sa

This author's current affiliation is United Technologies Research Center, East Hartford, CT, USA. E-mail: longq@utrc.utc.com.

This article has been accepted for publication and undergone full peer review but has not been through the copyediting, typesetting, pagination and proofreading process, which may lead to differences between this version and the Version of Record. Please cite this article as doi: 10.1002/nme.5211

1. INTRODUCTION

If the uncertainties in a parameter estimation problem are to be reduced, one of the key questions is how the experiments should be designed, such that the data would be as informative as possible and the effect of uncertainties in the experiments would be mitigated. We aim to design combustion kinetics experiments in a Bayesian framework [4, 2, 33], in which the solution of the inverse problem is a probability distribution instead of a deterministic value. The amount of information carried by the data from an experiment can be measured by the expected Kullback–Leibler divergence, namely the expected information gain or mutual information [19, 18, 11]. This measure is defined as the expected logarithm of the ratio between the posterior and prior probability density functions over all the possible model parameters and data values. Note that a trivial experiment would lead to zero information gain.

Although the expected information gain is a very useful utility function in OED, its estimation presents a computational challenge. The nested integrals required to compute the K-L divergence and its expectation over the parameter and data spaces [35] lead to a biased estimator, which can be computationally costly, especially when the evaluation of the model requires large systems of differential equations to be solved numerically. In [15], the authors used polynomial surrogates for their model to accelerate computation of the expected information gain in the context of combustion OED. In addition, the Laplace method has been used for fast estimation of the expected information gain [23, 24, 22]. Using the Laplace approximation, the two nested integrals are reduced to a single integral, significantly accelerating the computation. The novel aspect of the work presented in this paper is that we approximate the posterior pdf using a truncated Gaussian pdf defined on the bounded parametric domain. While the underlying Gaussian pdf can be obtained by truncating a Taylor expansion of the logarithmic posterior after the second order terms, the integration can be carried out by a highly efficient rejection sampling algorithm. Thus, we avoid the error of the Gaussian integral when a significant amount of posterior probability mass extends beyond the bounded parametric domain.

We apply the new technique as part of an OED aiming to select the initial temperature and fuel concentration in a series of shock-tube experiments. These experiments have been devised to measure the model parameters of a key reaction in hydrocarbon combustion, the elementary reaction $\text{H} + \text{O}_2 \rightleftharpoons \text{OH} + \text{O}$ [13]. We adopt the Arrhenius form of the reaction rate and assume a uniform prior pdf of the pre-exponential factor and activation energy with compact support. The observable is the maximum rate of increase of the water concentration in the shock-tube during the combustion of hydrogen/oxygen mixtures in argon diluent. The mapping from the parameters to the observable is constructed using a pseudo spectral polynomial surrogate [25, 10].

In Section 2, we first review the key concepts behind OED in a Bayesian framework and the Laplace method for computing the expected information gain. We then introduce a new approach to computing the expected information gain based on approximating the posterior pdf by a truncated Gaussian pdf. In Section 3, we describe the rejection sampling algorithm and briefly summarize the pseudo spectral polynomial surrogate. In Section 4, we apply the OED approach to the design of an optimal shock-tube experiment. We compare the convergence performance of the proposed approach against the direct Laplace approximation and the nested Monte Carlo sampling. We

analyse the expected information gains from different design scenarios and noise levels, and identify optimal design conditions. Conclusions are presented in Section 5.

2. EXPECTED INFORMATION GAIN IN A PROPOSED EXPERIMENT

2.1. Data, model and information gain

The model of the observable consists of a deterministic model and additive noise:

$$\mathbf{y} = \mathbf{g}(\boldsymbol{\theta}_0, \boldsymbol{\xi}) + \boldsymbol{\epsilon},$$

where $\mathbf{y}$ is an $(s \cdot n) \times 1$ total observable vector, n is the length of the vector for an individual observation, and s is the cardinality of the set of all the individual measurements. The function $\mathbf{g} : \mathcal{R}^d \times \mathcal{R}^c \rightarrow \mathcal{R}^{(s \cdot n)}$ is the deterministic model, $\boldsymbol{\theta}_0$ is the “true” parameter used to generate synthetic data, $\boldsymbol{\xi}$ is the vector of parameters describing the experimental setup, and $\boldsymbol{\epsilon}$ is an additive random noise with mean $\mathbf{0}$ and covariance matrix Σ_e .

Bayes theorem states that:

$$p(\boldsymbol{\theta}|\mathbf{y}, \boldsymbol{\xi}) = \frac{p(\mathbf{y}|\boldsymbol{\theta}, \boldsymbol{\xi})p(\boldsymbol{\theta})}{p(\mathbf{y}|\boldsymbol{\xi})},$$

where $p(\boldsymbol{\theta})$ and $p(\boldsymbol{\theta}|\mathbf{y}, \boldsymbol{\xi})$ are the prior and posterior pdfs of the parameter $\boldsymbol{\theta}$, respectively, $p(\mathbf{y}|\boldsymbol{\theta}, \boldsymbol{\xi})$ is the likelihood of the data $\mathbf{y}$ for given values of $\boldsymbol{\theta}$ and $\boldsymbol{\xi}$, and $p(\mathbf{y}|\boldsymbol{\xi})$ is the so-called evidence term. The Kullback–Leibler divergence is a measure of the amount of information gain from the data and is defined as

$$D_{KL}(\boldsymbol{\xi}, \mathbf{y}) := \int_{\Theta} \log \left[\frac{p(\boldsymbol{\theta}|\mathbf{y}, \boldsymbol{\xi})}{p(\boldsymbol{\theta})} \right] p(\boldsymbol{\theta}|\mathbf{y}, \boldsymbol{\xi}) d\boldsymbol{\theta}. \quad (1)$$

For a proposed experiment characterized by a choice of parameters $\boldsymbol{\xi}$, the expected information gain is the expected Kullback–Leibler divergence, i.e. the marginalization of the K-L divergence over all possible values of the data. This quantity is also called the expected information gain or mutual information between the parameters and the observables. The expected information gain $I(\boldsymbol{\xi})$ is a function of the design parameters and is defined as

$$I(\boldsymbol{\xi}) = \int_{\mathcal{Y}} D_{KL}(\boldsymbol{\xi}, \mathbf{y}) p(\mathbf{y}|\boldsymbol{\xi}) d\mathbf{y}.$$

In an optimal design problem, we seek to find the value of $\boldsymbol{\xi}$ which maximizes the expected information gain, i.e. $\boldsymbol{\xi}_{opt} := \arg \max_{\boldsymbol{\xi} \in \mathcal{X}} I(\boldsymbol{\xi})$, where $\mathcal{X}$ is the search space of the design parameters.

2.2. Laplace method for OED

We use a recently developed fast method to compute the expected information gain for both determined and under-determined experiments. This novel method relies on Laplace approximations

to estimate the integrals against the concentrated posterior distributions. In a determined experiment, the posterior pdf has a single dominant mode. In an under-determined experiment, by contrast, the posterior pdf is characterized by non-informative manifolds. Details of the mathematical derivation and error analysis were taken from Long et al. [23, 24].

The assumption underlying use of the Laplace approximation in experimental design is that the posterior pdf is concentrated around the maximum a posteriori probability (MAP) estimate, such that a Gaussian integral can be carried out to approximate any posterior integration asymptotically. In practice, this assumption is satisfied if a large amount of data is collected or the measurements can be conducted with low noise. In the classical scenario, the posterior pdf is approximated by a Gaussian distribution:

$$\tilde{p}(\boldsymbol{\theta}|\mathbf{y}, \boldsymbol{\xi}) = \frac{1}{\sqrt{2\pi}^d |\boldsymbol{\Sigma}|^{\frac{1}{2}}} \exp \left[-\frac{1}{2} \|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}\|_{\boldsymbol{\Sigma}}^2 \right],$$

where the posterior covariance, $\boldsymbol{\Sigma} = \mathbf{H}^{-1}$ with $\mathbf{H} = \nabla_{\boldsymbol{\theta}} \mathbf{g}^\top \boldsymbol{\Sigma}_\epsilon^{-1} \nabla_{\boldsymbol{\theta}} \mathbf{g}$, $\hat{\boldsymbol{\theta}}$ is the MAP estimate. Based on this approximation, the K-L divergence, D_{KL} , may be computed analytically, thereby reducing the overall computational cost significantly.

