

Multi-fidelity Design Optimization of Solid Oxide Fuel Cells Using a Bayesian Feature Enhanced Stochastic Collocation

Wei W. Xing², Akeel A. Shah^{1,*}, Guohao Dai³, Ziyang Zhang⁴, Ting Guo³, Hong Qiu³, Puiki Leung¹, Qian Xu⁵, Xun Zhu¹, Qiang Liao¹

¹*Key Laboratory of Low-grade Energy Utilization Technologies and Systems, Chongqing University, Ministry of Education, Chongqing 400044, PR China*

²*School of Integrated Circuit Science and Engineering, Beijing University of Aeronautics and Astronautics, Beijing, 100191, China*

³*College of Mechatronics and Control Engineering, Shenzhen University, Shenzhen, 518060, China*

⁴*College of Computer Science and Software Engineering, Shenzhen University, Shenzhen, 518060, China*

⁵*Institute for Energy Research, Jiangsu University, Zhenjiang 212013, China*

**Author to whom correspondence should be addressed. E-mail: akeelshah@cqu.edu.cn; Tel: +44 7979465001*

Abstract

We develop a multi-fidelity surrogate modelling approach to replace the complex and costly physics-based computer models that are often used in the optimization of solid oxide fuel cell (SOFC) performance, or the simplified models that are used in lieu of complex models. We extend multi-fidelity stochastic collocation through a feature engineering step, and eliminate the requirement for the exact low-fidelity output at the inference stage. In contrast to previous approaches, the surrogate model we develop provides detailed spatial information, rather than one or more scalar outputs. This allows for the incorporation of such information into the objective of the optimization study, with the flexibility to choose from more than one objective, such as a minimum, maximum or average. Furthermore, the detailed spatial information can be used for general design purposes, such as ensuring uniformity in reactant and potential distributions. From the results on a 3-d SOFC model, we demonstrate highly accurate predictions of multiple spatially distributed quantities at up to spatial 250,000 locations. The results are superior to state-of-the-art multi-fidelity approaches, particularly for low numbers of high fidelity training points. We use the surrogate model to optimize the SOFC performance with respect to different objectives (including with nonlinear constraints and multiple objectives), with results that are accurate and are obtained in a fraction of the time required for the full model.

Keywords: Solid oxide fuel cell optimization; spatially distributed; multi-fidelity; stochastic collocation; feature engineering; Bayesian.

1 Introduction

Fuel cells (FCs) are expected to play a major role in the future energy landscape, for standalone stationary power applications and for electric vehicle propulsion. FCs convert chemical energy into electricity [1, 2]

generating benign products. Solid oxide fuel cells (SOFC) are highly attractive due to their fuel flexibility, range of different geometries, and high temperatures, allowing for combined heat and power. The high temperatures lead to low costs since they eliminate the need for a catalyst and allow for fuels to be reformed internally in the anode. Modeling and simulation are important tools in the development of fuel cells, providing fundamental insights and lowering the costs associated with laboratory investigations, particularly for design and optimization [3,4].

Detailed fuel cell models, however, involve a complex system of nonlinear partial differential equations in a 3-d geometry comprising multiple components. They are often referred to as computational fluid dynamics (CFD) approaches, although in the general case they contain other laws such as charge and reactant conservation, and incorporate electrochemical reactions. Detailed reviews of SOFC and polymer electrolyte membrane (PEM) fuel cell modelling, including the various approaches and modelling assumptions, can be found in [5–7]. The time costs for solving such models, using finite differences, finite volumes or finite elements along with time-stepping schemes, can be prohibitive for applications that require many results at different parameter/input values, as for example, in optimization, sensitivity analyses and parameter identification. For small number of parameters [8] the time costs may be acceptable, but the number of computer model runs scales exponentially with the number of parameters, rendering complex models infeasible in many cases.

The most frequently employed strategy to circumvent this problem is to use simplified models. For example, Salva et al. [9] used a 1-d analytical model for optimizing PEM fuel cell performance with respect to operating conditions, while Hasanien et al. used a 0-d model for PEM fuel cell model parameter identification [10]. Pourkiaei et al. performed a multi-objective optimization of a tubular SOFC using a 0-d model combined with a genetic algorithm, using the electrical power and exergy destruction as objectives [11]. A similar study was performed by Mojaver et al. [12] for a methane-fueled SOFC, using the electrical energy efficiency, CO₂ emissions and electrical exergy efficiency as the objectives. A combined Taguchi/AHP/TOPSIS method was used for the optimization. A review of SOFC model parameter identification using similar models can be found in [13], based on a variety of optimization techniques.

Such approaches, however, will inevitably introduce crude simplifications that can lead to unsatisfactory results. Guo et al. used a detailed 3-d SOFC model that included charge, mass, momentum and heat balances. The model was developed in [14] and coupled to a mechanical-failure model in [15] to study the impact of the interconnector (IC) structure on the electrical performance and mechanical stability. Although the authors stated that their study was an optimization, in common with many other such studies, in actual fact it was a parametric study. Abdullah and Liu [16] performed a microstructural optimization for a low-temperature SOFC by coupling macroscopic and pore-scale models. To make the optimization feasible, the macroscopic model was restricted to 1-d.

An alternative approach is to use surrogate model approximations, usually based on machine learning

or multi-fidelity methods. In the context of electrochemical energy technologies, machine learning methods have been used extensively as surrogate models for sensitivity analyses [17], optimization [18, 19], fuel cell degradation [20], monitoring fuel cell performance [21], battery health monitoring [22, 23], inverse parameter estimation [24] and model identification [25]. Examples of the methods employed are linear regression [26], artificial neural networks (ANNs) [27–29], deep learning networks (DNNs) [30, 31] and Gaussian process (GP) models [21, 32]. ANNs [28, 29] and DNNs [33] have proven to be particularly popular given the availability of open-source and commercial codes.

In [34], Xu et al. assessed 19 different machine learning approaches, including ANNs, Boltzmann machines and support vector regression trained on data from a CFD model, which were used in combination with a genetic algorithm to optimize the power density output of a SOFC. Wang et al. [35] developed a response surface model based on linear regression using data from a finite element model to optimize geometric parameters for minimizing the long-term creep of a SOFC. Subotic et al., instead conducted an SOFC performance optimization using experimental polarization and impedance data, used for an ANN (although it was actual a parametric study) [29]. In general, however, ANNs require large volumes of data, which can be a hindrance when the complex computer model used to generate the data requires hours or even days to complete one simulation.

Almost exclusively, as in the examples cited above, machine-learning surrogate models have been used for scalar outputs, such as a peak power density or an efficiency. Ideally, however, the output of the surrogate model should contain the richness of information (spatial distributions of reactant concentrations, ionic and electronic potentials, temperature, and so on) found in the original computer model, in order to make fully-informed design and optimization decisions. For example, it is important to ensure uniform distributions of the reactants, temperature, flow velocity and overpotentials in the electrodes, which optimizes the performance and efficiency, and mitigates against degradation. This cannot be achieved by relying on a single or multiple scalar summaries from the model. Spatio-temporal information can also allow for optimization with respect to multiple objectives, and these objectives do not necessarily need to be decided beforehand. For example, we may use a minimum, maximum, average or some other functional summary. In this paper we therefore focus on the approximation of spatial outputs. Surrogate models for spatial or spatio-temporal outputs are particularly challenging due to the sizes of the output spaces. It is no surprise therefore that there have been few attempts to develop such models for fuel cells (and batteries), with exceptions being the reduced-dimensional GP approach of Shah and co-workers [36, 37], and a similar approach combined with a coregionalisation model for multi-task learning (multiple outputs learned together) [38].

Another paradigm for developing surrogates is multi-fidelity modelling, in which data from models of low fidelity (low accuracy and computational cost) is leveraged to aid the learning of outputs at a higher fidelity [39]. The great advantage of such approaches is that they typically require fewer runs of the high-fidelity model

than are required for pure machine learning approaches, which rely entirely on high-fidelity data. When the number of system parameters (those of interest in the application) is high, collecting the data for training a pure machine learning model can itself become an issue.

In [40], Liu et al. developed a dual-segmented approach for a PEM fuel cell stack in which interpolation was used to bridge the gap between coarse- and fine-level segmentations. The authors claimed that the high fidelity result was equivalent to a finite-element simulation in terms of accuracy, but presented no direct evidence of this. Voskuilen et al. [41] surveyed various battery models, classifying them in terms of the fidelity, which was typically based on the number and extent of simplifications (of the underlying physics). The authors did not, however, combine these models in any way, in spite of the title of their paper. More conventional and sophisticated multi-fidelity approaches have not been applied to fuel cells. Xing et al. [42] developed a residual Gaussian process model that was tested in only one example on an SOFC finite-element model.

Many multi-fidelity models involve building a surrogate model (usually based on machine learning) at each fidelity. The linear autoregressive (LAR) multi-fidelity model [43] (see also [44]) assumes a linear relationship or parametric mapping between the different fidelity levels. Nonlinear AR (NARGP) [45, 46] places GP priors over the unknown mappings between fidelities, allowing for greater model flexibility. For spatial and spatio-temporal problems, the major shortcoming of this method is that the solution at a fidelity f is used as an input to the GP model for fidelity $f + 1$. This presents an enormous challenge in terms of learning the model hyperparameters. Furthermore, tractability is lost and expensive sampling or variational methods are required for training and inference.

In multi-fidelity stochastic collocation (SC) [47], a greedy algorithm is used to identify inputs for a low number of high-fidelity simulations, while the low-fidelity results are used to approximate the coefficients in a high-fidelity SC approximation. SC and related approaches require out-of-sample low-fidelity model results in order to make predictions, which may be justified in some design contexts but renders them unusable for optimization, sensitivity analysis and uncertainty quantification when the low fidelity model is itself moderately expensive, as is frequently the case.

Greedy NAR [48] provides a bridge between NARGP and SC, appealing to the advantages of both methods. The high-fidelity model is written as a linear map of the low-fidelity model in an abstract (feature) space, with the feature map specified implicitly by marginalising over the weight matrix appearing in the map and using kernel substitution. Active learning is then used to efficiently identify the inputs for the data generation at all fidelities greater than the lowest. Greedy NAR outperformed LAR, NARGP and SC, but was subsequently outperformed by ResGP [42], which uses GP models for the residuals between fidelities, leading to a tractable posterior and scalable hyperparameter learning.

In this paper we develop a multi-fidelity surrogate model for a SOFC that is capable of approximating detailed spatial information and can be used for optimization purposes. We extend the SC approach to make

it more accurate and apply it to a detailed 3-d SOFC model that includes mass, momentum and charge balances in the channels, electrodes and electrolyte. The model capacity of SC is enhanced by introducing a novel feature engineering step, and the requirement of low-fidelity data for predictions is eliminated by using a GP model. The approach is shown to yield highly accurate predictions of multiple spatially distributed quantities, including the electrolyte current density, ionic potential, volumetric anode current density and cathode overpotential, at up to 250,000 spatial locations. We note that these quantities are chosen for the purposes of illustration, and that there is no restriction on the type of field output, which can be related to the electrochemical performance, flow field, reactant concentrations or temperature. Indeed, SC has been successfully applied to frequency-modulated heat driven cavity flows [49], acoustics [50], molecular dynamics [51], airfoil design [52], discrete-space probability density evolution [53], and turbulence modelling [54].

