



Original Research Article

Hyper-differential sensitivity analysis for inverse problems governed by ODEs with application to COVID-19 modeling

Mason Stevens, Isaac Sunseri¹, Alen Alexanderian^{*}

Department of Mathematics, North Carolina State University, Raleigh, NC, United States of America



ARTICLE INFO

Keywords:

Inverse problems
Sensitivity analysis
Uncertainty quantification
Design of experiments
Computational epidemiology

ABSTRACT

We consider inverse problems governed by systems of ordinary differential equations (ODEs) that contain uncertain parameters in addition to the parameters being estimated. In such problems, which are common in applications, it is important to understand the sensitivity of the solution of the inverse problem to the uncertain model parameters. It is also of interest to understand the sensitivity of the inverse problem solution to different types of measurements or parameters describing the experimental setup. Hyper-differential sensitivity analysis (HDSA) is a sensitivity analysis approach that provides tools for such tasks. We extend existing HDSA methods by developing methods for quantifying the uncertainty in the estimated parameters. Specifically, we propose a linear approximation to the solution of the inverse problem that allows efficiently approximating the statistical properties of the estimated parameters. We also explore the use of this linear model for approximate global sensitivity analysis. As a driving application, we consider an inverse problem governed by a COVID-19 model. We present comprehensive computational studies that examine the sensitivity of this inverse problem to several uncertain model parameters and different types of measurement data. Our results also demonstrate the effectiveness of the linear approximation model for uncertainty quantification in inverse problems and for parameter screening.

1. Introduction

Mathematical modeling has become an indispensable tool in analyzing and predicting natural phenomena. Inverse problems [1–5] arise when we are confronted with model parameters whose values we cannot determine directly. Instead, we estimate them by collecting measurements of the real-world system being modeled and finding the combination of parameter values for which the model output is most consistent with the measured data. We then use this information to quantify characteristics of the real-world system and make model-based predictions.

We call the model parameters that are estimated by solving an inverse problem the *inversion parameters*. In practice, the estimation of the inversion parameters is affected by various sources of uncertainty. In particular, in many applications, the governing model has several parameters that are not subject to parameter estimation, but are needed for a full model specification. Such parameters, which we refer to as *auxiliary parameters*, are fixed at some nominal values prior to solving the inverse problem. In practice, however, the auxiliary parameters are known only approximately and are thus uncertain. Aside from the auxiliary parameters, the solution of an inverse problem is also influenced greatly by the quantity and quality of the measurement data.

Therefore, the parameters defining the experimental conditions, which we refer to as *experimental parameters*, have an important impact on the solution of an inverse problem as well. Examples of experimental parameters include measurement error tolerances or parameters that control the design of experiments. The data measurements themselves can also be considered experimental parameters.

Both the auxiliary parameters and the experimental parameters, which we name collectively the *complementary parameters*, introduce uncertainty into the inverse problem. These issues are typical in areas such as mathematical biology where it is common to solve inverse problems governed by complex models with various source of uncertainties. In such cases, it is crucial to understand the sensitivity of the solution of the inverse problem to the uncertainties in the auxiliary and experimental parameters. Such an analysis is facilitated by hyper-differential sensitivity analysis (HDSA); see [6,7] and Section 2 where we review some basics regarding inverse problems and HDSA.

In [6,7], HDSA was introduced as a technique to efficiently compute the partial derivatives of the solution of a deterministic (as opposed to statistical) inverse problem with respect to the complementary parameters; these derivatives were then used to define suitable sensitivity indices. These works build on previous efforts such as [8–15] that

^{*} Corresponding author.

E-mail addresses: mtsteve3@ncsu.edu (M. Stevens), ipsunser@ncsu.edu (I. Sunseri), alexanderian@ncsu.edu (A. Alexanderian).

¹ Present address: Department of Mathematics, Penn State University, State College, PA.

developed methods for assessing the sensitivity of the solution of optimal control problems to perturbations in model parameters. In this article, our focus will be on deterministic inverse problems governed by ordinary differential equations (ODEs). We extend previous research to develop tools that offer insight on the sensitivity of the solution of such inverse problems to the complementary parameters and to quantify the uncertainty in the solution of an inverse problem.

Inverse problems governed by ODEs, which are common in mathematical biology, have several features that make them challenging. Although the dimension of the inversion parameter vector is typically not very large, the models are often highly nonlinear and might exhibit stiffness [16]. Also, in such problems one often has a number of auxiliary parameters such as initial states, rate constants, or source terms that are uncertain. HDSA is vital for identifying auxiliary parameters to which the inverse problem is very sensitive. This information can be used in various ways; for instance, a practitioner would take extra care in obtaining accurate values for important auxiliary parameters. Moreover, if possible, one may redesign the inverse problem by including some of the important auxiliary parameters in the set of inversion parameters. In the class of inverse problems under study, it is also common to have access to data measurements for more than one state variable. In such cases, HDSA with respect to measurement data can help identify sources of data the inverse problem is most sensitive to, hence guiding design of experiments. This is crucial in cases where it is unclear which sources of data are most informative to the inversion parameters of interest.

Performing HDSA with respect to complementary parameters provides important insight regarding an inverse problem. However, it is also desirable to quantify the uncertainty in the solution of an inverse problem caused by uncertainty in the complementary parameters. To this end, in Section 3, we outline an approach based on computing a local linear approximation to the solution of an inverse problem, considered as a function of the complementary parameters. This linear approximation can be obtained by leveraging the computations carried out when performing HDSA and provides an efficient surrogate model that can be used for uncertainty quantification in cases where the complementary parameters vary in small ranges around their nominal values. We also provide theoretical insight that relates the accuracy of the linear approximation to the variance of the complementary parameters.

As a driving application, we consider a model inverse problem governed by a system of ODEs modeling the spread of COVID-19, which has received particular attention over the past couple years due to the COVID-19 pandemic. The model inverse problem is detailed in Section 5. We deploy the HDSA methodology in the inverse problem under study and perform a detailed inspection of the sensitivities with respect to auxiliary parameters and different types of measurement data; see Sections 6.2–6.3. This provides not only important insight into the present application, but also a template for deploying HDSA in inverse problems from different application domains.

In Section 6.4, we demonstrate that the linear approximation model for the solution of the inverse problem provides an effective and efficient tool for quantifying the uncertainty in the solution of an inverse problem. In particular, we consider the auxiliary parameters as random variables distributed uniformly in small ranges around their nominal values, and quantify the resulting uncertainty in the estimated parameters. We also explore the use of the linear approximation to obtain approximate *global sensitivity information* in the auxiliary parameter domain in Section 6.5.