A further extension was made in [24] to deal with a posterior pdf which does not have a dominant single mode, but instead a MAP manifold. The underlying assumption is that the posterior pdf is concentrated in directions orthogonal to the MAP manifold, i.e. the kernel of the Jacobian matrix of the model output with respect to the parameters. Therefore, the Laplace approximation in these directions leads to a simplified version of the expected information gain.

The main results obtained in [23, 24] are as follows.

Theorem 1

Under the assumption that the smallest singular value of the Jacobian matrix is bounded from zero by a small constant and the model output $\mathbf{g}(\boldsymbol{\theta}, \boldsymbol{\xi})$ has continuous second derivatives, the expected information gain can be approximated by

$$I(\boldsymbol{\xi}) = \int_{\Theta} \left[-\frac{1}{2} \log(|\boldsymbol{\Sigma}(\boldsymbol{\theta}_0, \boldsymbol{\xi})|) - \frac{\text{tr}(\boldsymbol{\Sigma}(\boldsymbol{\theta}_0, \boldsymbol{\xi}) \mathbf{H}_h(\boldsymbol{\theta}_0))}{2} - \frac{d}{2} - \frac{d}{2} \log(2\pi) \right] p(\boldsymbol{\theta}_0) d\boldsymbol{\theta}_0 + \mathcal{O}\left(\frac{1}{s}\right), \quad (2)$$

where $\mathbf{H}_h$ is the Hessian of $h(\boldsymbol{\theta}) = \log[p(\boldsymbol{\theta})]$.

Theorem 2

Under the assumption that $r < d$ and that the model output $\mathbf{g}(\boldsymbol{\theta}, \boldsymbol{\xi})$ has continuous second derivatives, the expected information gain can be approximated by

$$I(\boldsymbol{\xi}) = \int_{\Theta} \left[-\frac{1}{2} \log |\boldsymbol{\Sigma}_{proj}(\boldsymbol{\theta}_0, \boldsymbol{\xi})| - \log \left[\int_T p_{s,t}(\mathbf{0}, \mathbf{t}) d\mathbf{t} \right] - \frac{r}{2} - \frac{r}{2} \log(2\pi) \right] p(\boldsymbol{\theta}_0) d\boldsymbol{\theta}_0 + \mathcal{O}\left(\frac{1}{s}\right), \quad (3)$$

where the projected covariance matrix is $\boldsymbol{\Sigma}_{proj} = (\mathbf{U}^\top \boldsymbol{\Sigma}^{-1} \mathbf{U})^{-1}$, $\mathbf{U}$ contains the basis spanning the orthogonal space of the Jacobian kernel, r is the rank of $\boldsymbol{\Sigma}$, $\int_T p_{s,t}(\mathbf{0}, \mathbf{t}) d\mathbf{t}$ is a manifold integration of the prior after a change of parameter, $\mathbf{t}$ is the new parameter which parameterizes the manifold and s is the new parameter which parameterizes the orthogonal directions.

Remark 1

The marginal prior on the MAP manifold can be accurately approximated by a linearization and a change of variable [22].

2.3. Truncated Gaussian Approximation

Approximating the posterior pdf using a Gaussian is a powerful approach to compute the expected information gain efficiently [23, 24, 22]. Nonetheless, the Gaussian approximation may lead to erroneous results when the support of the prior pdf is compact. An illustration is provided in Figure 1. The prior pdf is a uniform distribution $\mathcal{U}(\theta_l, \theta_u)$ with $\theta_l = [-1, -1]^\top$ and $\theta_u = [1, 1]^\top$. The “ground truth” posterior pdf shown in color is a truncated Gaussian pdf, i.e. θ has a Gaussian distribution $\mathcal{N}(\bar{\theta}, \Sigma)$, with

$$\bar{\theta} = [-0.4, 0.3]^\top \text{ and } \Sigma = \begin{bmatrix} 1.0 & 0.0 \\ 0.0 & 0.7 \end{bmatrix},$$

conditional on

$$\theta_l \leq \theta \leq \theta_u.$$

Note that the mean vector and covariance matrix of the truncated Gaussian distribution are

$$[-0.12, 0.12]^\top \text{ and } \begin{bmatrix} 0.28 & 0.0 \\ 0.0 & 0.27 \end{bmatrix}.$$

On the same plot, the Gaussian distribution $\mathcal{N}(\bar{\theta}, \Sigma)$ is shown in black. The distribution extends considerably beyond the limit of the compact support. In this schematic example it is clear that using the Gaussian distribution to approximate the truncated Gaussian distribution for computing the expected information gain leads to erroneous results.

The following theorem holds.

Theorem 3

Under the assumption that the smallest singular value of the Jacobian matrix is bounded from zero by a small constant and the model output $g(\theta, \xi)$ has continuous second derivatives, the expected information gain can be approximated by

$$I(\xi) = \int_{\Theta} \tilde{D}_{KL}(\theta_0, \xi) p(\theta_0) d\theta_0 + \mathcal{O}\left(\frac{1}{s}\right), \quad (4)$$

with

$$\tilde{D}_{KL}(\theta_0, \xi) = \int_{\Theta} \frac{\phi(\theta|y)}{p(\theta)} \phi(\theta|y) d\theta \quad \text{and} \quad \phi(\theta|y) = \frac{\tilde{p}(\theta|y, \xi)}{\int_{\Theta} \tilde{p}(\theta|y, \xi) d\theta}.$$

Proof

The proof is very similar to that of Theorem 1 in [23] and only an outline is provided here. Let

$$\phi(\theta) = p(\hat{\theta}|y) \exp\left(-\frac{(\theta - \hat{\theta})^\top \Sigma (\theta - \hat{\theta})}{2}\right) = \frac{p(\theta|y)}{\exp\left(\mathcal{O}_{\mathcal{P}}(|\theta - \hat{\theta}|^3)\right)},$$

where $\hat{\theta}$ is the global optimum of $p(\theta|\mathbf{y})$.

The Kullback-Leibler divergence can be decomposed into several terms as follows:

$$\begin{aligned} D_{KL} &= \int_{\Theta} \log \frac{p(\theta|\mathbf{y})}{p(\theta)} p(\theta|\mathbf{y}) d\theta \\ &= \int_{\Theta} \log \frac{\phi(\theta|\mathbf{y})}{p(\theta)} \phi(\theta|\mathbf{y}) d\theta + \int_{\Theta} \log \frac{p(\theta|\mathbf{y})}{\phi(\theta|\mathbf{y})} \phi(\theta|\mathbf{y}) d\theta + \int_{\Theta} \log \frac{p(\theta|\mathbf{y})}{p(\theta)} (p(\theta|\mathbf{y}) - \phi(\theta|\mathbf{y})) d\theta. \end{aligned} \quad (5)$$

By Taylor expansion it can be shown that the second and third terms on the right hand side of (5) are of order $\mathcal{O}_P\left(\frac{1}{s^2}\right)$ (see the Appendix of [23] for details). Hence, one may approximate the expected information gain using $\phi(\theta|\mathbf{y})$:

$$\begin{aligned} I &= \int_{\mathbf{y}} D_{KL} p(\mathbf{y}|\theta_0) p(\theta_0) d\theta_0 \\ &= \int_{\mathbf{y}} \int_{\Theta} \log \frac{\phi(\theta|\mathbf{y})}{p(\theta)} \phi(\theta|\mathbf{y}) d\theta p(\mathbf{y}|\theta_0) d\mathbf{y} p(\theta_0) d\theta_0 + \mathcal{O}\left(\frac{1}{s^2}\right). \end{aligned} \quad (6)$$

Furthermore, the MAP estimate, $\hat{\theta}$, may be approximated using the true parameter θ_0 : $\hat{\theta} = \theta_0 + \mathbf{C}(\mathbf{y} - \mathbf{g}(\theta_0)) + \mathcal{O}_P\left(\frac{1}{s}\right)$, where $\mathbf{C} \propto s^{-1}$ is a matrix associated with the Jacobian $\mathbf{J}$, which is also independent of data. Therefore, given θ_0 , the influence of ϵ_e on $\phi(\theta|\mathbf{y})$ is asymptotically negligible. In other words, the truncated Gaussian $\phi(\theta|\mathbf{y})$ is asymptotically independent of $\mathbf{y}$ for a given value of θ_0 . The integration over $\mathbf{y}$ in (6) is removed to obtain

$$I = \int_{\Theta} \int_{\Theta} \frac{\phi(\theta|\theta_0)}{p(\theta)} \phi(\theta|\theta_0) d\theta p(\theta_0) d\theta_0 + \mathcal{O}\left(\frac{1}{s}\right). \quad (7)$$

□

Theorem 3 can be extended to the under-determined scenario as detailed in the following theorem.