The results are compared to LAR, NARGP, ResGP and Greedy NAR, demonstrating superior performance, particularly for low numbers of high fidelity training points. We also compare the computational complexities of these methods with that of ours. We then apply the surrogate model to the optimization of the SOFC performance, considering different objectives, objectives with nonlinear constraints and multiple objectives. The level of detail furnished by the surrogate model we develop allows for flexibility in deciding the objectives a-posteriori, in contrast to approaches that focus on scalar targets. We reiterate the novel aspects of our work: (i) as far as we are aware, we develop the first genuine multi-fidelity approach for SOFC that requires only a few runs of a complex (time-intensive), physics-based model to generate data; (ii) we are able to predict spatial outputs, allowing for more informed design and optimization decisions, with some flexibility to change objectives a-posteriori; (c) the method we develop significantly extends the original stochastic collocation formulation using feature enhancement and, moreover, replaces the expensive low-fidelity simulation required for prediction with an extremely rapid GP approximation.

2 Physics-based high and low fidelity models of a SOFC

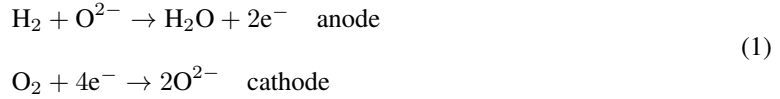
We consider a detailed steady-state, isothermal 3-d solid oxide fuel cell model, which was implemented in the COMSOL Multiphysics software package ¹. We present only the essential details of the model and refer the reader to the documentation for the full set of equations and parameter values. A two-fidelity setting was considered. COMSOL Multiphysics uses the finite element method to solve the governing equations. Additionally, a V cycle geometric multigrid is used in this particular implementation. The domain comprises single anode and cathode channels, the electrodes and a ceramic separator, depicted in Fig. 1. The spatial variable is denoted $\mathbf{x} = (x, y, z)^T$. The dimensions of the channels are 1 cm \times 0.5 mm \times 0.5 mm in the $x \times y \times z$ directions. The dimensions of the electrodes are 1 cm \times 1 mm \times 0.1 mm, and the dimensions of

¹<https://www.comsol.com/model/current-density-distribution-in-a-solid-oxide-fuel-cell-514>

the separator are $1 \text{ cm} \times 1 \text{ mm} \times 0.1 \text{ mm}$. The geometry is constructed such that the anode inlet is placed at $x = 0$, while the cathode inlet is placed at $x = 1 \text{ cm}$.

The model is based on hydrogen and air at the anode and cathode, respectively. Humidified hydrogen (hydrogen and water vapour) is supplied at the anode inlet, while humidified air (oxygen, water vapour and nitrogen) is supplied at the cathode inlet. The model includes charge balances in the electrolyte and electrodes (based on Ohm's law with fixed conductivity values) to calculate the electronic and ionic potentials. The incompressible Navier-Stokes equations along with continuity are used for the gas channel flow, while Brinkman's equation is used for the porous electrode flow. Mass balances of all species are solved in the channels and electrodes, incorporating Maxwell-Stefan diffusion and bulk motion by convection for the multi-component transport.

The electrochemical reactions are:



Butler-Volmer laws are assumed for these reactions, with current densities (Am^{-2}) (and associated reaction rates via stoichiometry) given by:

$$\begin{aligned} i_a &= i_{0,a} \left(\frac{P_{\text{H}_2}}{P_{\text{H}_2,\text{ref}}} \exp\left(\frac{F\eta}{2RT}\right) - \frac{P_{\text{H}_2\text{O}}}{P_{\text{H}_2\text{O},\text{ref}}} \exp\left(-\frac{3F\eta}{2RT}\right) \right) \quad \text{anode} \\ i_c &= i_{0,c} \left(\exp\left(\frac{3.5F\eta}{RT}\right) - \frac{P_{\text{O}_2}}{P_{\text{O}_2,\text{ref}}} \exp\left(-\frac{F\eta}{2RT}\right) \right) \quad \text{cathode} \end{aligned} \quad (2)$$

in which $i_{0,a}$ and $i_{0,c}$ are the anode and cathode exchange current densities (Am^{-2}), P_j and $P_{j,\text{ref}}$ denote the partial pressures and reference pressures of species j (Pa), F is Faraday's constant (Cmol^{-1}), R is the universal gas constant ($\text{Jmol}^{-1}\text{K}^{-1}$), T is the cell temperature (assumed constant) and η is the overpotential (V). The latter is defined by:

$$\eta = \phi_s - \phi_e - E_{eq,a} \quad (3)$$

in which ϕ_s and ϕ_e are the solid phase and electrolyte potentials, respectively, and E_{eq} is the equilibrium potential (assumed constant) for the anode or cathode ($j \in \{a, c\}$).

The cell is assumed to operate in potentiostatic mode. At anode channel rib, where the electrode interfaces with the end plate, a reference potential of zero is imposed, while the potential is set to the cell voltage E_{cell} at the cathode channel rib, where again the electrode interfaces with the end plate. At the walls of the gas channels and the porous electrode, insulation (zero mass and charge fluxes) is assumed. The compositions (weight fractions of species) are specified at the inlets, while the outlet conditions assume a convective flux. The default parameter values are specified in the documentation and are not repeated here to conserve space. Aside from the chosen inputs (specified below) all parameters were set to their default values.

Various settings can be adjusted in order to generate solutions of different fidelity. The low-fidelity solution in this case was defined by 360 domain elements, 474 boundary elements, 212 edge elements (corresponding to Fig. 1) and a relative tolerance of 0.1, while the high-fidelity solution was defined by 883601 domain elements, 93067 boundary elements, 3979 edge elements and a relative tolerance of 0.001, with other settings fixed, making the latter far more accurate and costly. The low fidelity model took 3 minutes and 3 seconds to compute the 128 solutions, while the high-fidelity model took 14 hours, 17 minutes and 8 seconds on a Macbook pro 2.3 GHz, 8-Core Intel Core i9 with 64 GB 2667 MHz DDR4 RAM.




Figure 1: Computational domain for the COMSOL model.

Four inputs were chosen for the experiments:

1. Electrode porosities in the range $\epsilon \in [0.4, 0.85]$
2. Cell voltage in the range $E_{\text{cell}} \in [0.3, 0.95]$ V
3. Cell temperature (assumed constant) in the range $T \in [973, 1273]$ K
4. Channel pressures (assumed equal) in the range $P \in [0.5, 2.5]$ atm.

and are collected inside an input vector:

$$\boldsymbol{\xi} = (\epsilon, E_{\text{cell}}, T, P)^T \quad (4)$$

The ranges above define the feasible input space $\mathcal{D} \subset \mathbb{R}^p$, $p = 4$.

$N = 128$ values of each input (vector) were selected using a Sobol sequence design-of-experiment, and are labelled $\boldsymbol{\xi}_n \in \mathcal{D}$, $n = 1 \dots, N$. The chosen outputs were: (1) the electrolyte current density $i_e(\mathbf{x}; \boldsymbol{\xi})$ (A m^{-2}); and (2) the ionic (electrolyte) potential $\phi_e(\mathbf{x}; \boldsymbol{\xi})$ (V), both in the $x - z$ plane at the centres of the

channels in the y direction (see Fig. 1). This notation indicates that these outputs are functions of \mathbf{x} , and are parameterised by the inputs $\boldsymbol{\xi}$.

Each of the outputs was recorded at locations \mathbf{x}_j , $j = 1, \dots, d = 5000$ on a uniform grid, and were subsequently vectorised to form vectors $\mathbf{y}_n^f \in \mathbb{R}^d$, where the subscript n indicates that the input $\boldsymbol{\xi}_n$ was used, and f denotes the fidelity, namely $f = F$ for the high fidelity output, and $f = L$ for the low fidelity output. Note that we use a generic symbol \mathbf{y}_n^f to denote both outputs in order to keep notation to a minimum.

Later in the optimization section 4.3 we also consider the volumetric current density in the anode $i_t(\mathbf{x}; \boldsymbol{\xi})$ (Am^{-3}) and the cathode overpotential $\eta_c(\mathbf{x}; \boldsymbol{\xi})$ (V), both recorded at 250,000 locations. We may think of the outputs as values of mappings (functions) $\mathbf{y}^f(\boldsymbol{\xi})$ of the input; that is, $\mathbf{y}_n^f = \mathbf{y}^f(\boldsymbol{\xi}_n)$. Our ultimate goal is to find an estimator for the high-fidelity mapping $\mathbf{y}^F(\boldsymbol{\xi})$ so that rapid and accurate predictions can be made for an arbitrary $\boldsymbol{\xi}$. This estimator can then be used for optimization, or any other application requiring repeated evaluations of, and predictions from the original computer model, e.g., sensitivity analysis and uncertainty quantification. The method we employ is an adaptation of the stochastic collocation (SC) approach [47], overcoming two of its limitations.

3 Stochastic collocation multi-fidelity fusion

The multi-fidelity SC methodology is based on a combination of sampling refinement and multi-fidelity data fusion. It has been successfully applied to numerous problems in science and engineering [50, 52, 54, 55]. Consider low- and high-fidelity models that define mappings $\mathbf{y}^f : \mathcal{D} \rightarrow \mathbb{R}^d$ from a p -dimensional input parameter space \mathcal{D} to low and high-fidelity output spaces lying in \mathbb{R}^d . In essence, the SC approach [47] finds an estimator $\mathbf{y}^F(\boldsymbol{\xi})$ of the high-fidelity mapping in the form:

$$\mathbf{y}^H(\boldsymbol{\xi}) = \sum_{n=1}^M c_n(\boldsymbol{\xi}) \mathbf{y}_n^H \quad (5)$$

in which M is the number of training points and $c_n(\boldsymbol{\xi})$ are unknown coefficients that depend on the input $\boldsymbol{\xi}$. Clearly, this approach amounts to an interpolation, in which the coefficients (usually assumed to be polynomials) are estimated using values of $\mathbf{y}^H(\boldsymbol{\xi})$ at judiciously selected collocation points, such as the Gauss-Lobatto-Legendre or Fekete points [56]. Due to the cost associated with calculating values of $\mathbf{y}^H(\boldsymbol{\xi})$, however, the low-fidelity solutions are instead employed to approximate the $c_n(\boldsymbol{\xi})$. Given $\mathbf{y}^L(\boldsymbol{\xi})$, corresponding to a general $\boldsymbol{\xi}$, the coefficients are determined by the conditions:

$$\sum_{n=1}^M c_n(\boldsymbol{\xi}) (\mathbf{y}_n^L)^T \mathbf{y}_n^L = (\mathbf{y}^L(\boldsymbol{\xi}))^T \mathbf{y}_n^L, \quad \forall \mathbf{y}_n^L \quad (6)$$

This is an interpolation or projection of $\mathbf{y}^L(\boldsymbol{\xi})$ onto the space $\text{span}\{\mathbf{y}_n^L\}$, i.e., all linear combinations of the low fidelity data points. In practice, the SC approach is equivalent to a Gaussian process model with a linear kernel (Bayesian linear regression), as demonstrated in [48].