We summarize the key contributions of this article as follows: (i) we introduce a linear approximation model as a computationally efficient tool for estimating the uncertainty in the inverse problem solution due to the uncertainty in the complementary parameters; (ii) we present a systematic framework for HDSA of inverse problems governed by ODEs that contain inversion and complementary parameter with different units; and (iii) we perform comprehensive computational studies

for an inverse problem governed by a complex COVID-19 model. These studies provide important insight into the specific COVID-19 inverse problem used, and serve as a guide for similar analyses for inverse problems in different application domains within mathematical biology and beyond. Additionally, we detail key computational considerations regarding derivative computations, as needed in HDSA, in Section 4. In that section, we also discuss a practical approach that uses tools from complex-step differentiation [17,18] to improve the accuracy and stability of derivative computation.

2. Preliminaries on sensitivity analysis of inverse problems

In this section, we summarize the background materials as well as the basic notations needed throughout the article.

2.1. Inverse problems

As mentioned in the introduction, the process of estimating unknown parameters in a given model based on measurements of the system being described is called solving an inverse problem. In this work, the COVID-19 model we study is governed by a system of ODEs. We denote the vector of inversion parameters by $\mathbf{m}$ and the vectors of auxiliary and experimental parameters by θ_a and θ_e , respectively. The vector of the complementary parameters is given by

$$\theta = \begin{bmatrix} \theta_a \\ \theta_e \end{bmatrix}.$$

It is important to note here how we model uncertainty in the complementary parameters. For a scalar uncertain parameter κ we model our uncertainty by defining

$$\kappa = \tilde{\kappa}(1 + a\theta),$$

where κ is the parameter of interest, $\tilde{\kappa}$ is its nominal value, a is a scaling coefficient, and θ is the perturbation itself. In this work we allow $a = .05$ to scale the perturbation by 5% of the nominal parameter value. With this model we consider dimensionless perturbations θ of the uncertain complementary parameters. For simplicity of exposition, throughout the article we often refer to both the uncertain parameter itself (κ) as well as the perturbation (θ) as the complementary parameter.

Next, let $\mathcal{A}$ be the nonlinear mapping from the model inputs $(\theta_a, \mathbf{m})$ to the solution of the corresponding ODE system, let $\mathcal{Q}$ be an operator that observes the value of this model solution at a fixed set of observation times, and let $\mathbf{d}$ represent the physical measurements of the state variables taken at those observation times. The inverse problem seeks to find parameters that minimize the following objective function that measures the data misfit:

$$J(\theta, \mathbf{m}) = \|\mathcal{Q}\mathcal{A}(\theta_a, \mathbf{m}) - \mathbf{d}(\theta_e)\|_2^2, \quad (1)$$

where $\|\cdot\|_2$ denotes the vector 2-norm. The solution of the inverse problem is the parameter $\mathbf{m}^*$ that minimizes this objective function. Note that, if necessary, a regularization or penalty term can be incorporated in the definition of the objective function J in (1).

2.2. Hyper-differential sensitivity analysis of inverse problems

Hyper-differential sensitivity analysis (HDSA) [6,7] provides a systematic and computationally efficient framework to assess the sensitivity of the solution $\mathbf{m}^*$ of an inverse problem to the complementary parameters. In this approach, we think of the solution to the inverse problem, $\mathbf{m}^*$, as a function of the complementary parameters,

$$\mathbf{m}^*(\theta) = \underset{\mathbf{m}}{\operatorname{argmin}} J(\theta, \mathbf{m}). \quad (2)$$

We seek to determine the change in $\mathbf{m}^*$ due to perturbations in θ .

Let $\mathbf{g}(\theta, \mathbf{m})$ be the gradient of J with respect to $\mathbf{m}$ and let $\bar{\theta}$ be the vector of nominal values of the complementary parameters. In practice,

the inverse problem is solved using this set of nominal complementary parameters. Because $\mathbf{m}^*(\bar{\theta})$ is a minimum of J , we know that it satisfies,

$$\mathbf{g}(\bar{\theta}, \mathbf{m}^*(\bar{\theta})) = 0. \quad (3)$$

Under the reasonable assumptions [6,7] that $\mathbf{g}$ is continuously differentiable and that the matrix of partial derivatives of $\mathbf{g}$ with respect to $\mathbf{m}$ is nonsingular, we can apply the Implicit Function Theorem to $\mathbf{g}$ and conclude that $\mathbf{m}^*(\theta)$ is a continuously differentiable function in a neighborhood of $(\bar{\theta}, \mathbf{m}^*(\bar{\theta}))$. To find its partial derivatives, we implicitly differentiate through (3) and apply the Chain Rule to obtain,

$$\frac{\partial \mathbf{g}}{\partial \mathbf{m}} \frac{\partial \mathbf{m}^*}{\partial \theta} + \frac{\partial \mathbf{g}}{\partial \theta} = 0, \quad (4)$$

where the partial derivatives are evaluated with θ and $\mathbf{m}$ set to $\bar{\theta}$ and $\mathbf{m}^*(\bar{\theta})$, respectively. The derivative $\frac{\partial \mathbf{m}^*}{\partial \theta}$ can be used to quantify the sensitivity of the estimated parameters with respect to the complementary parameters. We denote

$$\mathbf{D} = \frac{\partial \mathbf{m}^*}{\partial \theta}.$$

Note that $\frac{\partial \mathbf{g}}{\partial \mathbf{m}}$ in (4) is just the Hessian of J with respect to $\mathbf{m}$, which we denote by $\mathbf{H}$. Also, we denote $\frac{\partial \mathbf{g}}{\partial \theta}$ by the symbol $\mathbf{B}$. Using these notations, (4) reads $\mathbf{H}\mathbf{D} + \mathbf{B} = \mathbf{0}$, and thus,

$$\mathbf{D} = -\mathbf{H}^{-1}\mathbf{B}. \quad (5)$$

Note that $\mathbf{D} \in \mathbb{R}^{n_m \times n_\theta}$, $\mathbf{H} \in \mathbb{R}^{n_m \times n_m}$, and $\mathbf{B} \in \mathbb{R}^{n_m \times n_\theta}$, where n_m is the number of inversion parameters and n_θ is the number of complementary parameters.

2.3. HDSA indices

The so-called *pointwise sensitivity indices* are given by,

$$S_p[i, j] = \left| \frac{1}{m_i^*} D_{ij} \right|, \quad i = 1, \dots, n_m, j = 1, \dots, n_\theta, \quad (6)$$

where m_i^* is the i th component of $\mathbf{m}^*(\bar{\theta})$. We interpret the pointwise sensitivity index $S_p[i, j]$ as how significantly a perturbation in the j th complementary parameter will affect the estimation of the i th inversion parameter. Note that each sensitivity is normalized by the estimated value of the corresponding inversion parameter m_i^* .² This normalization will result in dimensionless indices that can be used to compare the sensitivity of different inversion parameters to one another.