Theorem 4

Under the assumption that $r < d$ and that the model output $\mathbf{g}(\theta, \xi)$ has continuous second derivatives, the expected information gain can be approximated by

$$I(\xi) = \int_{\Theta} D_{s|t}(\theta_0) p(\theta_0) d\theta_0 + \mathcal{O}\left(\frac{1}{s}\right), \quad (8)$$

with

$$\begin{aligned} D_{s|t}(\theta_0) &= \underbrace{\int_{S(\theta_0)} \left[-\frac{s^T \Sigma_{proj}^{-1} s}{2} - \nabla \log [p_{s,t}(\mathbf{0}, t)] s - \frac{s^T H_p(\mathbf{0}, t) s}{2} \right] \phi_s(s|t, \mathbf{y}) ds}_{D_1} \\ &\quad - \underbrace{\log \left[\int_T \int_{S(\mathbf{0}, t)} p(\mathbf{0}, t) \exp \left(-\frac{s^T \Sigma_{proj}^{-1} s}{2} \right) ds dt \right]}_{D_2} \end{aligned}$$

and

$$\phi_s(\mathbf{s}|\mathbf{t}, \mathbf{y}) = \frac{\tilde{p}(\mathbf{s}|\mathbf{t}, \mathbf{y})}{\int_{\mathbf{S}(\boldsymbol{\theta}_0)} \tilde{p}(\mathbf{s}|\mathbf{t}, \mathbf{y}) d\mathbf{s}},$$

where $H_p(\mathbf{s}, \mathbf{t})$ is the Hessian matrix of the logarithm of the transformed prior distribution, $\tilde{p}(\mathbf{s}|\mathbf{t}, \mathbf{y})$ is the pdf of a Gaussian distribution with mean $\mathbf{0}$ and covariance matrix Σ_{proj} , $\mathbf{S}(\boldsymbol{\theta}_0)$ is the range of parameter $\mathbf{s}$ at $\boldsymbol{\theta}_0$, and $\mathbf{S}(\mathbf{0}, \mathbf{t})$ is the range of parameter $\mathbf{s}$ at new coordinate $(\mathbf{0}, \mathbf{t})$.

Proof

$D_{\mathbf{s}|\mathbf{t}}$ is the conditional K-L divergence of the new parameter $\mathbf{s}$, given the value of parameter $\mathbf{t}$, which reads

$$D_{\mathbf{s}|\mathbf{t}}(\boldsymbol{\theta}_0) = \int_{\mathbf{S}(\boldsymbol{\theta}_0)} \log \left(\frac{\tilde{p}(\mathbf{s}, \mathbf{t}|\bar{\mathbf{y}})}{\tilde{p}(\mathbf{s}, \mathbf{t})} \right) \phi_s(\mathbf{s}|\mathbf{t}, \mathbf{y}) d\mathbf{s}, \quad (9)$$

where

$$\tilde{p}(\mathbf{s}, \mathbf{t}|\bar{\mathbf{y}}) = p(\mathbf{0}, \mathbf{t}|\bar{\mathbf{y}}) \exp \left[-\frac{\mathbf{s}^T \Sigma_{proj}^{-1} \mathbf{s}}{2} \right], \quad (10)$$

and

$$\tilde{p}(\mathbf{s}, \mathbf{t}) = p(\mathbf{0}, \mathbf{t}) \exp \left[\nabla \log [p(\mathbf{0}, \mathbf{t})] \mathbf{s} + \frac{\mathbf{s}^T H_p(\mathbf{0}, \mathbf{t}) \mathbf{s}}{2} \right]. \quad (11)$$

Substituting (10) and (11) into $\log \left[\frac{\tilde{p}(\mathbf{s}, \mathbf{t}|\bar{\mathbf{y}})}{\tilde{p}(\mathbf{s}, \mathbf{t})} \right]$ leads to

$$\log \left[\frac{\tilde{p}(\mathbf{s}, \mathbf{t}|\bar{\mathbf{y}})}{\tilde{p}(\mathbf{s}, \mathbf{t})} \right] = \log \left[\frac{p(\mathbf{0}, \mathbf{t}|\bar{\mathbf{y}})}{p(\mathbf{0}, \mathbf{t})} \right] - \frac{\mathbf{s}^T \Sigma_{proj}^{-1} \mathbf{s}}{2} - \nabla \log [p(\mathbf{0}, \mathbf{t})] \mathbf{s} - \frac{\mathbf{s}^T H_p(\mathbf{0}, \mathbf{t}) \mathbf{s}}{2}. \quad (12)$$

As with the derivations related to Equation (23) in ref. [24], we obtain

$$\log \left[\frac{p(\mathbf{0}, \mathbf{t}|\bar{\mathbf{y}})}{p(\mathbf{0}, \mathbf{t})} \right] = -\log \left[\int_{\mathbf{T}} \int_{\mathbf{S}(\mathbf{0}, \mathbf{t})} p(\mathbf{0}, \mathbf{t}) \exp \left(-\frac{\mathbf{s}^T \Sigma_{proj}^{-1} \mathbf{s}}{2} \right) d\mathbf{s} d\mathbf{t} \right] + \mathcal{O}_P \left(\frac{1}{s} \right). \quad (13)$$

Substituting Equation (13) into Equation (12), we have

$$\begin{aligned} \log \left[\frac{\tilde{p}(\mathbf{s}, \mathbf{t}|\bar{\mathbf{y}})}{\tilde{p}(\mathbf{s}, \mathbf{t})} \right] &= -\log \left[\int_{\mathbf{T}} \int_{\mathbf{S}(\mathbf{0}, \mathbf{t})} p(\mathbf{0}, \mathbf{t}) \exp \left(-\frac{\mathbf{s}^T \Sigma_{proj}^{-1} \mathbf{s}}{2} \right) d\mathbf{s} d\mathbf{t} \right] \\ &\quad - \frac{\mathbf{s}^T \Sigma_{proj}^{-1} \mathbf{s}}{2} - \nabla \log [p(\mathbf{0}, \mathbf{t})] \mathbf{s} - \frac{\mathbf{s}^T H_p(\mathbf{0}, \mathbf{t}) \mathbf{s}}{2} + \mathcal{O}_P \left(\frac{1}{s} \right). \end{aligned} \quad (14)$$

Equation (14) combined with Eq. (9) leads to the full expression of $D_{\mathbf{s}|\mathbf{t}}$ as shown in the statement of the theorem. The expression of I as an integral of $D_{\mathbf{s}|\mathbf{t}}$ against the prior pdf is derived in [24]. $\square$

Remark 2

The results in Theorem 4 involve nested integrals. Nevertheless, the integrals in $\mathbf{s}$ are extremely

efficient. D_2 can be rewritten as:

$$-\frac{1}{2} \log |\Sigma_{proj}| - \frac{r}{2} \log [2\pi] - \log \left[\int_T \int_{S(\mathbf{0}, \mathbf{t})} \tilde{p}(s|\mathbf{t}, \mathbf{y}) p(\mathbf{0}, \mathbf{t}) ds dt \right].$$

On the other hand, the range of the parameter s is constant in T in many practical applications. Hence, the expression of D_2 can be simplified to

$$-\frac{1}{2} \log |\Sigma_{proj}| - \frac{r}{2} \log [2\pi] - \log \left[\int_T p(\mathbf{0}, \mathbf{t}) dt \right] - \log \left[\int_{S(\theta_0)} \tilde{p}(s|\mathbf{t}, \mathbf{y}) ds \right].$$

3. NUMERICAL METHODS

In this section, we introduce the numerical methods used to estimate the expressions of Eq. (4) and (8).