Assume that we are able to sample at random from the parameter space \mathcal{D} and obtain N low-fidelity model evaluations. An $N \times N$ Gramian matrix \mathbf{G}^L can be constructed as follows:

$$[G^L]_{l,n} = k(\mathbf{y}^L(\boldsymbol{\xi}_l), \mathbf{y}^L(\boldsymbol{\xi}_n)), \quad l, n = 1, \dots, N \quad (7)$$

where $k(\cdot, \cdot) : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ is any valid kernel function. In the original formulation [47], a linear kernel is used, i.e.:

$$k(\mathbf{y}^L(\boldsymbol{\xi}_l), \mathbf{y}^L(\boldsymbol{\xi}_n)) = (\mathbf{y}^L(\boldsymbol{\xi}_l))^T \mathbf{y}^L(\boldsymbol{\xi}_n) \quad (8)$$

as seen in Eq. (6). A sampling refinement procedure to select the collocation points is then employed, involving the application of an LU decomposition with complete pivoting on the low-fidelity Gramian matrix \mathbf{G}^L . An ordering of the low-fidelity data points based on their ‘importance’ (explained in the next section) is obtained via the calculation of a matrix \mathbf{Q} from:

$$\mathbf{P}\mathbf{G}^L\mathbf{Q} = \mathbf{L}\mathbf{U} \quad (9)$$

where the matrices \mathbf{L} and \mathbf{U} are lower and upper triangular factorization matrices, respectively, while \mathbf{P} and \mathbf{Q} are permutation matrices that reorder the rows and columns of the Gramian matrix, respectively. This procedure leads to the selection of $m \ll N$ sampling points with indices n_1, \dots, n_m , for determining the $c_n(\boldsymbol{\xi})$ and evaluating the high-fidelity output at any $\boldsymbol{\xi} \in \mathcal{D}$. The latter is given by Eq. (5) applied at the sampling points:

$$\mathbf{y}^H(\boldsymbol{\xi}) = \sum_{i=1}^m c_i(\boldsymbol{\xi}) \mathbf{y}^H(\boldsymbol{\xi}_{n_i}) \quad (10)$$

in which the $c_i(\boldsymbol{\xi})$ are computed via a least-squares projection onto the low-fidelity model data:

$$\begin{pmatrix} G_{i_1, i_1}^L & \cdots & G_{i_1, i_m}^L \\ \vdots & \ddots & \vdots \\ G_{i_m, i_1}^L & \cdots & G_{i_m, i_m}^L \end{pmatrix} \begin{pmatrix} c_1(\boldsymbol{\xi}) \\ \vdots \\ c_m(\boldsymbol{\xi}) \end{pmatrix} = \begin{pmatrix} k(\mathbf{y}^L(\boldsymbol{\xi}), \mathbf{y}^L(\boldsymbol{\xi}_{i_1})) \\ \vdots \\ k(\mathbf{y}^L(\boldsymbol{\xi}), \mathbf{y}^L(\boldsymbol{\xi}_{i_m})) \end{pmatrix} \quad (11)$$

The $\mathbf{y}^H(\boldsymbol{\xi})$ in Eq. (10) is of course not the true latent function but an estimator. In order to avoid notational clutter, we do not make the distinction unless it is necessary. The same applies to the estimator of $\mathbf{y}^L(\boldsymbol{\xi})$ introduced later.

3.1 Explanation of the point selection strategy in multi-fidelity stochastic collocation

Selection of the inputs to generate the high-fidelity training data, i.e., the interpolation nodes $\boldsymbol{\xi}_{n_i}$ used in Eq. 10, is a weak greedy procedure. Consider a general description, in which we have a low-fidelity model $u^L(\boldsymbol{\xi}) : \mathcal{D} \rightarrow \mathcal{V}^L$ in some low-fidelity solution space \mathcal{V}^L that is assumed to be Hilbert space with inner

product $\langle \cdot, \cdot \rangle^L$. Then the distance between any subspace $\mathcal{W} \subset \mathcal{V}^L$ and a function $v(\boldsymbol{\xi}) \in \mathcal{V}^L$ can be defined as the greatest lower bound of the distance between v and \mathcal{W} , measured in terms of the norm induced by the inner product $\|\cdot\|^L = \langle \cdot, \cdot \rangle^L$:

$$d^L(v, \mathcal{W}) = \inf_{w \in \mathcal{W}} \|(I - P_{\mathcal{W}})v\|^L \quad (12)$$

where $P_{\mathcal{W}}$ denotes the orthogonal projection operator onto \mathcal{W} and I is the identity operator. In our case, $u(\boldsymbol{\xi}) = \mathbf{y}^L(\boldsymbol{\xi})$ and $\mathcal{V}^L \subset \mathbb{R}^d$, which is obviously a (finite dimensional) Hilbert space, for which we can use the standard inner product.

In radial basis methods (RBMs) for the (single-fidelity) solution $u(\boldsymbol{\xi}) : \mathcal{D} \rightarrow \mathcal{V}$ to a parameter-dependent PDE problem, a common greedy strategy is to select R query points $\boldsymbol{\xi}_1, \dots, \boldsymbol{\xi}_R$, to generate basis vectors $u(\boldsymbol{\xi}_1), \dots, u(\boldsymbol{\xi}_R)$, from the functional manifold $\{u(\boldsymbol{\xi}) : \boldsymbol{\xi} \in \mathcal{D}\}$ such that the distance (as defined above) between $\{u(\boldsymbol{\xi}) : \boldsymbol{\xi} \in \mathcal{D}\}$ and $\text{span}\{u(\boldsymbol{\xi}_1), \dots, u(\boldsymbol{\xi}_R)\} \subset \mathcal{V}$ is minimized.

We would like to perform a similar procedure to select query points for the high-fidelity data, but since the high-fidelity model is expensive, the low fidelity solution is instead used as a proxy. The selection of query points in RBMs as described above is implemented iteratively, by adding a point $\boldsymbol{\xi}_*$ to the optimal set Γ_n at step n such that the distance between $u^L(\boldsymbol{\xi}_*)$ and $\mathcal{U}^L(\Gamma_n) = \text{span}\{u(\boldsymbol{\xi}) \mid \boldsymbol{\xi} \in \Gamma_n\}$ is maximal:

$$\boldsymbol{\xi}_* = \underset{\boldsymbol{\xi} \in \overline{\mathcal{D}}}{\text{argmax}} d^L(u^L(\boldsymbol{\xi}), \mathcal{U}^L(\Gamma_n)), \quad \Gamma_{n+1} = \Gamma_n \cup \{\boldsymbol{\xi}_*\} \quad (13)$$

in which $\overline{\mathcal{D}} \subseteq \mathcal{D}$ and $\Gamma_0 = \{\}$, i.e., the empty set. Performing this process in the continuous space $\overline{\mathcal{D}} = \mathcal{D}$ is generally not computationally feasible, due to the potentially very high number of queries to the low fidelity model u^L allied with the lack of a closed form expression for $\nabla_{\boldsymbol{\xi}} u^L$, necessitating expensive numerical approximations. It is, however, achievable in a discrete finite-cardinality set $\overline{\mathcal{D}} \subset \mathcal{D}$.

As proven by Narayan et al. [47], full-pivoting LU decomposition of the low-fidelity Gramian matrix \mathbf{G}^L to obtain the m indices n_1, \dots, n_m , and corresponding sampling points $\Gamma_m = \{\boldsymbol{\xi}_{n_i}\}$ is one of the greedy solutions to the optimization problem defined in Eq. 13 on the finite-cardinality candidate set $\overline{\mathcal{D}} = \{\boldsymbol{\xi}_n\}_{n=1}^M$. This assumes that $\overline{\mathcal{D}}$ is dense enough to capture all important variations in the high-fidelity data, and that there are no restrictions on selecting point in \mathcal{D} .

3.2 Bayesian nonparametric model for the low-fidelity output when making predictions

In the original framework [47], a high-fidelity prediction $\mathbf{y}^H(\boldsymbol{\xi})$ for an unseen input $\boldsymbol{\xi}$ requires the true low-fidelity output $\mathbf{y}^L(\boldsymbol{\xi})$, which can still be very expensive to obtain in real-world applications. To build a more efficient multi-fidelity model, we use a GP approximation for the low-fidelity output by introducing a prior:

$$\mathbf{y}^L(\boldsymbol{\xi}) \sim \mathcal{GP}(\mathbf{y}^L(\boldsymbol{\xi}) \mid \mathbf{0}, k(\boldsymbol{\xi}, \boldsymbol{\xi}' \mid \boldsymbol{\theta}) \otimes \mathbf{Q} + \sigma^2 \mathbf{I} \otimes \delta(\boldsymbol{\xi}, \boldsymbol{\xi}')) \quad (14)$$

in which $\mathcal{GP}(\cdot | \cdot, \cdot)$ denotes a GP, with the first and second arguments denoting the mean and covariance functions. The symbol \sim signifies that the random variable or process to left is distributed according to the distribution to the right, \otimes is the Kronecker product, $\delta(\boldsymbol{\xi}, \boldsymbol{\xi}')$ is the Kronecker-delta function and \mathbf{Q} is a coregionalization matrix (treated as hyperparameter), containing the covariances between the components of $\mathbf{y}^L(\boldsymbol{\xi})$. $k(\boldsymbol{\xi}, \boldsymbol{\xi}' | \boldsymbol{\theta})$ is a kernel for the covariance across inputs, containing a vector of hyperparameter $\boldsymbol{\theta}$. The product structure $k(\boldsymbol{\xi}, \boldsymbol{\xi}' | \boldsymbol{\theta}) \otimes \mathbf{Q}$ is referred to as separable [57]. The term $\sigma^2 \mathbf{I} \otimes \delta(\boldsymbol{\xi}, \boldsymbol{\xi}')$ accounts for i.i.d. noise in the data, or equivalently acts as a regularization (penalty) term.

One of the most widely used kernels is the squared exponential automatic relevance determination (ARD) kernel:

$$k(\boldsymbol{\xi}, \boldsymbol{\xi}' | \boldsymbol{\theta}) = \exp \left(- \sum_{l=1}^d \frac{|\xi_{n,l} - \xi'_{n,l}|^2}{\theta_l^2} \right) \quad (15)$$

in which $\xi_{n,l}$ is the l -th component of $\boldsymbol{\xi}_n$, and θ_l are correlation lengths that form the vector $\boldsymbol{\theta}$. The total set of hyperparameters is $\{\boldsymbol{\theta}, \mathbf{Q}, \sigma^2\}$.

The simulation data can be collected in a matrix $\mathbf{Y}^L = [\mathbf{y}_1^L, \dots, \mathbf{y}_{N_L}^L]^T \in \mathbb{R}^{N_L \times d}$, where N_L is the number of low fidelity data points used for training. The log marginal likelihood \mathcal{L} for the model 14 is:

$$\mathcal{L} = \frac{1}{2} \ln |\boldsymbol{\Sigma}| - \frac{1}{2} \text{tr} (\text{vec}(\mathbf{Y}^L) \text{vec}(\mathbf{Y}^L)^T \boldsymbol{\Sigma}^{-1}) - \frac{dN_L}{2} \ln(2\pi) \quad (16)$$

in which $\text{vec}(\cdot)$ denotes a vectorization and tr is the trace of a matrix. $\boldsymbol{\Sigma} = \mathbf{Q} \otimes \mathbf{K} + \sigma^2 \mathbf{I} \in \mathbb{R}^{dN_L \times dN_L}$ is the covariance matrix, in which $\mathbf{K} \in \mathbb{R}^{N_L \times N_L}$ is given by $[K]_{ij} = k(\boldsymbol{\xi}_i, \boldsymbol{\xi}_j)$. A maximum likelihood estimate (MLE) can be used for the hyperparameters, which can be placed in the following posterior distribution for $\mathbf{y}^L(\boldsymbol{\xi})$ [58]:

$$\begin{aligned} \mathbb{E}[\mathbf{y}^L(\boldsymbol{\xi})] &= (\mathbf{Q} \otimes \mathbf{k}(\boldsymbol{\xi}))^T \boldsymbol{\Sigma}^{-1} \text{vec}(\mathbf{Y}^L) \\ \text{Var}(\mathbf{y}^L(\boldsymbol{\xi})) &= k(\boldsymbol{\xi}, \boldsymbol{\xi}) \mathbf{Q} - (\mathbf{Q} \otimes \mathbf{k}(\boldsymbol{\xi}))^T \boldsymbol{\Sigma}^{-1} (\mathbf{Q} \otimes \mathbf{k}(\boldsymbol{\xi})) \end{aligned} \quad (17)$$

in which $\mathbf{k}(\boldsymbol{\xi}) = [k(\boldsymbol{\xi}, \boldsymbol{\xi}_1), \dots, k(\boldsymbol{\xi}, \boldsymbol{\xi}_{N_H})]^T$ contains the correlations between the output at a test input $\boldsymbol{\xi}$ and the observations in the data set.