It is also desirable to compare the sensitivity of the entire solution vector $\mathbf{m}^*$ of the inverse problem to various types of complementary parameters. We accomplish this using *generalized sensitivity indices* [7]. First, we group related complementary parameters together into K subsets, each corresponding to a vector θ_k , $k = 1, \dots, K$. For example, when analyzing our COVID-19 model in Section 6, we group all measurements of a particular data compartment together to determine the general sensitivity of $\mathbf{m}^*$ to that data type. For each $k \in \{1, \dots, K\}$, let $\mathbf{T}_k$ be a selection operator that zeroes out components of θ that are not associated with θ_k . The generalized sensitivity index [7] corresponding to θ_k is given by,

$$S_g[k] = \max_{\|\theta\|_2=1} \|\mathbf{MDT}_k \theta\|_2, \quad (7)$$

where $\mathbf{M}$ is the diagonal (scaling) matrix

$$\mathbf{M} = \begin{bmatrix} \frac{1}{m_1^*} & & & \\ & \frac{1}{m_2^*} & & \\ & & \ddots & \\ & & & \frac{1}{m_{n_m}^*} \end{bmatrix}.$$

² If m_i^* is very close to zero, it might be necessary to replace the m_i^* in the denominator in (6) by $m_i^* + \epsilon_i$, for a user specified $\epsilon_i > 0$.

Note that $S_g[k]$ is the largest singular value of $\mathbf{MDT}_k$. We interpret this sensitivity as the maximum possible change in the entire solution vector $\mathbf{m}^*$ corresponding to a perturbation of the parameters in the k th subset. This provides a single measure of sensitivity for each subgroup, facilitating the comparison of relative importance across parameter subgroups.

3. Local uncertainty quantification in inverse problem solutions

The HDSA sensitivity indices described in Section 2 provide insight into the importance of the complementary parameters at their nominal values. In this section, we outline an approach, which uses the computations performed while calculating the HDSA indices, to enable estimating the uncertainty in the solution of an inverse problem. Specifically, we discuss the use of a local linear approximation to the mapping $\theta \mapsto \mathbf{m}^*(\theta)$ for a computationally efficient approach to quantifying the uncertainty in the solution of an inverse problem due to the uncertainty in the complementary parameters.

As before, we denote by $\bar{\theta}$ a vector of nominal values for the complementary parameters. Herein, we assume $\bar{\theta}$ is the expected value of the vector of complementary parameters. Let the nominal solution to the inverse problem be $\bar{\mathbf{m}}^* = \mathbf{m}^*(\bar{\theta})$, and consider,

$$\bar{\mathbf{D}} := \mathbf{D}(\bar{\theta}, \bar{\mathbf{m}}^*) = -\mathbf{H}(\bar{\theta}, \bar{\mathbf{m}}^*)^{-1}\mathbf{B}(\bar{\theta}, \bar{\mathbf{m}}^*).$$

The first order Taylor expansion of the mapping $\theta \mapsto \mathbf{m}^*(\theta)$, centered at $\bar{\theta}$, is given by,

$$\mathbf{m}^*(\theta) = \bar{\mathbf{m}}^* + \bar{\mathbf{D}}(\theta - \bar{\theta}) + o(\|\theta - \bar{\theta}\|).$$

Note that upon performing HDSA at $\bar{\theta}$, we have access to $\bar{\mathbf{m}}^*$ and $\bar{\mathbf{D}}$. For a θ in a neighborhood of $\bar{\theta}$ we use the approximation,

$$\mathbf{m}^*(\theta) \approx \hat{\mathbf{m}}^*(\theta) = \bar{\mathbf{m}}^* + \bar{\mathbf{D}}(\theta - \bar{\theta}) \quad (8)$$

Considering θ as a random vector, we next present a result on average error of local linear approximation of a function of θ below. This provides insight into the question of how the accuracy of the linear approximation depends on the distribution law of θ .

Proposition 1. *Let θ be a random variable that takes values in a convex set $\Theta \subseteq \mathbb{R}^m$ and has mean $\bar{\theta}$ and covariance matrix Σ . Assume $f : \Theta \rightarrow \mathbb{R}^n$ is a function that has continuous first and second order partial derivatives in the interior of Θ . Furthermore, assume that for each $i \in \{1, \dots, n\}$, there exists $K_i > 0$ such that*

$$\|\nabla^2 f_i(\theta)\|_2 \leq K_i, \quad \text{for all } \theta \in \Theta \quad (9)$$

where $\nabla^2 f_i$ denotes the Hessian of f_i , and $\|\nabla^2 f_i(\theta)\|_2$ denotes the operator 2-norm of $\nabla^2 f_i(\theta)$. Let $\hat{f}$ be the local linear approximation of f centered at $\bar{\theta}$. Then,

$$\mathbb{E}[\|f - \hat{f}\|_2] \leq K \text{trace}(\Sigma),$$

for a fixed constant $K > 0$.

Proof. By the multivariate Taylor formula, we have, for each $i \in \{1, \dots, n\}$,

$$f_i(\theta) = f_i(\bar{\theta}) + (\theta - \bar{\theta})^\top \nabla f_i(\bar{\theta}) + \frac{1}{2}(\theta - \bar{\theta})^\top \nabla^2 f_i(\xi)(\theta - \bar{\theta}),$$

for some ξ on the interior of the line segment joining θ and $\bar{\theta}$. Considering the linear approximation $\hat{f}_i(\theta) = f_i(\bar{\theta}) + (\theta - \bar{\theta})^\top \nabla f_i(\bar{\theta})$, we have that

$$\begin{aligned} |f_i(\theta) - \hat{f}_i(\theta)| &= \frac{1}{2} |(\theta - \bar{\theta})^\top \nabla^2 f_i(\xi)(\theta - \bar{\theta})| \\ &\leq \frac{1}{2} \|\nabla^2 f_i(\xi)(\theta - \bar{\theta})\|_2 \|\theta - \bar{\theta}\|_2 \\ &\leq \frac{1}{2} \|\nabla^2 f_i(\xi)\|_2 \|\theta - \bar{\theta}\|_2^2 \\ &\leq \frac{K_i}{2} \|\theta - \bar{\theta}\|_2^2. \end{aligned}$$

Therefore, letting $K = \sum_i K_i/2$,

$$\|f(\theta) - \hat{f}(\theta)\|_2 \leq \|f(\theta) - \hat{f}(\theta)\|_1 = \sum_{i=1}^n |f_i(\theta) - \hat{f}_i(\theta)| \leq K\|\theta - \bar{\theta}\|_2^2.$$

Hence, $\mathbb{E}[\|f(\theta) - \hat{f}(\theta)\|_2] \leq K\mathbb{E}[\|\theta - \bar{\theta}\|_2^2] = K\text{trace}(\Sigma)$. $\square$

Applying the above results to the mapping $\theta \mapsto m^*(\theta)$, we can see that the average error in linear approximation $\hat{m}^*(\theta)$ (centered at the nominal vector $\bar{\theta}$) is bounded by a multiple of the trace of the covariance operator of θ . To illustrate further, if θ_i 's are independent uniform random variables $\theta_i \sim U(\theta_i - \varepsilon_i, \theta_i + \varepsilon_i)$, with $\varepsilon_i > 0$, we have that $\mathbb{V}[\theta_i] = \frac{1}{12}(2\varepsilon_i)^2 = \varepsilon_i^2/3$. Thus,

$$\mathbb{E}[\|m^*(\theta) - \hat{m}^*(\theta)\|_2] \leq \frac{K}{3} \sum_{i=1}^{n_\theta} \varepsilon_i^2,$$

for some $K > 0$.