First, in order to sample from $\phi(\theta|\theta_0)$, i.e. the pdf of a truncated Gaussian distribution, we draw samples from $\tilde{p}(\theta|\theta_0)$ and reject those outside the bounded domain Θ . See the following algorithm for details on how to compute the sample average of $\tilde{D}_{KL}$.

3.1. Rejection sampling algorithm

The approximated K-L divergence, $\tilde{D}_{KL}(\theta_0, \xi)$, can be estimated using the rejection sampling algorithm described in Algorithm 1. Since samples from a Gaussian distribution with mean θ_0 and covariance matrix Σ are needed, the additional numerical cost is negligible.

Algorithm 1 Rejection sampling algorithm for the estimation of $\tilde{D}_{KL}(\theta_0, \xi)$

```

1: count = 0;  $\tilde{D}_{KL} = 0$ 
2: for j = 1:N do
3:   Draw sample  $\theta_j \sim \mathcal{N}(\theta_0, \Sigma)$ 
4:   if  $\theta_j \in \Theta$  then
5:      $\tilde{D}_{KL} = \tilde{D}_{KL} + \log \frac{\tilde{p}(\theta_j|\mathbf{y}, \xi)}{p(\theta_j)}$ 
6:     count = count + 1
7:   else
8:     goto the beginning of the for loop
9:   end if
10: end for
11:  $\tilde{D}_{KL} = \frac{\tilde{D}_{KL}}{\text{count}} - \log \frac{\text{count}}{N}$ 

```

Subsequently, based on the regularity of $\tilde{D}_{KL}$, quadrature or Monte Carlo sampling is used to compute the expected information gain:

$$I \approx \sum_{k=1:M} w_k \tilde{D}_{KL}(\theta_{0k}),$$

where w_k are the weights associated with the k^{th} quadrature point or random sample θ_{0k} . See [24, 22, 23] for details.

Let $\tilde{D}_{s|t}$ be the Monte Carlo estimator of $D_{s|t}$. The steps leading to its estimation are shown in Algorithm 2. Note that in practice we use the simplified version of D_2 as explained in Remark 2. Using a method similar to that in the determined case, in the under-determined case we estimate I as a sample average of $\tilde{D}_{s|t}$

$$I \approx \sum_{k=1:M} w_k \tilde{D}_{s|t}(\theta_{0k}).$$

Algorithm 2 Rejection sampling algorithm for the estimation of $D_{s|t}(\theta_0, \xi)$

```

1: count = 0;  $\tilde{D}_1 = 0$ ;  $\tilde{D}_2 = 0$ 
2: for  $j = 1:N$  do
3:   Draw sample  $s_j \sim \mathcal{N}(0, \Sigma_{proj})$ 
4:    $\theta_j = \theta_0 + U s_j$ 
5:   if  $\theta_j \in \Theta$  then
6:      $\tilde{D}_1 = \tilde{D}_1 - \frac{s_j^T \Sigma_{proj}^{-1} s_j}{2} - \nabla \log [p_{s,t}(\mathbf{0}, t)] s_j - \frac{s_j^T H_p(\mathbf{0}, t) s_j}{2}$ 
7:     count = count + 1
8:   else
9:     goto the beginning of the for loop
10:  end if
11: end for
12:  $\tilde{D}_1 = \frac{\tilde{D}_1}{count}$ 
13:  $\tilde{D}_2 = -\log \left[ \frac{count}{N} \sqrt{2\pi}^r |\Sigma_{proj}|^{1/2} \right] - \log \left[ \int_T p(\mathbf{0}, t) dt \right]$ 
14:  $\tilde{D}_{s|t} = \tilde{D}_1 + \tilde{D}_2$ 

```

3.2. Spectral polynomial surrogates

In order to enable efficient sampling of model quantities appearing in the expressions for the expected information gain, e.g. Σ , a surrogate of these quantities is constructed first. Our approach is based on polynomial chaos expansions [42, 3, 10, 25] of quantities of interest (QoI), namely in terms of an orthogonal family of polynomials defined over a suitable probability space. This approach has been adopted in numerous studies focusing on uncertain chemical systems [20, 28, 41, 39, 38, 28, 29, 1] and only a brief outline is provided here.

The Polynomial Chaos Expansion (PCE) of a generic QoI, Q , may be expressed as:

$$Q(\boldsymbol{\eta}) = \sum_0^P Q_k \Psi_k(\boldsymbol{\eta}), \quad (15)$$

where $\boldsymbol{\eta} \equiv (\eta_1, \eta_2, \dots, \eta_d)$ is a d -dimensional vector of independent, identically-distributed random variables used to parameterize the uncertain model inputs, $P + 1$ is the number of terms retained in the expansion, $\{\Psi_k\}_{k=0}^\infty$ is an orthogonal basis of multidimensional polynomials spanning the space of square-integrable functions over the range of $\boldsymbol{\eta}$, and the Q_k 's are deterministic coefficients.

To determine the coefficients in (15), a non-intrusive spectral projection is applied [9, 21]. The procedure evaluates the QoIs for selected values of the germ $\boldsymbol{\eta}$ based on independent simulations. The corresponding discrete values of the QoIs are then used in conjunction with quadrature formulas to evaluate the Q_k 's. To this end, we rely on an adaptive pseudo-spectral (aPSP) approach [6, 5, 43],

utilizing nested sparse grids based on the Gauss-Kronrod-Patterson rule [30, 7]. The adaptation of the grid is based on a greedy refinement algorithm [8], using the variance surplus as error indicator. The refinement is carried out until a minimum threshold is obtained. The aPSP algorithm offers the advantages that internal aliasing is avoided by construction and that a maximal basis can be determined based on the samples available. The effectiveness of the aPSP approach has been recently demonstrated in a high-dimensional setting focusing on methane ignition [44] and ion chemistry [17].

4. DESIGN OF SHOCK TUBE EXPERIMENTS

4.1. Introduction

The chain branching reaction $\text{H} + \text{O}_2 \rightleftharpoons \text{OH} + \text{O}$ is one of the most important elementary reactions in the oxidation of hydrocarbon fuels and several experimental studies have been carried out to measure its reaction rate constant [13, 26, 31, 40, 45, 36, 16]. Recently, Hong et al. [13] have performed a series of shock-tube experiments aiming at improving the accuracy with which the parameters of this reaction are known. They considered mixtures of hydrogen and oxygen in argon diluent. Temperatures were varied from 1100 to 1500 K, the pressure was approximately 2 atm, and the hydrogen/oxygen mixture was rich with initial volumetric concentrations of oxygen fixed at 0.1% and hydrogen at 2.9% (1119 – 1250 K) or 0.9% (1250 – 1500 K), depending on the experiment. The authors measured the time dependent concentration of water vapor H_2O via tunable diode laser absorption near 2.5 μm . Their data set features results from 20 experiments.

The combustion of hydrogen/oxygen mixtures involves many other reactions in addition to the chain-branching reaction of interest. Since the purpose of the shock-tube experiments is to measure the rate of $\text{H} + \text{O}_2 \rightleftharpoons \text{OH} + \text{O}$, it is common to search for the mixture equivalence ratio, i.e. the relative concentration of H_2 and O_2 in the unreacted mixture, that maximizes the sensitivity of the quantity of interest to the rate constant of the reaction. In the case of the chain-branching reaction $\text{H} + \text{O}_2 \rightleftharpoons \text{OH} + \text{O}$, the rate of increase of the water concentration is most sensitive to the rate constant of the reaction for rich mixtures and shock-tube experiments are usually carried out in this limiting regime. The sensitivity analysis is carried out using the forward model (see Section 4.2) evaluated at the nominal values of the model parameters.

In the study by Hong et al. [13], an initial temperature was selected and the water vapor time dependent evolution was measured for each of the 20 experiments. They evaluate the forward model multiple times and the rate constant of the target reaction is adjusted until a good match between the time history of the experimental and model quantities is found. A set of 20 pairs of temperatures and rate constants is assembled and a least squares approach is used to obtain the rate parameters (pre-exponential and activation energy) that best fit the data for the Arrhenius form of the rate.