3.3 Efficient computation using a noise-free assumption

For multi-fidelity fusion problems, the noise term can, and in fact usually is, taken to be zero by virtue of the fact that the low- and high-fidelity models are deterministic [47]. With this assumption, the predictive mean in

Eq. 17 simplifies as follows:

$$\begin{aligned}
\mathbb{E}[\mathbf{y}^L(\boldsymbol{\xi})] &= (\mathbf{Q} \otimes \mathbf{k}(\boldsymbol{\xi}))^T (\mathbf{Q} \otimes \mathbf{K})^{-1} \text{vec}(\mathbf{Y}^L) \\
&= \mathbf{Q} \otimes (\mathbf{k}(\boldsymbol{\xi}))^T (\mathbf{Q}^{-1} \otimes \mathbf{K}^{-1}) \text{vec}(\mathbf{Y}^L) \\
&= (\mathbf{Q}\mathbf{Q}^{-1}) \otimes (\mathbf{k}(\boldsymbol{\xi}))^T \mathbf{K}^{-1} \text{vec}(\mathbf{Y}^L) \\
&= \mathbf{I} \otimes (\mathbf{k}(\boldsymbol{\xi}))^T \mathbf{K}^{-1} \text{vec}(\mathbf{Y}^L) \\
&= \text{vec} \left((\mathbf{k}(\boldsymbol{\xi}))^T \mathbf{K}^{-1} \mathbf{Y}^L \mathbf{I}^T \right) \\
&= (\mathbf{Y}^L)^T \mathbf{K}^{-1} \mathbf{k}(\boldsymbol{\xi})
\end{aligned} \tag{18}$$

demonstrating that the value of \mathbf{Q} does not affect the mean predictions. What is more, outputs in high-dimensional spaces (large d) can be handled without the need for the difficult task of learning the $d(d+1)/2$ components of \mathbf{Q} . To illustrate the difficulty in the optimization of (16), in the examples we consider, $d \in \{5000, 250000\}$, yielding 1.25025×10^7 or 3.125×10^{10} entries. This would severely compromise the accuracy of the maximum likelihood estimate of the hyperparameters.

Therefore, we set $\mathbf{Q} = \mathbf{I}$ to obtain:

$$\mathbf{y}^L(\boldsymbol{\xi}) \sim \mathcal{GP}(\mathbf{y}^L(\boldsymbol{\xi}) \mid \mathbf{0}, k(\boldsymbol{\xi}, \boldsymbol{\xi}') \otimes \mathbf{I}) \tag{19}$$

and the likelihood:

$$\mathcal{L} = \frac{d}{2} \ln |\mathbf{K}| - \frac{1}{2} \text{tr} \left((\mathbf{Y}^L)^T \mathbf{K}^{-1} \mathbf{Y}^L \right) - \frac{N_L d}{2} \ln(2\pi) \tag{20}$$

The computational complexity associated with this likelihood is now $\mathcal{O}(N_L^3 d)$, in contrast to the $\mathcal{O}(N_L^3 d^3)$ complexity of the original likelihood maximization, making it significantly more scalable to high values of d . The corresponding predictive posterior is as follows:

$$\begin{aligned}
\mathbb{E}[\mathbf{y}^L(\boldsymbol{\xi})] &= \mathbf{k}(\boldsymbol{\xi})^T \mathbf{K}^{-1} \mathbf{Y}^L \\
\text{Var}(\mathbf{y}^L(\boldsymbol{\xi})) &= (k(\boldsymbol{\xi}, \boldsymbol{\xi}) - \mathbf{k}(\boldsymbol{\xi})^T \mathbf{K}^{-1} \mathbf{k}(\boldsymbol{\xi})) \otimes \mathbf{I}
\end{aligned} \tag{21}$$

which constitutes the estimator for the true low fidelity mapping.

3.4 Feature enhancement for the low-fidelity Gramian: an efficient nonlinear kernel

We note that the predictive equations in Eqs. (10) and (11) are very similar to those in Eq. (21). As we pointed out in [48], SC is a special case of NAR when using a linear kernel. However, in practice, we find that introducing a nonlinear kernel, e.g., an ARD kernel, can only improve the performance when sufficient data is provided to accurately learn the hyperparameters introduced; for ARD there are $\mathcal{O}(d)$ hyperparameters, which is typically very high for the types of models under consideration. In a practical situation in which the high-fidelity observations are sparse, performance will therefore be compromised.

There are possible remedies via approximate or fully Bayesian approaches, e.g., variational Bayes or MCMC, which, however, are computationally expensive and can defeat the original purpose of finding an inexpensive surrogate for the high fidelity model. To overcome this problem, we can use a feature enhancement technique to augment the low-fidelity observations in order to make more accurate high-fidelity predictions.

Specifically, we design a feature transformation of the low-fidelity observations via a linear combination of kernels as follows:

$$[G^L]_{l,j} = \langle \mathbf{y}^L(\boldsymbol{\xi}_l), \mathbf{y}^L(\boldsymbol{\xi}_j) \rangle + \omega_1 \langle (\mathbf{y}^L(\boldsymbol{\xi}_l))^2, (\mathbf{y}^L(\boldsymbol{\xi}_j))^2 \rangle + \omega_2 \langle \exp(\mathbf{y}^L(\boldsymbol{\xi}_l)), \exp(\mathbf{y}^L(\boldsymbol{\xi}_j)) \rangle + \omega_3 \langle \ln(\mathbf{y}^L(\boldsymbol{\xi}_l)), \ln(\mathbf{y}^L(\boldsymbol{\xi}_j)) \rangle + \omega_4 \langle \sin(\mathbf{y}^L(\boldsymbol{\xi}_l)), \sin(\mathbf{y}^L(\boldsymbol{\xi}_j)) \rangle + \dots \quad (22)$$

where $\langle \cdot, \cdot \rangle$ denotes the standard Euclidean inner product, ω_k are weights and the operations $(\cdot)^2$, $\exp(\cdot)$, etc. are performed element-wise. Other operations included in this work are $\sqrt{\cdot}$, $|\cdot|$ and $\cos(\cdot)$, although any arbitrary number Z of functions can be incorporated. In order to avoid overfitting, we place a Laplace prior over each weight:

$$p(\omega_k) \propto \exp(-|\omega_k|) \quad (23)$$

and set $\boldsymbol{\omega} = (\omega_1, \dots, \omega_Z)^T$. The Laplace prior with a maximum-likelihood point estimate of $\boldsymbol{\omega}$ is then equivalent to minimizing the L_2 loss function between the data points and their SC approximations with an L_1 regularization term:

$$\boldsymbol{\omega}^* = \underset{\boldsymbol{\omega}}{\operatorname{argmin}} \frac{1}{2} \sum_{n=1}^{N_H} \|\mathbf{y}^H(\boldsymbol{\xi}_n) - \mathbf{y}_n^H\|^2 + \lambda \|\boldsymbol{\omega}\|_1 \quad (24)$$

in which: $\|\cdot\|_1$ is the L_1 norm; N_H is the number of high fidelity samples available ($N_H = m$ with the greedy strategy); $\mathbf{y}^H(\boldsymbol{\xi}_n)$ are the SC predictions (functions of $\boldsymbol{\omega}$) via Eqs. (10), (11) and (22) of the high-fidelity observations \mathbf{y}_n^H ; and λ is the regularization constant, which is determined using a leave-one-out cross-validation.

3.5 High-fidelity predictive posterior using Monte-Carlo integration

Given $\boldsymbol{\theta}$ and $\boldsymbol{\omega}$, we can derive the high-fidelity posterior distribution. Despite the fact that the low-fidelity prediction (21) for $\mathbf{y}^L(\boldsymbol{\xi})$ is Gaussian, due to the feature enhancement, the high-fidelity prediction $\mathbf{y}^H(\boldsymbol{\xi})$ is no longer tractable, which necessitates integration:

$$p(\mathbf{y}^H(\boldsymbol{\xi})) = \int p(\mathbf{y}^L(\boldsymbol{\xi})) p(\mathbf{y}^H(\boldsymbol{\xi}) | \mathbf{y}^L(\boldsymbol{\xi})) d\mathbf{y}^L(\boldsymbol{\xi}) \quad (25)$$

in which $p(\mathbf{y}^L(\boldsymbol{\xi}))$ is the predictive posterior in Eq. (21) and $p(\mathbf{y}^H(\boldsymbol{\xi}) | \mathbf{y}^L(\boldsymbol{\xi}))$ is the nonlinear deterministic mapping given by Eqs. (10), (11) and (22). If d is low, says less than 10, we can use an efficient Gaussian quadrature to approximate the integral. For observations in high-dimensional spaces, we can use a Monte-Carlo integration, which is more expensive but still efficient since sampling from a standard Gaussian distribution is straightforward and quick using reparameterization.

Table 1: Model complexity comparison. p is the dimension of the input space, d is the dimension of the output space, N_H and N_L are the number of high and low fidelity training points, respectively, and Z is the number of features used in Eq. (22).

| Method | Complexity | Number of parameters |
|------------|---|----------------------|
| Our method | $\mathcal{O}(N_L^3) + \mathcal{O}(N_H^3)$ | $(p + 1) + Z$ |
| NARGP | $\mathcal{O}(N_L^3) + \mathcal{O}(N_H^3)$ | $2(p + 1) + d$ |
| LAR | $\mathcal{O}((N_L + N_H)^3)$ | $2(p + 2) - 1$ |
| Greedy NAR | $\mathcal{O}(N_L^3) + \mathcal{O}(N_H^3)$ | $(p + 1) + d$ |
| ResGP | $\mathcal{O}(N_L^3) + \mathcal{O}(N_H^3)$ | $2(p + 1)$ |