The result in Proposition 1 and the following discussion indicate that if the mapping $\theta \mapsto m^*(\theta)$ satisfies (9) with K_i 's that are not too large and the ranges of uncertainty in entries of θ are small, one could expect the local linear approximation to be a suitable surrogate model for purposes of uncertainty analysis. Note that the condition (9) concerns the size of the second order partial derivatives over the parameter domain Θ . In Section 6, we demonstrate the performance of the linear approximation to quantify the uncertainty in the solution of an inverse problem governed by a COVID-19 model. Specifically, we consider the uncertainty in auxiliary parameters, which are assumed to vary in small neighborhoods of their nominal values. In that section, we also explore ways in which this local linear model can be used for approximate global sensitivity analysis.

4. Implementation

In this section, we discuss numerical computation of the HDSA indices and neighborhood sensitivities. Our discussion focuses mainly on computing derivatives required in the process. Various approaches for computing derivatives can be used, such as adjoint-based methods [19,20], automatic differentiation [21], or finite-differences. For inverse problems governed by differential equations, one can also use a direct sensitivity analysis approach [22] to compute the sensitivities of the state variables appearing in the governing model with respect to parameters and use those derivatives in computing the HDSA indices. In this article, we focus on inverse problems governed by systems of ODEs with low-dimensional inversion and complementary parameters. Therefore, we focus on computing the derivatives via finite-differences.

Let $\{h_i\}_{i=1}^{n_m}$ and $\{\tilde{h}_i\}_{i=1}^{n_\theta}$ be stepsizes and let $\{e_i\}_{i=1}^{n_m}$ and $\{\tilde{e}_i\}_{i=1}^{n_\theta}$ be the standard bases in $\mathbb{R}^{n_m}$ and $\mathbb{R}^{n_\theta}$, respectively. Also, denote $v_i = h_i e_i$, $i = 1, \dots, n_m$ and $\tilde{v}_i = \tilde{h}_i \tilde{e}_i$, $i = 1, \dots, n_\theta$. We may use the finite difference formulas,

$$H_{jk} \approx \frac{J(m + v_j + v_k, \theta) - J(m + v_j - v_k, \theta) - J(m - v_j + v_k, \theta) + J(m - v_j - v_k, \theta)}{4h_j h_k},$$

for $j, k \in \{1, \dots, n_m\}$, and,

$$B_{jk} \approx \frac{J(m + v_j, \theta + \tilde{v}_k) - J(m + v_j, \theta - \tilde{v}_k) - J(m - v_j, \theta + \tilde{v}_k) + J(m - v_j, \theta - \tilde{v}_k)}{4h_j \tilde{h}_k},$$

for $j \in \{1, \dots, n_m\}$ and $k \in \{1, \dots, n_\theta\}$, to compute the matrices of second-order partial derivatives $\mathbf{H}$ and $\mathbf{B}$.

A drawback of using finite difference formulas is that the result might become very sensitive to the choice of steps sizes. This is due to cancellation errors introduced by subtraction. To increase the numerical stability of the finite difference calculations, we incorporate complex-step differentiation [23]. This technique extends the calculation into the complex plane to reduce the number of subtractions inherent in the formula, significantly reducing the risk of cancellation.

The first derivative of J with respect to the j th inversion parameter can be approximated using complex step as,

$$\frac{\partial J}{\partial m_j} \approx \frac{\text{Im}[J(m + i v_j)]}{h_j}, \quad (10)$$

where i is the imaginary unit, $i = \sqrt{-1}$, and $v_j = h_j e_j$ as before. The complex step approach can also be used for computing second order derivatives; see e.g., [24]. This, however, requires the implementation of multicomplex numbers, which are not native to most standard programming environments and would necessitate significant implementation efforts. Namely, for systems governed by ODEs, the ODE solvers need to know how to manipulate multicomplex numbers. To avoid this, while leveraging the compatibility of programming environments with complex numbers, we take a hybrid approach that combines the first-order complex step formulas with a first-order finite difference. This offers improved numerical stability compared to the standard finite difference approach. The augmented formulas for finding the components of $\mathbf{H}$ and $\mathbf{B}$ are,

$$H_{jk} \approx \frac{\text{Im}[J(m + v_j + i v_k, \theta)] - \text{Im}[J(m - v_j + i v_k, \theta)]}{2h_j h_k}, \quad (11)$$

$$B_{jk} \approx \frac{\text{Im}[J(m + v_j, \theta + i \tilde{v}_k)] - \text{Im}[J(m - v_j, \theta + i \tilde{v}_k)]}{2h_j \tilde{h}_k}. \quad (12)$$

This hybrid derivative implementation is used to improve the accuracy of our results presented in Section 6.

5. Model problem

In this section, we discuss a model inverse problem governed by a system of ODEs, which we use to illustrate the methods outlined in the previous sections. We focus on an SE(A)IR model, introduced in [25], that models the spread of COVID-19. This model is a variant of the classical Susceptible–Exposed–Infectious–Recovered (SEIR) epidemic model that is modified to better account for the effect of individuals who can spread the disease but are asymptomatic. This is especially relevant in modeling the spread of diseases such as COVID-19.

5.1. The forward problem

The SE(A)IR model organizes the population under study into four compartments. The majority of the population is initially in the Susceptible (S) group, which is for healthy individuals who have never been infected by the disease. The Exposed-Asymptomatic (E_A) compartment distinguishes the SE(A)IR model from the traditional SEIR model because it contains not only the individuals who have been exposed to the disease and are not yet contagious, but also those that are contagious and do not show symptoms. The members of the population that do show symptoms and are contagious are considered part of the Infectious (I) compartment. Finally, those who overcome the disease and become healthy again form the Recovered (R) group.

The movement of individuals between compartments in a given population is modeled by the following system of ODEs (13):

$$\frac{dS}{dt} = -\beta \frac{pE_A + qI}{N} S + kN - kS, \quad (13a)$$

$$\frac{dE_A}{dt} = \beta \frac{pE_A + qI}{N} S - \eta E_A - \gamma E_A - kE_A, \quad (13b)$$

$$\frac{dI}{dt} = \eta E_A - \gamma I - \mu I - kI, \quad (13c)$$

$$\frac{dR}{dt} = \gamma E_A + \gamma I - kR, \quad (13d)$$

$$N = S + E_A + I + R. \quad (13e)$$

In Table 1, we describe the various parameters in the model along with their respective nominal values.