In this optimal experimental design (OED) exercise, we seek to measure the Arrhenius parameters of the chain-branching reaction $\text{H} + \text{O}_2 \rightleftharpoons \text{OH} + \text{O}$. The relevant parameters describing the experimental setup are the initial temperature T_0 and the initial volumetric concentration of hydrogen, $[\text{H}_2]_0$. The oxygen concentration is fixed at 0.1%, while the temperature and initial hydrogen concentration are allowed to vary as $1100 < T_0 < 1500 \text{ K}$ and $0.5 < [\text{H}_2]_0 < 5\%$.

Consistent with experimental practice, the observable is the maximum rate of increase of the water vapor concentration during ignition of the reactive mixture.

4.2. Forward model

The forward model requires the solution of a system of nonlinear ordinary differential equations describing the evolution of the species, temperature, and pressure in a zero-dimensional reactor. The combustion process in a shock-tube is described as an adiabatic, isochoric ignition event as detailed below.

Let Y_i indicate the mass fraction of the i th species, T indicate temperature, and p pressure. The differential equations describing the evolution of the thermochemical state of the mixture are

$$\frac{dY_i}{dt} = \frac{W_i}{\rho} \dot{\omega}_i, \quad (i = 1, \dots, N) \quad (16)$$

$$\frac{dT}{dt} = -\frac{1}{c_v} \sum_{k=1}^N \frac{dY_k}{dt} u_k, \quad (17)$$

$$\frac{dp}{dt} = p \left[\frac{1}{T} \frac{dT}{dt} + \left(\sum_{k=1}^N \frac{dY_k}{dt} W_k^{-1} \right) \left(\sum_{k=1}^N Y_k W_k^{-1} \right)^{-1} \right], \quad (18)$$

where ρ and c_v are the density and specific heat at constant volume of the mixture, while W_i , $\dot{\omega}_i$ and u_i is the molecular mass, the molar production rate, and the internal energy of species i , respectively. In Eq. (16) and Eq. (17), the molar production rate $\dot{\omega}_i$ is a function of the thermochemical state of the mixture, which is fully specified by T , p , and Y_k ($k = 1, \dots, N$). The molar production rate of species i is due to the contribution of R reactions as follows:

$$\dot{\omega}_i = \sum_{j=1}^R C_j [\nu_{ji}'' - \nu_{ji}'] \left[k_j^f \prod_{l=1}^N [X_l]^{\nu_{jl}'} - k_j^r \prod_{l=1}^N [X_l]^{\nu_{jl}''} \right], \quad (19)$$

where C_j is a parameter accounting for any third-body and/or pressure effects, ν_{ji}'' and ν_{ji}' are the reactant and product stoichiometric matrices for the coefficient of species i in reaction j , k_j^f and k_j^r are the forward and backward reaction rate constants, and $[X_l]$ is the molar concentration of species l . The forward rate constant is expressed in Arrhenius form as

$$k_j^f = A_j T^{b_j} \exp(-E_j/\mathcal{R}T). \quad (20)$$

In the Arrhenius form of the reaction rate constant, A_j is the pre-exponential factor, E_j is the activation energy, b_j is the temperature exponent, and $\mathcal{R}$ is the universal gas constant. The backward rate for a reversible reaction is related to the forward rate as $k_j^r = k_j^f / K_j^c$, where K_j^c is the equilibrium constant for reaction j and is a function of the temperature and thermodynamic properties of the species taking part in the reaction. More details on the ordinary differential equations (ODE) describing the ignition of a reactive mixture under adiabatic and isochoric conditions may be found in [37].

The thermodynamic properties and rate parameters required to evaluate the rate constants and species properties are taken from the combustion mechanism in [14]. The solution of the nonlinear

system of autonomous ODE (Eq. (16), Eq. (17), and Eq. (18)) is obtained from the initial thermochemical state using the open-source chemical kinetics library TChem [37], which relies on the CVODE stiff ODE integrator [12]. The library returns the time evolution of the thermochemical state vector during ignition.

We use additive Gaussian noise to represent observation errors in the experiments. We select the noise standard deviations to be 0.25 and 0.025, corresponding to 5% and 0.5% of the model output at the nominal parameters, respectively. While in this study we model experimental errors as additive noise, more complex error models, e.g. noisy experimental setup, will be considered in future work.

4.3. Prior distribution

When formulated in the Arrhenius form, the rate constant has three parameters, A , b , and E , as shown in Eq. (20) and related commentary. In this work, we assume that the temperature exponent, b , is zero. This assumption is consistent with the existing literature and recommendations [14].

The two model parameters A and E are described as independent random variables and non-informative priors are derived as follows. We assume A has the following relationship with a uniform random variable $\eta_1 \sim \mathcal{U}(-1, 1)$:

$$A = A_0 f^{\eta_1}, \quad (21)$$

where $f = 2$ and $A_0 = 1.04 \times 10^{14}$ is the recommended nominal pre-exponential value in the rate constant used in the forward model [14]. The range of E is chosen according to the following linear mapping:

$$E = E_0(1 + d_r \eta_2), \quad (22)$$

where $d_r = 0.1$ and $E_0 = 15,286$ cal/mol is the nominal value [14], and $\eta_2 \sim \mathcal{U}(-1, 1)$.

In this connection, we mapped the unknown parameters A and E to new independent random parameters η_1 and η_2 , which are distributed uniformly between -1 and 1 . We will use η_1 and η_2 in the computations. In the following, the symbol k_0 denotes the rate constant value computed with A_0 and E_0 . Given the range of the distributions of A and E , the rate constant varies from -44% to $+80\%$ of the k_0 value.

4.4. Convergence properties of the numerical approximations

Figure 2(a) shows the convergence of the expected information gain at a temperature $T_0 = 1500$ K and initial fuel concentration $[\text{H}_2]_0 = 5\%$. We use both sparse quadratures and Monte Carlo samples to carry out the integral in (3). For both types of numerical integrations, the Laplace method outperforms the nested Monte Carlo sampling in terms of number of samples by several orders of magnitude. Note that we reuse the samples in the outer and inner sample averages of the nested Monte Carlo sampler. Nevertheless, the number of likelihood evaluations, which accounts for most of the computational cost, is the square of the number of samples for the nested Monte Carlo.

As we intend to compute $\tilde{D}_{KL}$ for design parameter vectors of large dimensionality under uniform priors for the model parameters, the convergence properties of the approximations are assessed in the case of three experiments, each identified by a choice of initial temperature and hydrogen concentration, $\xi = [T_0, [\text{H}_2]_0]$. In Fig. 2(b), we report the results on the convergence of

the newly proposed method, which approximates the posterior pdf as a truncated Gaussian pdf. We draw 10^6 samples from the truncated Gaussian distribution to estimate $\tilde{D}_{KL}$. The integration against the prior in Eq. (4) is carried out by Gauss Legendre sparse quadratures. As shown in Fig. 2(b), the truncated Gaussian method converges to the right solution within less than 100 quadratures, i.e. around 100 times faster than the nested Monte Carlo sampler in terms of the number of samples. We should also note that the actual computing time is dominated by the number of likelihood pdf evaluations, when the number of samples becomes large. Most importantly, the expected information gain is underestimated by a factor of $\approx 50\%$ using the direct Laplace method, resulting in an unacceptable error in the context of optimal experimental design.

Fig. 3 shows a comparison between the Laplace method, the truncated Gaussian method and the nested Monte Carlo sampler under different values of measurement noise σ_e . The Laplace approximation can be arbitrarily accurate, when the value of σ_e decreases. When the σ_e is smaller than 0.025 the error of Laplace method is to the order of 1%. When σ_e is larger than 0.05, the error of Laplace approximation can be barely acceptable in engineering applications. On the other hand, the truncated Gaussian method is always consistent with the reference in the tested range of σ_e . Note that adequate number of quadratures and random samples are used to obtain negligible numerical integration error, i.e., 381 Gaussian quadratures are used in the truncated Gaussian and Laplace methods, 10^4 random samples are used in the nested Monte Carlo sampler. The statistical error bars represent 95% confidence intervals.