To this end, we first sample Kd points γ_i from a normal distribution $\mathcal{N}(0, 1)$ and reshape them into a matrix $\mathbf{\Gamma} \in \mathbb{R}^{d \times K}$, with columns $\gamma_k \in \mathbb{R}^d$, $k = 1, \dots, K$. We can then approximate the mean of the high-fidelity predictive distribution as:

$$\mathbb{E}[\mathbf{y}^H(\boldsymbol{\xi})] \approx \sum_{k=1}^K (\mathbf{y}_k^H(\boldsymbol{\xi}) | \mathbf{y}_k^L(\boldsymbol{\xi})) = \frac{1}{K} \sum_{k=1}^K \sum_{l=1}^m c_l^k(\boldsymbol{\xi}) \mathbf{y}^H(\boldsymbol{\xi}_l) \quad (26)$$

in which

$$\mathbf{y}_k^L(\boldsymbol{\xi}) = \mathbf{k}(\boldsymbol{\xi})^T \mathbf{K}^{-1} \mathbf{Y}^L + \gamma_k \sqrt{k(\boldsymbol{\xi}, \boldsymbol{\xi}) - \mathbf{k}(\boldsymbol{\xi})^T \mathbf{K}^{-1} \mathbf{k}(\boldsymbol{\xi})} \quad (27)$$

are samples from the low-fidelity predictive posterior (21). The coefficients $c_l^k(\boldsymbol{\xi})$ are obtained by substituting the low-fidelity samples into Eq. (11) with the kernel defined in Eq. (22). Similarly, the variance of the predictive distribution is obtained from:

$$\text{Var}(\mathbf{y}^H(\boldsymbol{\xi})) = \frac{1}{K} \sum_{k=1}^K \left(\sum_{l=1}^m c_l^k(\boldsymbol{\xi}) \mathbf{y}^H(\boldsymbol{\xi}_l) - \mathbb{E}[\mathbf{y}^H(\boldsymbol{\xi})] \right) \left(\sum_{l=1}^m c_l^k(\boldsymbol{\xi}) \mathbf{y}^H(\boldsymbol{\xi}_l) - \mathbb{E}[\mathbf{y}^H(\boldsymbol{\xi})] \right)^T \quad (28)$$

3.6 Computational complexity

Table 1 compares the time/computational complexity for the training of our method with those of state-of-the-art methods. For a GP model with N_L training points, the computational complexity of the training is $\mathcal{O}(N_L^3)$, by virtue of the requirement to invert the $N_L \times N_L$ covariance matrix when implementing the maximum log-likelihood solution [58]. For SC, the complexity is $\mathcal{O}(N_H^3)$ due to the inversion of an $N_H \times N_H$ Gramian matrix in the loss function (24). This is the same complexity as that of all the other methods in Table 1, except for linear autoregression (LAR) [43], which is considerably higher. The number of model parameters is

comparable to ResGP [42] and LAR for typical values of p , but much lower than for NARGP [45] and Greedy NAR [48].

4 Results and discussion

In the comparisons in this section, NARGP, LAR, Greedy NAR and ResGP are implemented according to their original formulations and specifications. The GP based methods assume a zero-mean function after centering the data and ARD kernels are used in all of these methods, except for NARGP. For NARGP, the fidelity-1 kernel is an ARD kernel, and the fidelity 2 kernel is of the form used by Perdikaris et al. in their original implementation [45]. Three ARD kernels are combined together by factoring the dependence on the inputs and the low fidelity solution, and adding another kernel that represents a bias term. All methods were implemented in both Python (PyTorch) and Matlab. For the optimization, Matlab was preferred due to the availability of the Optimization Toolbox.

The prediction accuracy of all methods is measured using the root square mean error (RMSE)

$$\text{RMSE} = \sqrt{\frac{1}{N_t} \sum_{n=1}^{N_t} \|\mathbf{y}_{p,n} - \mathbf{y}_n\|^2} \quad (29)$$

in which N_t is the number of test points, $\mathbf{y}_{p,n}$ is the prediction and \mathbf{y}_n is the test point. All numerical experiments were conducted five times with randomization of the data, and the average value of the RMSE is presented. For individual results we use the square error (SE), $\|\mathbf{y}_{p,n} - \mathbf{y}_n\|^2$.

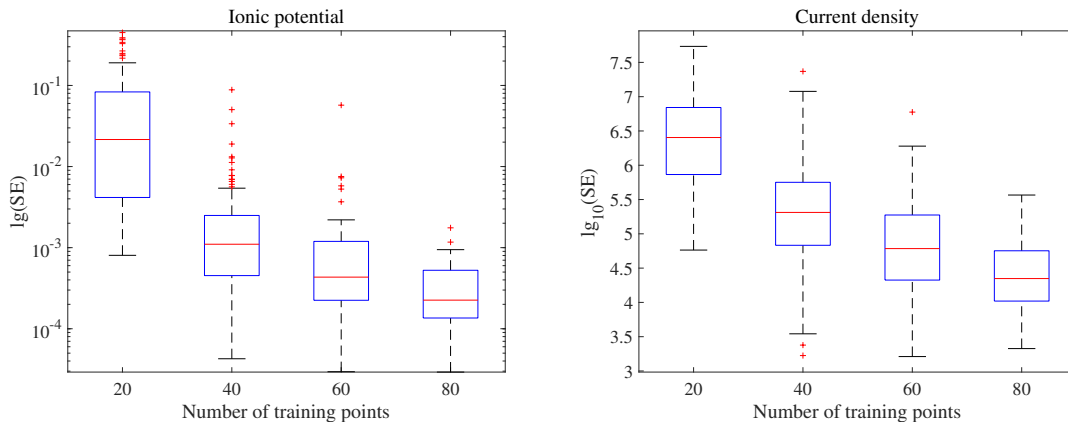


Figure 2: Box plots of the square errors (SE) against the test examples for the prediction of the low fidelity ionic potential $\phi_e(\mathbf{x}; \boldsymbol{\xi})$ (V) and current density $i_e(\mathbf{x}; \boldsymbol{\xi})$ (A m^{-2}) using Eq. (21), for an increasing number of training points.

Before moving onto the multi-fidelity results, we first show (Fig. 2) box plots of the square errors (SE)

against the test examples for the prediction of the low fidelity ionic potential $\phi_e(\mathbf{x}; \boldsymbol{\xi})$ (V) and current density $i_e(\mathbf{x}; \boldsymbol{\xi})$ (A m^{-2}) using Eq. (21), for an increasing number of training points. There is a clear decrease in the SE values as the number of training points increases, and the results are very accurate even at 40 training points.

To give an indication of the significance of the sizes of the errors, we first note that the mean value of the current density data was $1.058 \times 10^3 \text{ A m}^{-2}$. The square root of the median square error per coordinate of the output ($d = 5000$ in this case) at 40 training points is 5.963 A m^{-2} , while the square root of the maximum square error (an outlier) is 50.178 A m^{-2} . The corresponding errors are significantly lower at 80 training points. For the ionic potential, the mean value of the data was -0.0924 V . The square root of the median square error per coordinate at 40 training points is $4.472 \times 10^{-4} \text{ V}$, while the square root of the maximum square error (an outlier) is 1.890×10^{-3} . Again, the corresponding errors decrease significantly at 80 training points.

4.1 Prediction of the ionic potential

The RMSE values (with the five-fold cross validation) relating to the discretised ionic potential $\phi_e(\mathbf{x}; \boldsymbol{\xi})$ (V) in the $x - z$ plane at the channel centres are shown in Fig. 3, for increasing numbers of training samples, with the high- to low-fidelity sample ratios fixed to 1:1, 1:2, 1:3, and 1:4. In each case the remaining high fidelity data points were used for testing. Comparisons to the other methods are also shown. When the numbers of low and high fidelity samples are equal, all methods yield a similar level of accuracy, with the exception of NARGP at low training point numbers. This is not surprising considering the number of hyperparameters involved in NARGP.

The accuracy of our method is very high, with average RMSE values below $1.65 \times 10^{-5} \text{ V}$ per each of the 5000 coordinates of the output when the number of low fidelity samples is more than around 25 (for all training data ratios). This is ca. 0.017% of the mean value -0.0970 V of the ionic potential across the data set. For around 55 training points, the RMSE is below $1 \times 10^{-5} \text{ V}$ per coordinate. As the ratio of high-fidelity to low-fidelity samples decreases, our method shows a significant improvement over the other methods, even at a low to mid-range number of training points (40 to 60). NARGP performs poorly except when a large data set is available; around 85 for a ratio of 3:1 and around 100 for a ratio of 1:4, in which case it approaches the accuracy of our method. As already explained, this is to be expected for high values of d .

The improvement over the other methods for decreasing high-fidelity to low-fidelity sample ratios is significant because it allows for fewer high-fidelity training points. For example, at a ratio of 4:1, with 72 low-fidelity and only 18 high-fidelity training points, our method is approximately 26% more accurate than LAR, the next best method. LAR, moreover, has roughly double the computational complexity compared to our method for this number of low and high fidelity samples.

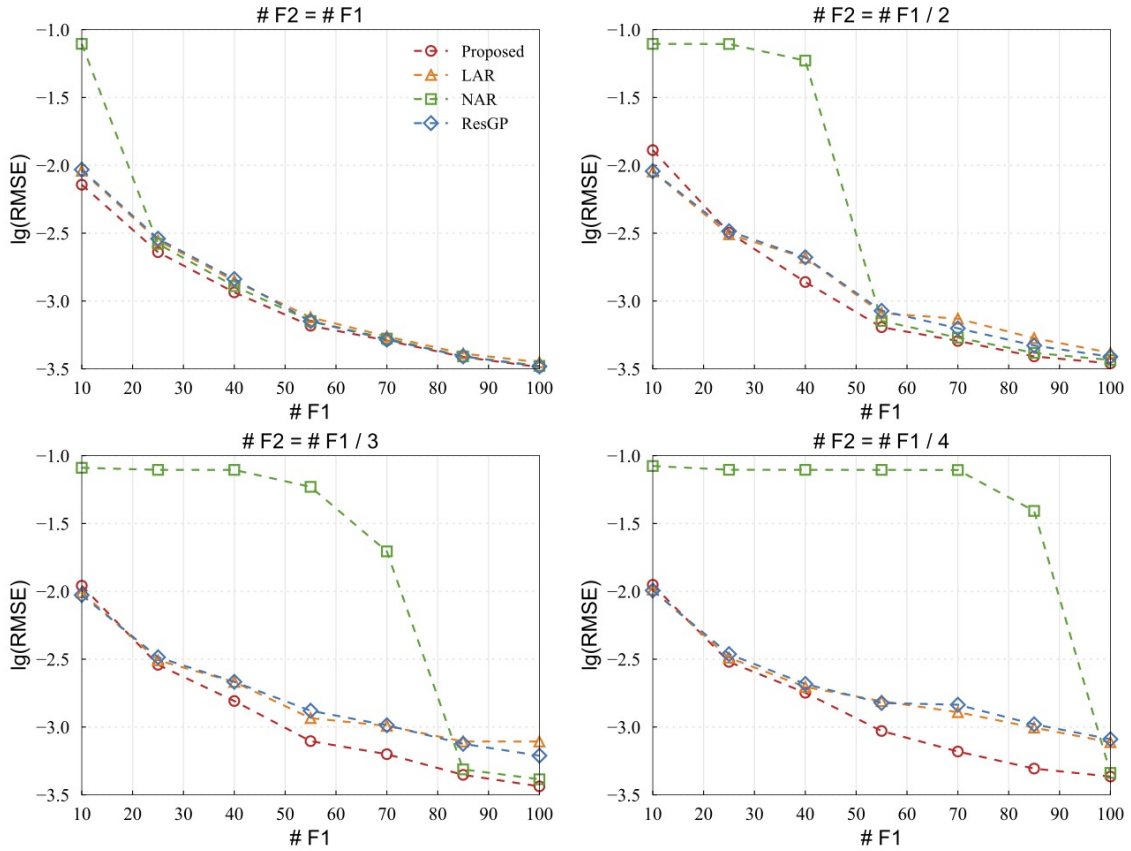


Figure 3: Log root mean square errors (RMSE) against the test examples for our proposed method, corresponding to different numbers and ratios of low- (F1) and high-fidelity (F2) ionic potential $\phi_e(\mathbf{x}; \xi)$ (V) training data points. Included are comparisons to LAR, NARGP, ResGP, and Greedy NAR.




Figure 4: A comparison of log RMSE for stochastic collocation (SC) and our proposed method with the feature enhancement in Eq. (22). In this case, the exact low fidelity data is provided for prediction.