At the beginning of the observed time period, a small portion of the population is in the Infectious compartment, with everyone else in the Susceptible one. When healthy individuals are exposed to the disease by contagious individuals, they move from the Susceptible compartment to the Exposed-Asymptomatic compartment. Once the pathogen leaves its incubation cycle and the individual becomes contagious, there are two possibilities. If they show symptoms, they

Table 1

The parameters used in the SE(A)IR model along with their nominal values. The nominal values for parameters η through q are adapted from [25]. For more details on the interpretations of each of these parameters, also see [25]. The values for k and N_0 are specific to the numerical experiments in the present work.

Parameter	Symbol	Value
Inverse of incubation period scaled by the probability of being symptomatic	η	0.1429
Transmission rate	β	0.3000
Recovery rate	γ	0.0476
Disease-related death rate	μ	0.0040
Proportion of E_A that is asymptomatic as opposed to in the incubation period	p	0.7000
Fraction of usual social contacts maintained after noticing symptoms	q	0.0700
Initial infectious population	I_0	100
Natural birth and death rate	k	0.0100
Initial total population	N_0	1×10^6

join the Infectious compartment. If they are asymptomatic, however, they stay in the Exposed-Asymptomatic group. Some symptomatic individuals may grow seriously ill and even die due to the disease; however, all asymptomatic individuals and the surviving portion of the Infectious compartment eventually recover and transition to the Recovered compartment. Members of the population are assumed to have permanent immunity after recovering from the disease. Note that we have modified the model proposed in [25] slightly by treating the transmission rate as a constant and by incorporating the natural birth and death rates, which are assumed equal here.

5.2. The inverse problem setup

The inverse problem formulation follows the general setup described in Section 2.1. For clarity, we detail the definition of the objective function J , which we seek to minimize, for the present inverse problem. To do so, we first define some notation. Recall that we take $\mathcal{A}(\theta_a, \mathbf{m})$ to represent the model solution for a given set of parameter values and $\mathbf{d}(\theta_e)$ to be the measurements collected from some physical system; see (1). For $j = 1, \dots, n_d$, let $S(t_j)$ represent the Susceptible component of $\mathcal{A}$ when evaluated at time t_j and let d_j^S represent the Susceptible component of the j th observation of the state variables, taken at time t_j . Finally, let $\bar{d}^S$ denote the average of the observations of the Susceptible compartment. We extend this notation to the other compartments of the SE(A)IR model and define the objective function J of the inverse problem as follows:

$$J(\theta, \mathbf{m}) = \sum_{j=1}^{n_d} \left[\left(\frac{S(t_j) - d_j^S}{\bar{d}^S} \right)^2 + \left(\frac{E_A(t_j) - d_j^{E_A}}{\bar{d}^{E_A}} \right)^2 + \left(\frac{I(t_j) - d_j^I}{\bar{d}^I} \right)^2 + \left(\frac{R(t_j) - d_j^R}{\bar{d}^R} \right)^2 \right].$$

Note that, in the present study, we consider the case where one uses data from all of the state variables to estimate the inversion parameters of interest. Note also that since the solution operator $\mathcal{A}$ depends on θ_a and $\mathbf{m}$, so do the observations of the model. Similarly, since $\mathbf{d}$ depends on θ_e , so do the data measurements. In fact, in the present study, we take $\mathbf{d} = \theta_e$, so there is no difference between our data measurements and our experimental parameters.

When we solve the SE(A)IR inverse problem, we specifically seek to estimate η , β , γ , and μ , which are all characteristics inherent to the disease we are studying. These are the inversion parameters in our experiments. The remaining parameters, p , q , I_0 and k , are taken to be the auxiliary parameters in our inverse problem setup. Thus, the vectors $\mathbf{m}$ and θ_a of the inversion and auxiliary parameters are given by,

$$\mathbf{m} = \begin{bmatrix} \eta \\ \beta \\ \gamma \\ \mu \end{bmatrix} \text{ and } \theta_a = \begin{bmatrix} p \\ q \\ I_0 \\ k \end{bmatrix}.$$

Note that taking p and q as auxiliary parameters was a choice we made in the present inverse problem setup and one could alternatively consider them as inversion parameters.

Next, we discuss the data measurements that will be used in solving the inverse problem. As mentioned earlier, we use data measurements from each of the S , E_A , I , and R compartments when solving the inverse problem and executing our HDSA methodology. While it would be difficult to obtain measurements of all the state variables in practice, our purpose is to demonstrate the utility of HDSA in providing insight into the relative importance of different measurements and measurement types.

In the present study, we synthesize data by solving the SE(A)IR model solution for the nominal set of parameter values listed in Table 1, over the time discretization $t = 0, 1, \dots, 365$, where t is measured in days. Specifically, after obtaining the solution, we took solution values at 120 approximately equidistant observation times. Each data measurement includes unbiased additive Gaussian noise whose standard deviation is equal to 3% of the solution's magnitude.

All results discussed in Section 6 are specific to this COVID-19 inverse problem. The exact results will vary for other epidemiological models, but the same methodologies presented in this paper can be applied to extract meaningful insights on other inverse problems.

6. Results and discussion

In this section, we examine the sensitivity of the inverse problem described in Section 5 to the complementary parameters. Additionally, we demonstrate the effectiveness of the linear approximation model discussed in Section 3 for uncertainty quantification of the inverse problem and parameter screening. We begin by solving the inverse problem in Section 6.1. Then, we apply the HDSA methodologies to understand the sensitivity of the estimated parameter vector to perturbations in the complementary parameters. We consider the pointwise sensitivity indices with respect to the complementary parameters in Section 6.2, followed by generalized sensitivity indices in Section 6.3. Next, we compute the linear approximation model presented in Section 3 to estimate the solution of the inverse problem for realizations of the auxiliary parameters within a neighborhood of their nominal values. Sampling this linear approximation provides a means to quantify the uncertainty in the solution of the inverse problem due to the uncertainty in the auxiliary parameters; see Section 6.4. Finally, in Section 6.5, we explore the use of the linear model for approximate derivative-based global sensitivity analysis.

6.1. Solving the inverse problem

As the first step, we solve the inverse problem described in Section 5. We obtain the following parameter estimate:

$$\mathbf{m}^* = [\eta^* \quad \beta^* \quad \gamma^* \quad \mu^*]^\top \approx [0.1435 \quad 0.3006 \quad 0.0476 \quad 0.0042]^\top.$$




Fig. 1. A graph of the solution to the inverse problem against the synthesized measurement data.




Fig. 2. The pointwise sensitivities of the solution to the inverse problem with respect to the auxiliary parameters. Note that the y-axis is logarithmic.

The corresponding model solution is reported in Fig. 1. We first note that with the estimated m^* , the model trajectory closely fits the synthesized data measurements. From an epidemiological point of view, we see typical behavior. The disease initially spreads slowly. After day 100, the percentage of the population in the E_A and I compartments begins growing, with the epidemic peaking around day 240. The portion of the population in the S compartment nears 75% at its minimum, then starts to increase again as the epidemic slows down and new susceptible individuals are born into the population.