4.5. Single experiment

We consider a shock-tube experiment with initial temperature T_0 and initial hydrogen concentration $[\text{H}_2]_0$. Recall that for any single experiment, the observable is the maximum rate of increase of the water vapor concentration $[\text{H}_2\text{O}]$ during ignition. Although a single experiment does not allow the two model parameters A and E to be estimated and constitutes an underdetermined problem, it does allow us to calibrate the reaction rate constant k .

The surface of expected information gain computed using two levels of Gauss Legendre sparse quadratures for the model surrogate is shown in Fig. 4. The expected information gain grows as T_0 increases for any given value of $[\text{H}_2]_0$. The slope of the expected information gain surface along the T_0 axis is approximately constant over the design space. Conversely, the expected information gain is not very sensitive to variations in $[\text{H}_2]_0$ beyond a small region, where $[\text{H}_2]_0 < 1.0\%$ for any given value of temperature. As $[\text{H}_2]_0$ drops below 1.0%, the expected information gain decreases dramatically.

In the scenario of a single shock-tube experiment, the initial temperature T_0 has a more significant effect on the information gain than the initial fuel concentration $[\text{H}_2]_0$ does. The highest information gain is achieved at the highest T_0 of 1500 K and the highest $[\text{H}_2]_0$ of 5%, and the lowest information gain is obtained at the lowest T_0 of 1100 K and the lowest $[\text{H}_2]_0$ of 0.5%. In Fig. 4, the extrema of the information gain surface are marked as points A and B , respectively.

To validate the computed expected information gain, the two designs of the highest and lowest expected information gains, corresponding to designs A and design B in Fig. 4, are selected to generate synthetic data. A_0 and E_0 are set at their nominal values and the adaptive Markov Chain Monte Carlo (MCMC) [27] is applied to obtain the posterior samples, which are plotted in Fig. 5(a). The two sets of samples represent the posterior according to the experiment that has the highest (A)

and lowest (B) expected information gain. Note that the range of each axis is identical to the range in the design parameter space and that the axis of A is in log-scale. Thus, the prior is uniform along both axes of the two independent random variables and is also uniform in the plot. Qualitatively, the samples from the high information gain design (A) show much less scatter compared to the low information gain design (B). The values of the information gain D_{KL} , estimated using the posterior samples, are 2.6 and 0.7, respectively, and are close to the expected information gain values in Fig. 4.

In order to explain the trends in Fig. 5(a), let us consider the linear relation between $\log k$ and the pre-exponential factor and activation energy:

$$\log k = \log A - \frac{E}{RT}. \quad (23)$$

From the equation above, it is clear that, for a given temperature T , k is constant along a line in the plane $\log A$ - E and the straight line has a slope proportional to the temperature. Thus, the MCMC samples from the higher information gain case are distributed around a narrow strip along a line for which $k(1500) = \text{const}$ (recall that $T_0 = 1500$ K for design A). The posterior samples from the experiment with the lower information gain are distributed along another strip that is wider and oriented along a line for which $k(1100) = \text{const}$. As discussed above, the broader extent of the region occupied by the posterior samples indicates a less informative choice of experimental parameters.

The posterior samples of A and E from MCMC are used to produce the samples of k at each of the two temperatures using the Arrhenius form of the rate constant (Eq. (20)) and the results are shown in Fig. 5(b). The more efficient design with the higher information gain results in a very sharp distribution of k at 1500 K. At 1100 K the posterior pdf is less sharp, but improved against the uniform prior. Conversely, the design with the lower information gain has wide posterior distributions at both temperatures.

4.6. Multiple experiments

In this subsection, we address the case of optimal experimental design for multiple shock-tube experiments based on the expected information gain. We first consider the case where two shock-tube experiments are carried out. In this design scenario, the initial hydrogen concentration is fixed at 5%, i.e. the maximum concentration in the stipulated range, while the initial temperatures in the two experiments, T_{01} and T_{02} , are allowed to vary, so that the optimum expected information gain can be determined. In the second and last design scenario considered, we treat initial temperatures and initial hydrogen concentrations as design variables and show how the optimal expected information gain increases and saturates in a sequence of eight experiments.

The two-run experimental design problem features a two-dimensional design space and the observable is two-dimensional also. The surfaces of the expected information gain are shown in Figure 6 for two values of measurement noise, i.e. $\sigma_e = 0.25$ and 0.025.

In the case of large noise (Figure 6(a)), the I surface is rather flat and the maximum expected information gain is reached when both T_{01} and T_{02} are maximal and equal to the upper bound on the admissible temperature range. Interestingly, the information gain measure indicates that there is no penalty for repeating both experiments at the same temperature, but that low temperatures should

be avoided. Note that the lowest information gain is found when both temperatures are 1100 K, the lower bound on the temperature range.

The small noise scenario, shown in Fig. 6(b), differs significantly. As expected, the response surface is symmetric with respect to T_{01} and T_{02} . The expected information gain indicates clearly that repeating the two experiments at the same temperature should be avoided, and the optimal set of two experiments consists in one experiment at the maximum temperature $T_{01} = 1500$ K followed by a second experiment at a somewhat lower temperature $T_{02} = 1226$ K.

Comparing the low and high noise scenarios, it is apparent that the expected information gain is lower overall for the high noise case than the low noise case, regardless of the choice of design parameters T_{01} and T_{02} . Furthermore, as already noted, repeating the experiment with identical temperatures is not optimal when the noise level is low but is optimal when the noise level is high, provided that the temperature is set as high as possible.

We now consider the case of multiple experiments, for which the design parameters are optimized sequentially. We used the function *ga* in MATLAB to carry out the optimization, which requires only the values of the expected information gain. Where the gradient values are available, we can use gradient based stochastic optimization methods [34, 32]. In each experiment, the initial temperature and hydrogen concentration are allowed to vary to seek the maximum information gain. Once a maximum is found, we fix the experimental setup at the optimal value and search for the optimal design of the next experiment. In this way, we constrain the dimension of the searching space, although a global optimum is usually not obtained. The expected information gain associated with the optimal conditions is shown in Figure 7 as a function of the number of experiments. For comparison, the expected information gain for a sequence of experiments where the design conditions are chosen randomly is also shown. Data for the two levels of noise ($\sigma_e = 0.25$ and 0.025) are reported.

It is apparent that a lower number of experiments are needed to reach a given value of I if the optimal design strategy is used, as opposed to selecting experimental conditions randomly. This conclusion holds for both noise levels. For example, if the case where $\sigma_e = 0.025$, $I_{\text{opt}}(2) = I_{\text{random}}(4)$, i.e. two optimized experiments are equivalent to four experiments in which the conditions are chosen at random, according to a uniform distribution in the allowed range of the design space.

In the case of low noise, it appears that not much more information can be gained by repeating the experiment more than 4 times as I reaches an asymptote and saturates. Similar conclusions hold for the randomized experiments, even though the asymptotic behavior is shifted to a larger number of experiments in this case.

4.7. Validation of the DOE using legacy data

In order to validate the proposed optimal design approach further using real data, we obtained the observable, i.e. the maximum slope of the time history of $[\text{H}_2\text{O}]$, from the supplementary material in Ref. [13]. In that paper, the maximum slope of the measured water concentration is not given. The observable may be reconstructed from the values of the estimated rate constant at the various temperatures together with the initial experimental conditions of the fuel mixture. Using the rate constant and the initial conditions, we performed multiple simulations to reconstruct the observable,

which we then used for the analysis described below. Therefore it should be noted that the data used here could be different from the raw data in [13]. The estimated data are shown in Fig. 8

First, we apply MCMC to the data to obtain a posterior sample of the parameters A and E . Furthermore, MCMC provides a sample of a hyper-parameter, the variance σ_ϵ^2 of the random experimental noise ϵ . The posterior samples are shown in Fig. 9 (a) and the pdf of the variance σ_ϵ^2 in (b). The plot shows that the posterior samples has a high correlation along the diagonal direction like the previous MCMC sample in Fig. 5 and the posterior distribution is not centered at the nominal values of $A/A_0 = 1$ and $E/E_0 = 1$. The sample points are slightly biased to the lower left of the center and the sample has the highest probability when $A/A_0 = 0.88$ and $E/E_0 = 0.98$. Both values are smaller than the parameter values suggested in [13]. One reason for the difference in the parameter estimates is that we have reconstructed a scalar observable to represent the time series, which is possibly low-rank. Nevertheless, the differences between our estimates and the nominal values are within the region of high confidence [13].