We now compare our method to SC, in which both methods are provided with the exact low fidelity data for prediction. This allows us to assess the effect of the added features. The number of low and high fidelity training points is the same for both methods. As can be seen in Fig. 4, the methods are roughly the same for training point numbers below 25, but for higher numbers of training points our method is significantly more accurate. For training point numbers above 70, the RMSE is ca. 25% lower. Even at 55 training points it is more than 11% lower. This underlines the benefit of the added features.

We note also that in Fig. 4 the RMSE is significantly lower for both methods than the errors in Fig. 3. This is because the low fidelity data is provided for both methods, whereas in Fig. 3 the GP model (21) is used to predict the low-fidelity result in order to generate the high-fidelity approximation. Thus, it is not a fair comparison since execution of the low fidelity model involves considerable computational costs. For the optimization presented later (and other applications) it is usually infeasible to make repeated calls to the low fidelity model.




Figure 5: Three predictions of the ionic potential $\phi_e(\mathbf{x}; \boldsymbol{\xi})$ (V) in the $x - z$ plane at the channel centres (Fig. 1) for 80 low-fidelity and 20 high-fidelity training points. These predictions were chosen such that the top row exhibited the lowest square error, the middle row close to the median square error and the bottom row the highest square error. The first column contains the prediction, the second column the tests (ground truths) and the third column the point-wise absolute differences.

Examples of the predictions using our method are shown in Fig. 5, corresponding to the case with the lowest square error, a case in the region of the median square error and the case with the highest square error. Also shown are the tests (ground truths) and the point-wise absolute differences between the prediction and ground truth. 80 low-fidelity and 20 high-fidelity training points were used to generate these results. It can be seen that even in the case of the highest square error, the prediction is highly accurate, both qualitatively and quantitatively.

4.2 Prediction of the current density

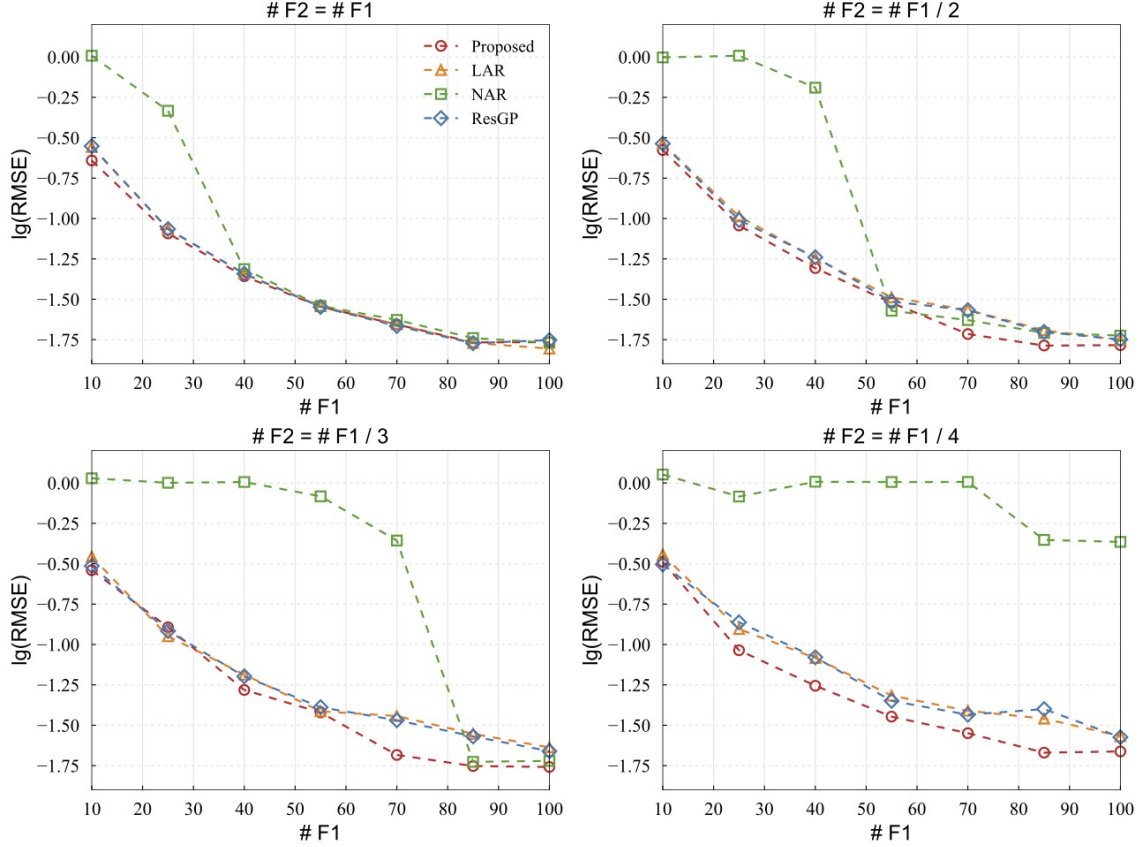


Figure 6: Log root mean square errors (RMSE) against the test examples for our proposed method, corresponding to different numbers and ratios of low- (F1) and high-fidelity (F2) current density $i_e(\mathbf{x}; \xi)$ (A m^{-2}) training data points. Included are comparisons to LAR, NARGP, ResGP, and Greedy NAR.

The RMSE values (with the five-fold cross validation) relating to the discretised current density $i_e(\mathbf{x}; \xi)$ (A m^{-2}) in the $x - z$ plane at the channel centres are shown in Fig. 6, for increasing numbers of training samples and different high- to low-fidelity sample ratios. In each case the remaining high fidelity data points were used for testing. Again, comparisons to the other methods are provided. As with the ionic potential, the RMSE values are roughly the same when the numbers of low and high-fidelity samples are equal. This, however, is unrealistic. The more realistic scenarios of 1:2, 1:3, and 1:4 show again that our method yields the lowest RMSE in most cases, and the differences are especially noticeable at the two lowest ratios. It must be said, however, that the differences are not as pronounced as in the case of the ionic potential. At a training ratio of 1:3, the RMSE is, nevertheless, ca. 20% lower than it is for ResGP, the next best method.

Examples of the predictions using our method are shown in Fig. 7, corresponding to the case with the lowest square error, a case with a square error in the vicinity of the median square error and the case with the highest square error, together with the ground truths and the point-wise absolute differences. 80 low-fidelity and 20 high-fidelity training points were used. Again, even in the case of the highest square error, the prediction is highly accurate.




Figure 7: Three predictions of the current density $i_e(\mathbf{x}; \boldsymbol{\xi})$ (A m^{-2}) in the $x - z$ plane at the channel centres (Fig. 1) for 80 low-fidelity and 20 high-fidelity training points. These predictions were chosen such that the top row exhibited the lowest square error, the middle row close to the median square error and the bottom row the highest square error. The first column contains the prediction, the second column the tests (ground truths) and the third column the point-wise absolute differences.

4.3 Surrogate assisted optimization of an SOFC

The results of the previous section demonstrate that our method has sufficient accuracy to be used as surrogate model for optimization. In this section we employ the method to optimize the SOFC performance with respect to several objectives. We note that the objectives are illustrative and serve as examples and tests, rather than being physically inspired.

We start with the aim of maximizing the current density in the anode. For this we extracted values of the volumetric current density $i_t(\mathbf{x}; \xi_n)$ for each of the 128 inputs ξ_n . The values were taken at 250,000 locations \mathbf{x}_j and vectorised to form outputs \mathbf{y}_n as described in section 2. The objective is:

$$\xi_* = \operatorname{argmax}_{\xi} \frac{1}{V} \int_{\mathbf{x}} i_t(\mathbf{x}; \xi) d\mathbf{x}$$

(30)

subject to

$$0.2 \leq \xi_1 \leq 0.9, \quad 0.1 \text{ V} \leq \xi_2 \leq 0.95 \text{ V} \quad 800 \text{ K} \leq \xi_3 \leq 1500 \text{ K}, \quad 0.5 \text{ atm} \leq \xi_3 \leq 3 \text{ atm}$$

where V is the volume of the anode. We note that the lower and upper bounds on the input are deliberately chosen to fall outside the design space \mathcal{D} used to train the model, so that problem (30) forms a good test for the reliability of the model. We use our multi-fidelity surrogate model for $i_t(\mathbf{x}; \xi_n)$, namely, for $\mathbf{y}(\xi)$, with a trapezoidal rule to estimate the integral in (30), based on the components of $\mathbf{y}(\xi)$.

The optimization method used is a primal-dual interior point method, i.e., both the primal and the dual problems are simultaneously solved, implemented in Matlab (the 'fmincon' function). This particular interior point method [59] is a variant of predictor-corrector infeasible-interior-point algorithm by Mehrotra [60], which calculates a Newton (prediction) direction, followed by a corrector direction, in which a barrier is introduced.

The work flow is as follows:

1. Load the data set (inputs and outputs for low and high fidelity)
2. Use Eq. (20), (21) to obtain the low-fidelity posterior prediction for $\mathbf{y}^L(\xi)$ for a given number of low fidelity training points
3. The high fidelity predictive posterior mean for $\mathbf{y}^H(\xi)$ is given by Eqs. (26) and (27), using Eqs. (11), (21) and (22), for a given number of low and high fidelity training points
4. Perform the optimization using the predictive posterior mean for $\mathbf{y}^H(\xi)$ to estimate the objective function using a trapezoidal rule

Using 80 low-fidelity and 40 high-fidelity training points, the results are summarised in Table 2, averaged over 5 experiments. The average time taken to estimate ξ_* (the entire procedure described above) was 22.95 s,

Table 2: Solutions to the optimization problems (30)-(33), together with the function evaluation counts and the times taken.

| Objective | ξ_1 | ξ_2/V | ξ_3/K | ξ_4/atm | Time/s | Function evals |
|-----------|---------|-----------|-----------|--------------------|--------|----------------|
| (30) | 0.9 | 0.1 | 800 | 3 | 22.95 | 121 |
| (31) | 0.9 | 0.1 | 800 | 3 | 12.11 | 67 |
| (32) | 0.2 | 0.1 | 800 | 3 | 17.27 | 316 |
| (33) | 0.2009 | 0.6957 | 1100.4 | 2.9992 | 17.27 | 387 |

and yielded $\xi_* = (0.9, 0.1 \text{ V}, 800 \text{ K}, 3 \text{ atm})^T - \epsilon$, in which each component of ϵ is $\mathcal{O}(10^{-7})$. The function evaluation count ranged from 87 to 153, underlining the benefit of replacing the high-fidelity solver with the surrogate. The RMSE for the high-fidelity prediction against 88 high-fidelity test points was $1.334 \times 10^{-4} \text{ A cm}^{-3}$ per component of \mathbf{y} ($d = 250, 000$), while the mean value of $i_t(\mathbf{x}; \xi_n)$ in the data set was 19.754 A cm^{-3} .

The result in Table 2 is of course correct, since for the highest current density we need to use the highest porosity and channel inlet pressures, together with the lowest cell voltage and temperature. Note that in the COMSOL model, the inlet mole fractions of species are fixed and the channel inlet concentrations C follow the ideal gas law $C = P/(RT)$. Thus, higher temperatures lead to lower concentrations for a fixed channel pressure, which overwhelms any other effects associated with higher temperatures, such as faster kinetics. The highest channel concentrations are attained when $P = 3 \text{ atm}$ and $T = 800\text{K}$. Fixing the channel inlet pressures and varying the concentrations independently would lead to the opposite result, namely that higher temperatures would lead to higher current densities.