6.2. Pointwise sensitivity indices

In this section, we analyze the sensitivity of the inversion parameters to the auxiliary parameters. We begin by computing the pointwise HDSA indices that quantify the sensitivity of the different components of the inversion parameter vector to each of the auxiliary parameters. The results are reported in Fig. 2.

It is interesting that, for this problem, the ordering of the importance of the auxiliary parameters is not the same for each of the inversion parameters. The natural birth and death rate k is consistently important. However, while k ranks as the most important auxiliary parameter when estimating γ and μ , it is narrowly the least important one for estimating η . It makes sense that μ and γ would be more sensitive to k than the other inversion parameters are, since k plays a similar role as μ and γ . For a fixed set of data, an increase or decrease in the given value of k would directly imply an increase or decrease in

the values of μ and γ to compensate. This relationship is not as direct for η and β , so the reduced sensitivity is a logical result.

We generally see the sensitivities with respect to p and q are similar, which we would expect since the two coefficients play similar roles in the model. However, the notable exception is in the case of the transmission rate β , for which the sensitivity to p is much higher. This result indicates that, for the given parameter values, knowing how likely someone infected is to be asymptomatic is more valuable in estimating the transmission rate of disease than the degree to which symptomatic individuals reduce their social contacts. While we expected the sensitivity of γ with respect to p and q to be small, the negligible magnitude of these sensitivities is surprisingly low. These observations are an example of how the HDSA results provide useful insight that would be difficult or impossible to obtain by simply looking at the model.

Finally, we were surprised by the relatively large magnitude of the sensitivity with respect to I_0 for each of the inversion parameters, and especially for μ . This could be due to the fact that the only part of the model in which μ appears is for dI/dt , and I_0 is the corresponding initial condition for this compartment. It seems reasonable that the sensitivities of η and β with respect to I_0 are similar and of moderate importance. This is because a similar trajectory of the epidemic early in time for a perturbed I_0 could be achieved by adjusting incubation period or the transmission rate to compensate for the difference. The fact that this relationship only holds early in the pandemic could explain why these sensitivities are lower than that for μ .

Next, we turn to the pointwise sensitivity indices with respect to the data measurements, displayed in Fig. 3. To begin, note from the bottom row of Fig. 3 that the peak of the epidemic is between times $t = 225$ and 275. The sensitivity to measurements of the Susceptible population generally increase over time, with the exception of a local maximum in the sensitivity of β right before the height of the epidemic, followed by a brief dip. Similarly, measurements of the Recovered population are relatively unimportant until the epidemic begins to peak, after which the sensitivities to those data measurements increase rapidly. Finally, the sensitivities to measurements of both the Exposed-Asymptomatic population and the Infectious population are at their maximum as the epidemic is reaching its height, after which the importance of these measurements begins to decline. It is interesting to note that the sensitivity to the measurements of the Infectious compartment quickly diminishes as we move past the height of the epidemic, and then rebounds in the following days.

The pointwise sensitivity analysis highlights the utility of HDSA in providing meaningful insight into which auxiliary parameters are most influential in reconstructing the different inversion parameters, as well as which measurements are most important to take at what times. Given the complexity of the pointwise sensitivity results, however, it is difficult to draw holistic conclusions. We leave this discussion for the next section, where we present the generalized sensitivity indices.

6.3. Generalized sensitivity indices

The detailed conclusions we draw from the pointwise sensitivity indices are useful, but we are also interested in the broader conclusions we can draw using the generalized sensitivity indices. For the auxiliary parameters, we choose our parameter groups to be the individual parameters; for the experimental parameters, we group them by the compartment to which they belong in the SE(A)IR model. Because we define our sensitivity indices to be unitless, we can actually compare the generalized sensitivity indices for both the auxiliary and experimental parameters. We show them next to each other in Fig. 4.

The natural birth and death rate stands out as the most important complementary parameter for estimating the inversion parameters. This is surprising, but could be explained by the heavy influence of this parameter on the shape of the model as the epidemic begins to fade after the 300th day. The parameters p and q are shown to be less




Fig. 3. The pointwise sensitivities of the solution to the inverse problem with respect to the measurement data from different state variables. In order, the rows correspond to the sensitivities of η , β , γ , and μ . The columns correspond to the sensitivities with respect to the S , E_A , I , and R compartments, respectively. The fifth row shows the state variable in the model for the compartment in the corresponding column, and the horizontal axis of each plot gives the number of days elapsed.




Fig. 4. The generalized sensitivities of the solution to the inverse problem with respect to the complementary parameters. Note that the y-axis is linear.

influential in finding the solution of the inverse problem, but still play a notable role. With regard to the sensitivities with respect to data, the Recovered compartment is the one to which the estimation of the inversion parameters is most sensitive. Measurements with respect to the E_A and I compartments are also significant, but less important than each of the auxiliary parameters.

It is interesting to note that the generalized sensitivity index for the auxiliary parameter I_0 is larger than that for the Infectious data compartment. A possible explanation is that the data values are the $d(\theta_e)$ term in (1), while the value of I_0 is incorporated into the $\mathcal{A}(\theta_a, m)$ term of (1). This difference of roles between the parameters in solving the inverse problem may explain the observed result.

6.4. Uncertainty quantification using the linear approximation model

In this section, we demonstrate the utility of the local linear approximation approach, described in Section 3, for quantifying the uncertainty in the solution of the inverse problem. Specifically, we focus on quantifying the impact of uncertainty in the auxiliary parameters on the solution of the inverse problem. We assume that the auxiliary parameters are distributed uniformly according to $\theta_i \sim U(-1, 1)$ for

$i = 1, \dots, 4$. Recall that, by our definition of the dimensionless θ , this corresponds to sampling values within 5% of the physical values of the auxiliary parameters.

First, we examine the accuracy of the local linear approximation to the solution of the inverse problem. To this end, we evaluate $m^*(\theta)$ and its approximation $\hat{m}^*(\theta)$ for 1000 realizations of the auxiliary parameters, which we then use to approximate the statistical properties of the relative error defined by,

$$\text{relative error} = \frac{\|\hat{m}^*(\theta) - m^*(\theta)\|_2}{\|m^*(\theta)\|_2}. \quad (14)$$

We present the probability density function (PDF) of the relative errors in Fig. 5.

These results demonstrate that, with high probability, $\hat{m}^*(\theta)$ is a close approximation for $m^*(\theta)$. Specifically, the probability of the relative error being greater than 0.3% is negligible. This shows that for the present problem, the linear approximation can be used reliably for quantifying the uncertainty in m^* , due to uncertainties in the auxiliary parameters. In particular, we can approximate the statistical properties




Fig. 5. The estimated probability density function for the 2-norm relative error between $\mathbf{m}^*(\theta)$ and $\hat{\mathbf{m}}^*(\theta)$. Note that the x-axis is logarithmic.

of $\mathbf{m}^*$ by sampling the surrogate model $\hat{\mathbf{m}}^*$. This is of great computational significance: the cost of evaluating $\hat{\mathbf{m}}^*$ is negligible compared to that of evaluating $\mathbf{m}^*$, which requires solving an inverse problem.