As shown in Fig. 9 (b), the estimated pdf of the noise variance σ_ϵ^2 has a mean value of 0.0211 s^{-2} and standard deviation of 0.145 s^{-1} . The estimated mean of the noise variance is smaller than that stemming from the high noise design scenario ($\sigma_\epsilon = 0.25 \text{ s}^{-1}$) used for the experimental design, but larger than that in the low noise scenario ($\sigma_\epsilon = 0.025 \text{ s}^{-1}$). The variance of the low noise is outside the high probability range of the estimated noise variance, while the high noise variance is comparable with the estimated experimental noise. Thus the high noise design scenario is more plausible. Fig. 9 (c) shows the two empirical probability densities of the rate constant k constructed from the MCMC posterior samples of parameters A and E . The normalized standard deviations are 5% and 2% at 1100 K and 1500 K, respectively.

Next, we use the reconstructed data-the maximum slope of $[\text{H}_2\text{O}]$ time history at the lowest (1091 K) and highest (1527 K) temperatures to estimate the parameters A and E separately, and compare the results under both experimental designs. The high noise design scenario ($\sigma_\epsilon = 0.25 \text{ s}^{-1}$) is used for MCMC. The MCMC samples of the parameters A and E and the probability densities of the rate constant k are shown in Fig. 10. The MCMC samples computed from the reconstructed legacy data are distributed along the diagonal direction in a very similarly way to those obtained using synthetic data in Fig. 5 (a). The high temperature design has a more concentrated posterior pdf than the low temperature design, and the computed values of the information gain D_{KL} are 2.7 and 0.9, respectively. The estimated values of the information gain D_{KL} are comparable with the values of the expected information gain I in Fig. 4, although the information gain values from the posterior samples are slightly smaller. The shapes of the empirical probability densities of the rate constant k in Fig. 10 (b) resemble those in Fig. 5 (b), although the MAP estimate is negatively biased with respect to the nominal value. Furthermore, we applied MCMC to the two-run designs: design one consists of two runs at the high temperature, while design two consists of one run at the high temperature followed by one at the low temperature. As we have seen in the experimental design section, the former design has a higher value of the expected information gain than the latter when the high noise model ($\sigma_\epsilon = 0.25 \text{ s}^{-1}$) is used. Using the reconstructed legacy data, the computed D_{KL} values for the two designs are 3.1 and 2.7 respectively, which are consistent with the values of expected information gains as shown in Fig. 6.

The analysis presented in this Section illustrates the validity and robustness of our design approach. Additionally, it is worth pointing out that the design would be erroneous if the prior is chosen incorrectly. Then, a careful data analysis and a redefinition of the prior are necessary.

5. CONCLUSION

The estimation of rate parameters in the field of combustion kinetics requires careful design of shock tube experiments. In this work, a new optimal experimental design framework [23, 24] is further developed and applied to the selection of optimal conditions for measuring the rate parameters of the chain branching reaction $\text{H} + \text{O}_2 \rightleftharpoons \text{OH} + \text{O}$. More precisely, by approximating the posterior pdf using a truncated Gaussian distribution, we have obtained a robust estimator of the expected information gain when the priors have compact supports, i.e. uniform priors in our study. This is usually the case in the study of combustion kinetics. The truncated Gaussian can be obtained by using an approach similar to the Laplace method. We have shown that the proposed method is accurate and computationally much more efficient than nested Monte Carlo method. We have also shown that the optimal design differs depending on the magnitudes of the observational noises. We obtained the most informative data from a single experiment when the temperature was at a maximum, but the situation becomes more complex when multiple experiments are planned. The information that can be extracted from the data saturates as a logarithmic function of the number of experiments, and few experiments are needed when they are conducted under the optimal experimental design conditions. Future work will focus on developing more realistic error models describing the noisy experimental setup and on designing practical laser-based diagnostics for combustion applications. We also plan to undertake further analyses treating the time history of the species concentrations during the ignition event as observables.

ACKNOWLEDGMENTS

The research reported in this publication was supported by King Abdullah University of Science and Technology (KAUST) and the Center for Uncertainty Quantification in Computational Science and Engineering funded by the Strategic Research Initiative (SRI).

REFERENCES

1. A. Alexanderian, O. P. Maître, H. N. Najm, M. Iskandarani, and O. M. Knio. Multiscale stochastic preconditioners in non-intrusive spectral projection. *Journal of Scientific Computing*, 50(2):306–340, 2012.
2. A. C. Atkinson, A. N. Donev, and R. D. Tobias. *Optimum Experimental Designs, with SAS*. Oxford University Press, 2007.
3. R. H. Cameron and W. T. Martin. The orthogonal development of non-linear functionals in series of Fourier-Hermite functionals. *Annals of Mathematics*, 48(2):385–392, 1947.
4. K. Chaloner and I. Verdinelli. Bayesian experimental design: a review. *Statistical Science*, 10(3):273–304, 1995.
5. P. R. Conrad and Y. M. Marzouk. Adaptive Smolyak pseudospectral approximations. *SIAM Journal on Scientific Computing*, 35(6):A2643–A2670, 2013.