Applying the same procedure to:

$$\xi_* = \operatorname{argmax}_{\xi} \frac{1}{A} \int_{\mathbf{x}} i_e(\mathbf{x}; \xi) d\mathbf{x} \quad (31)$$

subject to

$$0.2 \leq \xi_1 \leq 0.9, \quad 0.1 \text{ V} \leq \xi_2 \leq 0.95 \text{ V}, \quad 800 \text{ K} \leq \xi_3 \leq 1500 \text{ K}, \quad 0.5 \text{ atm} \leq \xi_4 \leq 3 \text{ atm}$$

in which we maximize the average of the electrolyte current density in the $x - z$ plane (with area A) at the channel centres, led to the same result (to within 10^{-6} for all components of ξ) for the same reasons, as summarised in Table 2.

We now turn our attention to maximizing the magnitude of the average electrolyte potential in the $x - z$

plane at the channel centres:

$$\xi_* = \operatorname{argmax}_{\xi} \left| \frac{1}{A} \int_{\mathbf{x}} \phi_e(\mathbf{x}; \xi) d\mathbf{x} \right|$$

subject to

(32)

$$0.2 \leq \xi_1 \leq 0.9, \quad 0.1 \text{ V} \leq \xi_2 \leq 0.95 \text{ V}, \quad 800 \text{ K} \leq \xi_3 \leq 1500 \text{ K}, \quad 0.5 \text{ atm} \leq \xi_3 \leq 3 \text{ atm}$$

The result (to within 10^{-4} for all components of ξ) is shown in Table 2. Here we see that the maximum occurs for the lowest temperature and highest pressure (i.e., the highest channel concentrations), the lowest cell voltage and the lowest porosity. This is as expected, since the lowest electrode porosity leads to the highest volume fraction of electrolyte. In this case the average number of function evaluations was 316, which would take around 35 hrs to complete with the original high-fidelity COMSOL model. We now add a nonlinear constraint to (32), namely that the minimum value of $\phi_e(\mathbf{x}; \xi)$ should not be less than -0.2V :

$$\xi_* = \operatorname{argmax}_{\xi} \left| \frac{1}{A} \int_{\mathbf{x}} \phi_e(\mathbf{x}; \xi) d\mathbf{x} \right|$$

subject to

(33)

$$-0.2 \text{ V} \leq \min_{\xi} \phi_e(\mathbf{x}; \xi)$$

$$0.2 \leq \xi_1 \leq 0.9, \quad 0.1 \text{ V} \leq \xi_2 \leq 0.95 \text{ V}, \quad 800 \text{ K} \leq \xi_3 \leq 1500 \text{ K}, \quad 0.5 \text{ atm} \leq \xi_3 \leq 3 \text{ atm}$$

The result is again shown in Table 2. The values were averaged over 5 runs, with differences between the values of each coordinate of ξ within 0.5%. Here, we find that to maximize the magnitude of the average electrolyte potential, the added constraint raises the optimal temperature to ca. 1100 K and raises the cell voltage to ca. 0.7 V, while the channel pressure and electrode porosity remain unchanged. Both a higher cell voltage and a higher temperature (which lowers the channel concentrations) will lower the rate of charge transfer, and therefore lower the magnitude of the average electrolyte potential.

The final experiment relates to multiple objectives, namely:

$$\xi_* = \operatorname{argmax}_{\xi} \mathbf{f}(\xi) = \left(\frac{1}{V} \int_{\mathbf{x}} i_t(\mathbf{x}; \xi) d\mathbf{x}, \quad -\frac{1}{V} \int_{\mathbf{x}} \eta_c(\mathbf{x}; \xi) d\mathbf{x} \right)^T$$

subject to

(34)

$$0.4 \leq \xi_1 \leq 0.8, \quad 0.2 \text{ V} \leq \xi_2 \leq 0.85 \text{ V}, \quad 973 \text{ K} \leq \xi_3 \leq 1273 \text{ K}, \quad 0.5 \text{ atm} \leq \xi_3 \leq 2.5 \text{ atm}$$

in which we simultaneously wish to maximize the average volumetric anode current density and minimize the magnitude of the average cathode overpotential $\eta_c(\mathbf{x}; \xi)$. As with the volumetric anode current density, the cathode overpotential was recorded at 250,000 spatial locations and vectorised to form outputs \mathbf{y}_n . The objectives in this case are clearly competing. To solve this problem we approximate the Pareto front. This is the first stage of a so-called Generate-First Choose-Later approach, in which a set of Pareto optimal solutions

is first generated as a candidate set, and from which the design choice can subsequently be made according to other considerations, experience and preferences. A solution is called Pareto optimal or non-dominated if none of the objectives can be improved upon with degrading one or more of the other objectives. The Pareto front is the set of all such solutions, and it was approximated using the genetic algorithm implemented in the ‘gamultiobj’ function in Matlab.

The average time taken to build the surrogate model and conduct the optimization (over 5 experiments) was 242.43 s, and on average 7350 function evaluations were required. The results are shown in Figure 8, in which the objective function values corresponding to the Pareto optimal solutions are plotted against each other. In total, a set of 20 Pareto optimal solutions were generated, all of which could be considered as candidates for achieving the objective.

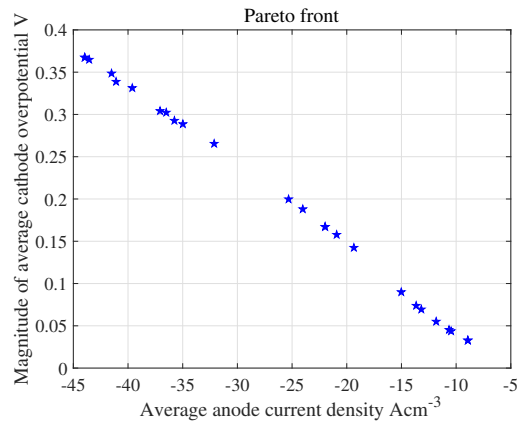


Figure 8: A plot of the objective function values corresponding to the Pareto optimal solutions to problem (34).

Although, as stated, the objectives chosen are for illustration rather than for any particular physical significance, most of the results above are readily explainable, which facilitates the validation of the approach. Other objectives (single or multiple) would not require a qualitatively different approach. The method is flexible and general enough to incorporate any summary objective or objectives, as long as the corresponding data is available from the multi-fidelity model.

5 Summary and Conclusions

Modelling and simulation play a major role in the development of fuel cell technologies. The complexity of computer models and the attendant times associated with their execution has led to a growth in surrogate modelling approaches for applications such as optimization. Very few such approaches are based on multi-fidelity methods, despite their advantages over pure machine learning models. In addition, few attempts have been made in the fuel cell literature to capture spatial variations in key quantities with surrogate models. In

this paper we extended the multi-fidelity stochastic collocation framework for spatially distributed outputs using a feature enhancement, and removed the requirement for exact low-fidelity outputs during inference by using another surrogate model. Our approach is Bayesian, so that estimates of uncertainty are possible.

We showed that the addition of features into the original stochastic collocation formulation can improve its accuracy. Experiments on a detailed SOFC model at two fidelities showed that our approach leads to a highly accurate surrogate model, outperforming several state-of-the-art multi-fidelity methods. It was used in several SOFC optimization problems, for a range of different objectives, for objectives with nonlinear constraints and for multiple objectives. The level of detail furnished by our approach allows for a great deal of flexibility in deciding the objectives a-posteriori, in contrast to focusing on specific scalar targets. It is also worth noting that using the original high-fidelity model would not be feasible for many problems of this nature, such as the multi-objective example we consider, in which over 7000 function evaluations were required.

Acknowledgment

This work was supported by the Creative Research Groups of the National Natural Science Foundation of China (No. 52021004).

References

- [1] Shambhu Singh Rathore, Saheli Biswas, Daniel Fini, Aniruddha P. Kulkarni, and Sarbjit Giddey. Direct ammonia solid-oxide fuel cells: A review of progress and prospects. *International Journal of Hydrogen Energy*, 46(71):35365–35384, 2021.
- [2] J.O. Abe, A.P.I. Popoola, E. Ajenifuja, and O.M. Popoola. Hydrogen energy, economy and storage: Review and recommendation. *International Journal of Hydrogen Energy*, 44(29):15072–15086, 2019.
- [3] Babak Ghorbani and Krishna Vijayaraghavan. A review study on software-based modeling of hydrogen-fueled solid oxide fuel cells. *International Journal of Hydrogen Energy*, 44(26):13700–13727, 2019.
- [4] Raj Kamal Abdul Rasheed, Quan Liao, Zhang Caizhi, and Siew Hwa Chan. A review on modelling of high temperature proton exchange membrane fuel cells (ht-pemfcs). *International Journal of Hydrogen Energy*, 42(5):3142–3165, 2017.
- [5] Kun Wang, Daniel Hissel, Marie-Cecile Péra, Nadia Steiner, Dario Marra, Marco Sorrentino, Cesare Pianese, Michela Monteverde, Pietro Cardone, and Jaakko Saarinen. A review on solid oxide fuel cell models. *International Journal of Hydrogen Energy*, 36(12):7212–7228, 2011.
- [6] A.A. Shah, K.H. Luo, T.R. Ralph, and F.C. Walsh. Recent trends and developments in polymer electrolyte membrane fuel cell modelling. *Electrochimica Acta*, 56(11):3731 – 3757, 2011.
- [7] Stefanos Tzelepis, Kosmas A. Kavadias, George E. Marnellos, and George Xydis. A review study on proton exchange membrane fuel cell electrochemical performance focusing on anode and cathode catalyst layer modelling at macroscopic level. *Renewable and Sustainable Energy Reviews*, 151:111543, 2021.
- [8] Wei-Zhuo Li, Wei-Wei Yang, Ning Wang, Yu-Hang Jiao, Yu Yang, and Zhi-Guo Qu. Optimization of blocked channel design for a proton exchange membrane fuel cell by coupled genetic algorithm and three-dimensional cfd modeling. *International Journal of Hydrogen Energy*, 45(35):17759–17770, 2020.
- [9] J. Antonio Salva, Alfredo Iranzo, Felipe Rosa, Elvira Tapia, Eduardo Lopez, and Fernando Isorna. Optimization of a pem fuel cell operating conditions: Obtaining the maximum performance polarization curve. *International Journal of Hydrogen Energy*, 41(43):19713–19723, 2016.
- [10] Hany M. Hasaniien, Mohamed A.M. Shaheen, Rania A. Turkey, Mohammed H. Qais, Saad Alghuwainem, Salah Kamel, Marcos Tostado-Véliz, and Francisco Jurado. Precise modeling of pem fuel cell using a novel enhanced transient search optimization algorithm. *Energy*, 247:123530, 2022.