Next, we study the effectiveness of the linear model for approximating the distribution of each component of the inverse problem solution. In Fig. 6, we illustrate the PDFs created from the solutions to the inverse problem and their approximations gathered from the previous experiment. As seen in this figure, the local linear approximation $\hat{\mathbf{m}}^*$ provides an effective and computationally efficient means to approximate the distribution of the components of $\mathbf{m}^*$. It is important to note, however, that some tail behavior of the true distributions will be missed by the local linear model, and the discrepancy would grow as the range of uncertainty in the auxiliary parameters grows.

The generally small magnitude of the errors present in the estimation of the solution to the inverse problem opens the door to applying the linear approximation model to perform analyses that may otherwise have been prohibitively expensive. In the next section, we consider one such application—approximate global sensitivity analysis of the inverse problem solution.

6.5. Approximate global sensitivity analysis using the linear model

Standard HDSA, which we have been using up to this point in this article, is local in the set of complementary parameters. In some problems one may have access to information regarding the ranges of likely values for the complementary parameters or more detailed statistical information. In such cases, it is desirable to obtain sensitivity measures that are global in the complementary parameter domain. This can be accomplished by computing averages of the pointwise sensitivity indices provided by HDSA. For example, considering the pointwise indices $S_p[i, j]$, defined in (6), we may consider a *global* sensitivity measure,

$$\langle S_p[i, j] \rangle = \text{median}(\{S_p^{(\ell)}[i, j]\}_{\ell=1}^n), \quad (15)$$

where $S_p^{(\ell)}[i, j]$ are realizations of the pointwise indices $S_p[i, j]$ corresponding to different realizations of the complementary parameters. This approach to obtain global HDSA measures, however, is computationally expensive. Specifically, the inverse problem needs to be solved n times, and an HDSA procedure must be conducted in each case to obtain the corresponding sensitivity indices. Here we explore an *approximate global HDSA* approach that uses the local linear approximation to build a computationally efficient parameter screening procedure.

As before, denote by $\bar{\theta}$ a vector of nominal values for the complementary parameters. In particular, thinking of each parameter as being drawn from a statistical distribution that describes its range of uncertainty, we set the nominal value of that parameter to its expected value. We follow the same procedure given in Section 3 to obtain $\hat{\mathbf{m}}^*(\theta) \approx \mathbf{m}^*(\theta)$. We then use $\hat{\mathbf{m}}^*(\theta)$ to approximate the matrix $\mathbf{D}(\theta)$ of partial derivatives as follows:

$$\mathbf{D}(\theta) = -\mathbf{H}(\theta, \mathbf{m}^*(\theta))^{-1} \mathbf{B}(\theta, \mathbf{m}^*(\theta)) \approx -\mathbf{H}(\theta, \hat{\mathbf{m}}^*(\theta))^{-1} \mathbf{B}(\theta, \hat{\mathbf{m}}^*(\theta)) =: \hat{\mathbf{D}}(\theta).$$

This enables computing approximate global sensitivities, which we call *neighborhood sensitivity indices*. Namely, letting $\{\theta^{(\ell)}\}_{\ell=1}^n$ be realizations of the complementary parameters, we define the neighborhood sensitivity indices by,

$$S_p[i, j] = \text{median}\left(\left\{\left|\frac{1}{\hat{m}_i^*} \hat{D}_{ij}(\theta^{(\ell)})\right|\right\}_{\ell=1}^n\right),$$

where $\hat{m}_i^*$ is the i th component of $\hat{\mathbf{m}}^*(\theta^{(\ell)})$. Given the ways in which error present in $\hat{\mathbf{m}}^*(\theta)$ can be compounded when performing HDSA at $(\theta, \hat{\mathbf{m}}^*(\theta))$, the neighborhood sensitivity indices might not be reliable as a surrogate for the exact global HDSA indices. However, as demonstrated below, these approximate indices can be used to consistently classify every complementary parameter as important or unimportant in accurately estimating each inversion parameter. In the present work, we implement the following screening procedure:

1. group the neighborhood sensitivity indices by inversion parameter;
2. for each inversion parameter's subgroup,
 - (a) define a threshold sensitivity level as z times the largest sensitivity index in that subgroup, for some $z \in (0, 1)$;
 - (b) label the parameters whose corresponding sensitivity indices are greater than or equal to the threshold level as “important” in estimating the current inversion parameter; label the others as “unimportant”.

We stress that this approximate parameter screening approach is appropriate only in cases where the ranges of uncertainty for the complementary parameters are small. In the present work we specifically use this methodology to screen the auxiliary parameters, but it could be applied to any subset of the complementary parameters. The choice of the threshold z is problem-dependent and should be specified by the experimenter; for the present study, we chose $z = 0.1$.

If we apply the classification procedure above to both $\langle S_p \rangle$ and S_p to create “true” and “approximate” sets of labels, respectively, we can generate a confusion matrix that illustrates the accuracy of the more computationally efficient approach; see Fig. 7.

The 93.7% classification accuracy for the neighborhood sensitivity indices demonstrates that this is a reasonable approach for parameter screening. If we were to map each label back to its corresponding sensitivity index, this neighborhood analysis allows us to make the stronger claim that the “important” and “unimportant” designations assigned to each one should hold even if the auxiliary parameter values deviate by up to 5% of their expected value. While we would like to achieve perfect classification accuracy, we are limited by the accuracy of our approximation to the solution of the inverse problem. More sophisticated approximation techniques could be used, which is the subject of future work.

It should be noted that the above classification procedure can be modified to accommodate the demands of specific applications. For example, an experimenter may wish to sort the sensitivity indices and determine importance by examining the difference between consecutive indices. Alternatively, all of the sensitivity indices could be classified together instead of being grouped by corresponding inversion parameter. Finally, an experimenter could choose to define the neighborhood sensitivity index using a different metric for average, such as mean or mode, instead of median.




Fig. 6. The estimated probability density functions of the components of the approximated (blue) and true (red) solutions to the inverse problem.




Fig. 7. The alignment between the importance classifications generated by the true global HDSA indices and by the approximate neighborhood sensitivity indices. Recall that our results are strictly for the sensitivity indices with respect to the auxiliary parameters, but could also be extended to the data measurements.

7. Conclusion

In this paper, we focused on sensitivity analysis and uncertainty quantification in inverse problems governed by ODEs. Specifically, we applied the HDSA framework to a COVID-19 model governed by a system of ordinary differential equations. The pointwise local sensitivities give insight into the key times over the course of an epidemic at which different types of measurement data are most useful in accurately estimating each of the inversion parameters. The generalized sensitivities provide a means to compare the overall relative importance of each complementary parameter with respect to the solution.