6. P. G. Constantine, M. S. Eldred, and E. T. Phipps. Sparse pseudospectral approximation method. *Computer Methods in Applied Mechanics and Engineering*, 229:1–12, 2012.
7. T. Gerstner and M. Griebel. Numerical integration using sparse grids. *Numerical Algorithms*, 18(3-4):209–232, 1998.
8. T. Gerstner and M. Griebel. Dimension-adaptive tensor-product quadrature. *Computing*, 71(1):65–87, 2003.
9. R. Ghanem, J. Red-Horse, and A. Sarkar. The modal properties of a space-frame with localized system uncertainties. In *8th ASCE Specialty Conference of Probabilistic Mechanics and Structural Reliability (PMC2000)*, 2000.
10. R. G. Ghanem and P. D. Spanos. *Stochastic Finite Elements: A Spectral Approach*. Springer New York, 1991.
11. J. Ginebra. On the measure of the information in a statistical experiment. *Bayesian Analysis*, 2(1):167–211, 2007.
12. A. Hindmarsh, P. Brown, K. Grant, S. Lee, R. Serban, D. Shumaker, and C. Woodward. SUNDIALS: Suite of nonlinear and differential/algebraic equation solvers. *ACM Transactions on Mathematical Software (TOMS)*, 31(3):363–396, 2005.
13. Z. Hong, D. F. Davidson, E. A. Barbour, and R. K. Hanson. A new shock tube study of the $\text{H} + \text{O}_2 \rightarrow \text{OH} + \text{O}$ reaction rate using tunable diode laser absorption of H_2O near $2.5 \mu\text{m}$. *Proceedings of the Combustion Institute*, 33:309–316, 2011.
14. Z. K. Hong, D. F. Davidson, and R. K. Hanson. An improved H_2/O_2 mechanism based on recent shock tube/laser absorption measurements. *Combustion and Flame*, 158(4):633–644, 2011.
15. X. Huan and Y. Marzouk. Simulation-based optimal Bayesian experimental design for nonlinear systems. *Journal of Computational Physics*, 232:288–317, 2013.
16. S. M. Hwang, S. O. Ryu, K. J. De Witt, and M. J. Rabinowitz. High temperature rate coefficient measurements of $\text{H}+\text{O}_2$ chain-branching and chain-terminating reaction. *Chemical Physics Letters*, 408(1-3):107–111, 2005.
17. D. Kim, F. Rizzi, K. W. Cheng, J. Han, F. Bisetti, and O. M. Knio. Uncertainty quantification of ion chemistry in lean and stoichiometric homogenous mixtures of methane, oxygen, and argon. *Combustion and Flame*, 162(7):2904–2915, 2015.
18. S. Kullback. *Information Theory and Statistics*. Wiley, 1959.
19. S. Kullback and R. A. Leibler. On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86, 1951.
20. O. P. Le Maître, H. N. Najm, P. P. Pébay, R. G. Ghanem, and O. M. Knio. Multi-resolution-analysis scheme for uncertainty quantification in chemical systems. *SIAM Journal on Scientific Computing*, 29(2):864–889, 2007.
21. O. P. Le Maître, M. T. Reagan, H. N. Najm, R. G. Ghanem, and O. M. Knio. Stochastic projection method for fluid flow - ii. random process. *Journal of Computational Physics*, 181(1):9–44, 2002.
22. Q. Long, M. Motamed, and R. Tempone. Fast Bayesian optimal experimental design for seismic source inversion. *Computer Methods in Applied Mechanics and Engineering*, 291:123–145, 2015.
23. Q. Long, M. Scavino, R. Tempone, and S. Wang. Fast estimation of expected information gains for Bayesian experimental designs based on Laplace approximations. *Computer Methods in Applied Mechanics and Engineering*, 259(0):24–39, 2013.
24. Q. Long, M. Scavino, R. Tempone, and S. Wang. A Laplace method for under-determined Bayesian optimal experimental design. *Computer Methods in Applied Mechanics and Engineering*, 285:849–876, 2015.
25. O. L. Maitre and O. Knio. *Spectral Methods for Uncertainty Quantification: with Applications to Computational Fluid Dynamics*. Springer Series in Scientific Computation. Springer, 2010.
26. D. A. Masten, R. K. Hanson, and C. T. Bowman. Shock-tube study of the reaction $\text{H} + \text{O}_2 \rightarrow \text{OH} + \text{O}$ using OH laser-absorption. *Journal of Physical Chemistry*, 94(18):7119–7128, 1990.
27. N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of state calculations by fast computing machines. *Journal of Chemical Physics*, 21(6):1087–1092, 1953.
28. H. N. Najm, B. J. Debusschere, Y. A. Marzouk, S. Widmer, and O. P. Le Maître. Uncertainty quantification in chemical systems. *International Journal for Numerical Methods in Engineering*, 80(6-7):789–814, 2009.
29. H. N. Najm and M. Valorani. Enforcing positivity in intrusive PC-UQ methods for reactive ODE systems. *Journal of Computational Physics*, 270:544–569, 2014.
30. T. Patterson. The optimum addition of points to quadrature formulae. *Mathematics of Computation*, 22(104):847–856, 1968.
31. A. N. Pirraglia, J. V. Michael, J. W. Sutherland, and R. B. Klemm. A flash-photolysis shock-tube kinetic-study of the H atom reaction with O_2 : $\text{H}+\text{O}_2 \rightleftharpoons \text{OH} + \text{O}$ ($962 \text{ K} \leq t \leq 1705 \text{ K}$) and $\text{H}+\text{O}_2 + \text{Ar} \rightarrow \text{HO}_2 + \text{Ar}$ ($746 \text{ K} \leq t \leq 987 \text{ K}$). *Journal of Physical Chemistry*, 93(1):282–291, 1989.
32. B. T. Polyak and A. B. Juditsky. Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization*, 30:838, 1992.
33. F. Pukelsheim. *Optimal Design of Experiments*. SIAM, 2006.

34. H. Robbins and S. Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, 22:400, 1951.
35. K. J. Ryan. Estimating expected information gains for experimental designs with application to the random fatigue-limit model. *Journal of Computational and Graphical Statistics*, 12(3):585–603, 2003.
36. S. O. Ryu, S. M. Hwang, and M. J. Rabinowitz. Shock-tube and modeling study of the $\text{H}+\text{O}_2\rightarrow\text{OH}+\text{O}$ reaction over a wide-range of composition, pressure, and temperature. *Journal of Physical Chemistry*, 99(38):13984–13991, 1995.
37. C. Safta, H. Najm, and O. Knio. Tchem - a software toolkit for the analysis of complex kinetic models. *Sandia Report, SAND2011-3282*, 2011.
38. D. A. Sheen and H. Wang. Combustion kinetic modeling using multispecies time histories in shock-tube oxidation of heptane. *Combustion and Flame*, 158(4):645–656, 2011.
39. D. A. Sheen and H. Wang. The method of uncertainty quantification and minimization using polynomial chaos expansions. *Combustion and Flame*, 158(12):2358–2374, 2011.
40. K. S. Shin and J. V. Michael. Rate constants for the reactions $\text{H}+\text{O}_2\rightarrow\text{OH}+\text{O}$ and $\text{D}+\text{O}_2\rightarrow\text{OD}+\text{O}$, over the temperature-range 1085–2278 K by the laser photolysis shock-tube technique. *Journal of Chemical Physics*, 95(1):262–273, 1991.
41. M. Villegas, F. Augustin, A. Gilg, A. Hmadi, and U. Wever. Application of the Polynomial Chaos Expansion to the simulation of chemical reactors with uncertainties. *Mathematics and Computers in Simulation*, 82(5):805–817, 2012.
42. N. Wiener. The homogeneous chaos. *American Journal of Mathematics*, 60:897–936, 1938.
43. J. Winokur, P. Conrad, I. Sraj, O. Knio, A. Srinivasan, W. C. Thacker, Y. Marzouk, and M. Iskandarani. A priori testing of sparse adaptive polynomial chaos expansions using an ocean general circulation model database. *Computational Geosciences*, 2013.
44. J. Winokur, D. Kim, F. Bisetti, O. L. Maître, and O. Knio. Sparse pseudo spectral projection methods with directional adaptation for uncertainty quantification. *Journal of Scientific Computing*, submitted, 2014.
45. T. Yuan, C. Wang, C. L. Yu, M. Frenklach, and M. J. Rabinowitz. Determination of the rate coefficient for the reaction $\text{H}+\text{O}_2\rightarrow\text{OH}+\text{O}$ by a shock-tube laser-absorption detailed modeling study. *Journal of Physical Chemistry*, 95(3):1258–1265, 1991.

Figure 1. The probability density function of a truncated Gaussian (colored lines) defined on $[-1, 1] \times [-1, 1]$, and the Gaussian distribution (black lines) with the same MAP estimate visualized in the domain of $[-2, 2] \times [-2, 2]$. See text for details of the parameters of the distributions.

Figure 2. (a) Convergence of the expected information gain of one experiment with $T_0 = 1500K$, $[H_2]_0 = 5\%$, and $\sigma_e = 0.25$. (b) Convergence of the expected information gain of three experiments with $\xi_1 = [1500, 5\%]^T$, $\xi_2 = [1100, 0.5\%]^T$, $\xi_3 = [1500, 0.5\%]^T$ and $\sigma_e = 0.25$. The statistical error bars represent 95% confidence intervals.

Figure 3. Comparison of the expected information gains computed under different values of measurement noise σ_e using the Laplace method, the truncated Gaussian method and the nested Monte Carlo sampler.

Figure 4. The expected information gain of a single experiment with $\sigma_e = 0.25$. Note that the ranges of T_0 and $[H_2]_0$ are normalized to $[-1, 1]$.

Figure 5. (a) Posterior samples of the parameters computed by MCMC. Red color indicates design A. Blue color indicates design B. (b) Pdf of $\log(k/k_0)$ conditioning on different temperatures and designs.

(a)

(b)

Figure 6. Expected information gain for the two-run experimental design problem. In both experiments, $[H_2]_0 = 5\%$. (a): $\sigma_e = 0.25$ and (b): $\sigma_e = 0.025$. The ranges of T_{01} and T_{02} are normalized to $[-1, 1]$.

Figure 7. Expected information gain for a sequence of optimized experiments using the optimal selection strategy or a randomized design strategy.

Figure 8. Maximum slopes extracted from data in [13].

Figure 9. (a) MCMC posterior sample (red) of A and E based on real data and the sample point (blue) of the highest probability. (b) Probability density of the variance of the experimental noise sampled by MCMC. (c) Posterior probability densities of the rate constant at the low (blue) and high (red) temperatures.

Figure 10. (a) Posterior samples of A and E based on real data from a single experiment: low (blue) and high (red) temperature designs. (b) The probability densities of k from the posterior samples at 1100 K and 1500 K.