- [11] Seyed Mohsen Pourkiaei, Fathollah Pourfayaz, Mehdi Mehrpooya, and Mohammad H Ahmadi. Multi-objective optimization of tubular solid oxide fuel cells fed by natural gas: an energetic and exergetic simultaneous optimization. *Journal of Thermal Analysis and Calorimetry*, 145(3):1575–1583, 2021.
- [12] Parisa Mojaver, Shahram Khalilarya, Ata Chitsaz, and Mohsen Assadi. Multi-objective optimization of a power generation system based sofc using taguchi/ahp/topsis triple method. *Sustainable Energy Technologies and Assessments*, 38:100674, 2020.
- [13] Bo Yang, Jingbo Wang, Mengting Zhang, Hongchun Shu, Tao Yu, Xiaoshun Zhang, Wei Yao, and Liming Sun. A state-of-the-art survey of solid oxide fuel cell parameter identification: Modelling, methodology, and perspectives. *Energy Conversion and Management*, 213:112856, 2020.
- [14] Meiting Guo, Dongqi Zhao, Qidong Xu, Zheng Li, Haoran Xu, and Meng Ni. New interconnector designs for electrical performance enhancement of solid oxide fuel cells: A 3d modelling study. *Journal of Power Sources*, 533:231373, 2022.
- [15] Meiting Guo, Dongqi Zhao, Qidong Xu, Zheng Li, Haoran Xu, and Meng Ni. New interconnector design optimization to balance electrical and mechanical performance of solid oxide fuel cell stack. *International Journal of Hydrogen Energy*, 48(8):3107–3121, 2023.
- [16] Taufiq Abdullah and Lin Liu. Simulation-based microstructural optimization of solid oxide fuel cell for low temperature operation. *International Journal of Hydrogen Energy*, 41(31):13632–13643, 2016.
- [17] Akos Kriston, Andreas Podias, Ibtissam Adanouj, and Andreas Pfrang. Analysis of the effect of thermal runaway initiation conditions on the severity of thermal runaway—numerical simulation and machine learning study. *Journal of The Electrochemical Society*, 167(9):090555, 2020.
- [18] Majdi I Radaideh, Mohammed I Radaideh, and Tomasz Kozlowski. Design optimization under uncertainty of hybrid fuel cell energy systems for power generation and cooling purposes. *International Journal of Hydrogen Energy*, 45(3):2224–2243, 2020.
- [19] Hossein Pourrahmani, Majid Siavashi, and Mahdi Moghimi. Design optimization and thermal management of the pemfc using artificial neural networks. *Energy*, 182:443–459, 2019.
- [20] L. Vichard, F. Harel, A. Ravey, P. Venet, and D. Hissel. Degradation prediction of pem fuel cell based on artificial intelligence. *International Journal of Hydrogen Energy*, 45(29):14953–14963, 2020.
- [21] Ke Sun, Iñaki Esnaola, Okechukwu Okorie, Fiona Charnley, Mariale Moreno, and Ashutosh Tiwari. Data-driven modeling and monitoring of fuel cell performance. *International Journal of Hydrogen Energy*, 46(66):33206–33217, 2021.

- [22] Jiangong Zhu, Yixiu Wang, Yuan Huang, R Bhushan Gopaluni, Yankai Cao, Michael Heere, Martin J Mühlbauer, Liuda Mereacre, Haifeng Dai, Xinhua Liu, et al. Data-driven capacity estimation of commercial lithium-ion batteries from voltage relaxation. *Nature communications*, 13(1):1–10, 2022.
- [23] Wenhua Li, Zhipeng Jiao, Le Du, Wenyi Fan, and Yazun Zhu. An indirect rul prognosis for lithium-ion battery under vibration stress using elman neural network. *International Journal of Hydrogen Energy*, 44(23):12270–12276, 2019.
- [24] Antonio Guarino, Riccardo Trincherio, Flavio Canavero, and Giovanni Spagnuolo. A fast fuel cell parametric identification approach based on machine learning inverse models. *Energy*, 239:122140, 2022.
- [25] Daniel Ritzberger, Johannes Höflinger, Zhang Peng Du, Christoph Hametner, and Stefan Jakubek. Data-driven parameterization of polymer electrolyte membrane fuel cell models via simultaneous local linear structured state space identification. *International Journal of Hydrogen Energy*, 46(21):11878–11893, 2021.
- [26] M. Peksen, L. Blum, and D. Stolten. Optimisation of a solid oxide fuel cell reformer using surrogate modelling, design of experiments and computational fluid dynamics. *International Journal of Hydrogen Energy*, 37(17):12540–12547, 2012.
- [27] Tabbi Wilberforce and AG Olabi. Proton exchange membrane fuel cell performance prediction using artificial neural network. *International Journal of Hydrogen Energy*, 46(8):6037–6050, 2021.
- [28] Jarosław Milewski and Konrad Świrski. Modelling the sofc behaviours by artificial neural network. *International Journal of Hydrogen Energy*, 34(13):5546–5553, 2009.
- [29] Vanja Subotić, Michael Eibl, and Christoph Hochenauer. Artificial intelligence for time-efficient prediction and optimization of solid oxide fuel cell performances. *Energy Conversion and Management*, 230:113764, 2021.
- [30] Xuexia Zhang, Jingzhe Zhou, and Weirong Chen. Data-driven fault diagnosis for pemfc systems of hybrid tram based on deep learning. *International Journal of Hydrogen Energy*, 45(24):13483–13495, 2020.
- [31] Jiahuan Lu, Rui Xiong, Jinpeng Tian, Chenxu Wang, Chia-Wei Hsu, Nien-Ti Tsou, Fengchun Sun, and Ju Li. Battery degradation prediction against uncertain future conditions with recurrent neural network enabled deep learning. *Energy Storage Materials*, 50:139–151, 2022.

- [32] Jaeyoo Choi, Yohan Cha, Jihoon Kong, Neil Vaz, Jaeseung Lee, Sang-Bum Ma, Jin-Hyuk Kim, Seung Woo Lee, Seung Soon Jang, and Hyunchul Ju. Multi-variate optimization of polymer electrolyte membrane fuel cells in consideration of effects of gdl compression and intrusion. *Journal of The Electrochemical Society*, 169(1):014511, 2022.
- [33] Haoran Xu, Jingbo Ma, Peng Tan, Bin Chen, Zhen Wu, Yanxiang Zhang, Huizhi Wang, Jin Xuan, and Meng Ni. Towards online optimisation of solid oxide fuel cell performance: Combining deep learning with multi-physics simulation. *Energy and AI*, 1:100003, 2020.
- [34] Guoping Xu, Zeting Yu, Lei Xia, Changjiang Wang, and Shaobo Ji. Performance improvement of solid oxide fuel cells by combining three-dimensional cfd modeling, artificial neural network and genetic algorithm. *Energy Conversion and Management*, 268:116026, 2022.
- [35] Yu Wang, Wenchun Jiang, Yun Luo, Ming Song, and Shan-Tung Tu. High temperature creep strength design and optimization of solid oxide fuel cell. *International Journal of Hydrogen Energy*, 47(50):21450–21461, 2022.
- [36] AA Shah. Surrogate modeling for spatially distributed fuel cell models with applications to uncertainty quantification. *Journal of Electrochemical Energy Conversion and Storage*, 14(1), 2017.
- [37] WW Xing, V Triantafyllidis, AA Shah, PB Nair, and Nicholas Zabarar. Manifold learning for the emulation of spatial fields from computational models. *Journal of Computational Physics*, 326:666–690, 2016.
- [38] W.W. Xing, F. Yu, P.K. Leung, X. Li, P. Wang, and A.A. Shah. A new multi-task learning framework for fuel cell model outputs in high-dimensional spaces. *Journal of Power Sources*, 482:228930, 2021.
- [39] M. Giselle Fernandez-Godino, Chanyoung Park, Nam-Ho Kim, and Raphael T. Haftka. Review of multi-fidelity models, 2016.
- [40] Yingxu Liu, Steffen Dirkes, Markus Kohn, Maximilian Wick, and Stefan Pischinger. A high-fidelity real-time capable dynamic discretized model of proton exchange membrane fuel cells for the development of control strategies. *Journal of Power Sources*, 537:231394, 2022.
- [41] Tyler G. Voskuilen, Harry K. Moffat, Benjamin B. Schroeder, and Scott A. Roberts. Multi-fidelity electrochemical modeling of thermally activated battery cells. *Journal of Power Sources*, 488:229469, 2021.
- [42] W.W. Xing, A.A. Shah, P. Wang, S. Zhe, Q. Fu, and R.M. Kirby. Residual gaussian process: A tractable nonparametric bayesian emulator for multi-fidelity simulations. *Applied Mathematical Modelling*, 97:36–56, 2021.

- [43] Marc C Kennedy and Anthony O’Hagan. Predicting the output from a complex computer code when fast approximations are available. *Biometrika*, 87(1):1–13, 2000.
- [44] Loic Le Gratiet. *Multi-fidelity Gaussian process regression for computer experiments*. PhD thesis, Université Paris-Diderot-Paris VII, 2013.
- [45] Paris Perdikaris, Maziar Raissi, Andreas Damianou, ND Lawrence, and George Em Karniadakis. Nonlinear information fusion algorithms for data-efficient multi-fidelity modelling. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 473(2198):20160751, 2017.
- [46] L. Parussini, D. Venturi, P. Perdikaris, and G.E. Karniadakis. Multi-fidelity Gaussian process regression for prediction of random fields. *Journal of Computational Physics*, 336(C):36–50, May 2017.
- [47] Akil Narayan, Claude Gittelsohn, and Dongbin Xiu. A stochastic collocation algorithm with multifidelity models. *SIAM Journal on Scientific Computing*, 36(2):A495–A521, 2014.
- [48] W. Xing, M. Razi, R.M. Kirby, K. Sun, and A.A. Shah. Greedy nonlinear autoregression for multifidelity computer models at different scales. *Energy and AI*, 1:100012, 2020.
- [49] Jerrad Hampton, Hillary Fairbanks, Akil Narayan, and Alireza Doostan. Parametric/stochastic model reduction: low-rank representation, non-intrusive bi-fidelity approximation, and convergence analysis. *arXiv preprint arXiv:1709.03661*, 2017.
- [50] Xueyu Zhu, Akil Narayan, and Dongbin Xiu. Computational aspects of stochastic collocation with multifidelity models. *SIAM/ASA Journal on Uncertainty Quantification*, 2(1):444–463, 2014.
- [51] M. Razi, A. Narayan, R. M. Kirby, and D. Bedrov. Fast predictive models based on multi-fidelity sampling of properties in molecular dynamics simulations. *Computational Material Science*, 152(C):125–133, 2018.
- [52] Ryan Skinner, Alireza Doostan, Eric Peters, John Evans, and Kenneth E Jansen. An evaluation of bi-fidelity modeling efficiency on a general family of NACA airfoils. In *35th AIAA Applied Aerodynamics Conference*, page 3260, 2017.
- [53] Mani Razi, Robert M Kirby, and Akil Narayan. Fast predictive multi-fidelity prediction with models of quantized fidelity levels. *Journal of Computational Physics*, 376:992–1008, 2019.
- [54] Lluís Jofre, Gianluca Geraci, Hillary Fairbanks, Alireza Doostan, and Gianluca Iaccarino. Multi-fidelity uncertainty quantification of irradiated particle-laden turbulence. *arXiv preprint arXiv:1801.06062*, 2018.

- [55] Jerrad Hampton, Hillary R Fairbanks, Akil Narayan, and Alireza Doostan. Practical error bounds for a non-intrusive bi-fidelity approach to parametric/stochastic model reduction. *Journal of Computational Physics*, 368:315–332, 2018.
- [56] Alessandra Lunardi. *Interpolation theory*, volume 9. Springer, 2009.
- [57] S. Conti and A. O’Hagan. Bayesian emulation of complex multi-output and dynamic computer models. *Journal of Statistical Planning and Inference*, 140:640–651, 2010.
- [58] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- [59] Yin Zhang. Solving large-scale linear programs by interior-point methods under the matlab environment. *Optimization Methods and Software*, 10(1):1–31, 1998.
- [60] Sanjay Mehrotra. On the implementation of a primal-dual interior point method. *SIAM Journal on Optimization*, 2(4):575–601, 1992.