We also extended HDSA by applying a local linear approximation to the solution of the inverse problem. This approximation allows us to efficiently capture the uncertainty in the solution of the inverse problem due to the uncertainty in the complementary parameters. We also explore the use of this linear approximation in performing global

sensitivity analysis in a neighborhood of the nominal parameter values. This approach provides approximate sensitivities that can be used for classifying the importance of each parameter. This approach has shown to be effective and accurate for the purposes of importance classification and serves to validate the local HDSA importance classifications in a neighborhood of the nominal parameter values.

There are several opportunities for future work. Our results with the local linear approximation were limited to sensitivities with respect to the auxiliary parameters, but the methodology could be applied to data measurements as well. Also, future work could explore use of higher order approximations for obtaining more accurate local approximations to the inverse problem solution. Another line of inquiry is the use of the linear approximation for fast estimation of variance-based sensitivity indices, which apportion the uncertainty in the solution of the inverse problem to different complementary parameters.

The developments and discussions in this article apply to broad classes of inverse problems within mathematical biology. In applications, obtaining informative and high-fidelity data is challenging. As seen in our computational results, HDSA with respect to data provides comprehensive information about the sensitivity of the inverse problem to different types of data at different measurement times. Such information provides vital insight that can guide design of experiments. HDSA also reveals sources of data that are not necessarily the most informative, but might be cheaper or easier to obtain. Using this information, one may strike a balance between the competing goals of obtaining the most informative data and controlling the cost of data acquisition. It is also important to understand the sensitivity of the inverse problem solution to the auxiliary model parameters. Specifying the values of the important auxiliary parameters would demand extra care. Moreover, if possible, the inverse problem can be reformulated so as to estimate these parameters in addition to the inversion parameters of interest. Finally, characterizing the uncertainty in the solution of an inverse problem, which can be achieved by using the local linear approximation model, provides important insight and can facilitate parameter estimation approaches that are robust with respect to model uncertainty.

Finally, the ideas of HDSA and uncertainty quantification of inverse problem solutions can be extended to analyze the sensitivity of quantities of interest, computed using estimated model parameters, to modeling or experimental uncertainties. As an example, one may use the estimated parameters in an epidemic model to compute the

basic reproduction number. One can then adapt the HDSA tools to understand the sensitivity of the computed basic reproduction number to the different complementary parameters.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was supported in part by National Science Foundation, United States of America under grant DMS-1745654. The work of A. Alexanderian was also supported in part by National Science Foundation, United States of America under grant DMS-2111044.

References

- [1] H. Engl, M. Hanke, A. Neubauer, *Regularization of Inverse Problems*, Kluwer Academic Publishers, 1996.
- [2] K. Ito, B. Jin, *Inverse Problems: Tikhonov Theory and Algorithms*, World Scientific, 2015.
- [3] J. Kaipio, E. Somersalo, *Statistical and Computational Inverse Problems*, Springer, 2005.
- [4] A. Tarantola, *Inverse Problem Theory and Methods for Model Parameter Estimation*, Society for Industrial and Applied Mathematics, 2005.
- [5] C. Vogel, *Computational Methods for Inverse Problems*, Society for Industrial and Applied Mathematics, 2002.
- [6] J. Hart, B. van Bloemen Waanders, R. Herzog, Hyper-differential sensitivity analysis of uncertain parameters in PDE-constrained optimization, *Int. J. Uncertain. Quantif.* 10 (3) (2020).
- [7] I. Sunseri, J. Hart, B. van Bloemen Waanders, A. Alexanderian, Hyper-differential sensitivity analysis for inverse problems constrained by partial differential equations, *Inverse Problems* 36 (12) (2020).
- [8] K. Brandes, R. Griesse, Quantitative stability analysis of optimal solutions in PDE-constrained optimization, *J. Comput. Appl. Math.* (2006).
- [9] C. Büskens, R. Griesse, Parametric sensitivity analysis of perturbed PDE optimal control problems with state and control constraints, *J. Optim. Theory Appl.* 131 (1) (2006) 17–35.
- [10] R. Griesse, Parametric sensitivity analysis in optimal control of a reaction-diffusion system – part II: practical methods and examples, *Optim. Methods Softw.* 19 (2) (2004) 217–242.
- [11] R. Griesse, Parametric sensitivity analysis in optimal control of a reaction diffusion system. I. Solution differentiability, *Numer. Funct. Anal. Optim.* 25 (1–2) (2004) 93–117.
- [12] R. Griesse, *Stability and Sensitivity Analysis in Optimal Control of Partial Differential Equations (Habilitation Thesis)*, Faculty of Natural Sciences, Karl-Franzens University, 2007.
- [13] R. Griesse, B. Vexler, Numerical sensitivity analysis for the quantity of interest in PDE-constrained optimization, *SIAM J. Sci. Comput.* 29 (1) (2007) 22–48.
- [14] R. Griesse, S. Volkwein, Parametric sensitivity analysis for optimal boundary control of a 3D reaction-diffusion system, in: G.D. Pillo, M. Roma (Eds.), *Nonconvex Optimization and its Applications*, Vol. 83, Springer, Berlin, 2006.
- [15] R. Griesse, A. Walther, Parametric sensitivities for optimal control problems using automatic differentiation, *Optim. Control Appl. Methods* 24 (2003) 297–314.
- [16] P. Deufhard, S. Rößlitz, *A Guide to Numerical Modelling in Systems Biology*, Vol. 12, Springer, 2015.
- [17] J.N. Lyness, C.B. Moler, Numerical differentiation of analytic functions, *SIAM J. Numer. Anal.* 4 (2) (1967) 202–210.
- [18] J. Martins, P. Sturdza, J. Alonso, The complex-step derivative approximation, *ACM Trans. Math. Softw.* 29 (3) (2003) 245–262.
- [19] Y. Cao, S. Li, L. Petzold, R. Serban, Adjoint sensitivity analysis for differential-algebraic equations: The adjoint DAE system and its numerical solution, *SIAM J. Sci. Comput.* 24 (3) (2003) 1076–1089.
- [20] M.D. Gunzburger, *Perspectives in Flow Control and Optimization*, SIAM, 2003.
- [21] L.B. Rall, G.F. Corliss, An introduction to automatic differentiation, in: *Computational Differentiation: Techniques, Applications, and Tools*, Vol. 89, SIAM, Philadelphia, PA, 1996.
- [22] T. Maly, L.R. Petzold, Numerical methods and software for sensitivity analysis of differential-algebraic systems, *Appl. Numer. Math.* 20 (1–2) (1996) 57–79.
- [23] W. Squire, G. Trapp, Using complex variables to estimate derivatives of real functions, *SIAM Rev.* 40 (1) (1998) 110–112.
- [24] J. Fike, J. Alonso, The development of hyper-dual numbers for exact second-derivative calculations, in: *49th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*, 2011, p. 886.
- [25] D. Calvetti, A. Hoover, J. Rose, E. Somersalo, Bayesian particle filter algorithm for learning epidemic dynamics, *Inverse Problems* 37 (11) (2021) 115008.