

A note on tools for prediction under uncertainty and identifiability of SIR-like dynamical systems for epidemiology

Chiara Piazzola^a, Lorenzo Tamellini^{a,*}, Raúl Tempone^{b,c}

^a*Consiglio Nazionale delle Ricerche - Istituto di Matematica Applicata e Tecnologie Informatiche “E. Magenes” (CNR-IMATI), Via Ferrata 5/A, 27100 Pavia, Italy*

^b*Alexander von Humboldt Professor in Mathematics for Uncertainty Quantification, RWTH Aachen University, Pontdriesch 14-16, 52062, Aachen, Germany*

^c*King Abdullah University of Science and Technology (KAUST) - Computer, Electrical and Mathematical Sciences & Engineering Division (CEMSE), Thuwal, 23955-6900, Saudi Arabia.*

Abstract

We provide an overview of the methods that can be used for prediction under uncertainty and data fitting of dynamical systems, and of the fundamental challenges that arise in this context. The focus is on SIR-like models, that are being commonly used when attempting to predict the trend of the COVID-19 pandemic. In particular, we raise a warning flag about identifiability of the parameters of SIR-like models; often, it might be hard to infer the correct values of the parameters from data, even for very simple models, making it non-trivial to use these models for meaningful predictions. Most of the points that we touch upon are actually generally valid for inverse problems in more general setups.

Keywords: Dynamical Systems, Mathematical Epidemiology, Uncertainty Quantification, Model Identifiability, Bayesian Inversion, Fisher Approximation

1. Introduction

This work provides an overview of the methods that can be used for prediction under uncertainty (also known as Uncertainty Quantification) and data fitting of dynamical systems, and of the fundamental challenges that arise in this context. While this work can be easily connected with the usage of SIR-like models for the COVID-19 pandemic, the discussion presented here is actually valid for compartmental models in epidemiology and for dynamical systems in general; most points would actually be valid also in the context of inverse problems with spatial inhomogeneities. We put particular emphasis on the issue of identifiability, whose possible lack might cause serious issues when attempting long-term forecasts. To make our case clearer, in this work we use synthetic data only, which gives us full control on the errors generated by the numerical identifiability procedure.

For the sake of compactness, we have chosen to not provide many technical details on the topics that we touch, but rather point the reader to the relevant bibliography. For the same reason, most of the bibliography for further reading is provided at the end of each section, rather than during the discussion. We chose, however, to keep a rather concrete register, therefore each section comes with one or two short examples. We use for this purpose simple models, with the understanding that the points raised by the examples will be even more valid for more complicated models. For a more bird’s eye view on data-informed modeling and identifiability, and their ramifications in the general society, see e.g. [1, 2]. For readers’ convenience, we report here the topic of each section and list the examples:

Section 2: SIR-like models in epidemiology

Section 3: Forward Uncertainty Quantification (UQ): tools to make predictions under uncertainty

*Corresponding author

Email addresses: chiara.piazzola@imati.cnr.it (Chiara Piazzola), tamellini@imati.cnr.it (Lorenzo Tamellini), tempone@uq.rwth-aachen.de, raul.tempone@kaust.edu.sa (Raúl Tempone)

Section 4: Sensitivity analysis: pinpointing what parameters we need to get right, and preliminary assessment of feasibility of inversion

Section 5: Inverse UQ (data fitting) as a preliminary step to tune the pdf of the parameters to the data

Section 6: An ideal UQ workflow, from data to predictions under uncertainty

Section 7: Structural identifiability

Section 8: Practical identifiability

Section 9: Discussion and conclusions: a revisited UQ workflow

Example 1: Forward UQ of a SIR model

Example 2: Computing the Sobol indices for a SIR model

Example 3: Inverse and posterior-based forward UQ of a SIR model

Example 4: Incorporating prior information on parameters in inverse UQ

Example 5: Inverse UQ when different data types have different noise levels

Example 6: Structural identifiability of a SIR model by differential algebra

Example 7: Structural identifiability of a SIR model by mapping approach

Example 8: Practical non-identifiability of a SIR model with unknown under-reporting factor

Example 9: Structural and practical identifiability of a SEIRD model

All the numerical results in the examples have been obtained with Matlab, and the source code is available at <https://sites.google.com/view/sparse-grids-kit>. Some of the examples rely on the functionalities of the Sparse Grids Matlab Kit, which is developed by some of the authors of this manuscript and can be downloaded from the same website.

2. SIR-like models in epidemiology

The recent COVID-19 pandemic has triggered an unprecedented effort among researchers worldwide¹. In the field of applied mathematics, a large share of this effort has been focusing on devising tools to forecast the trends of the epidemics.

The most widely used tools to this end are compartmental models, where individuals of a population are categorized in compartments (Infected, Recovered, Dead, etc.) and can transition from one compartment to another according to some “transition rates”. The origin of these models can be traced back to the work of Kermack and McKendrick [3]. The actual model in that paper was a system of integro-differential equations. Some simplifications allow to rewrite those equations as a non-linear system of ordinary differential equations (ODEs), whose simplest form is the SIR model:

$$\begin{cases} \dot{S} = -\frac{\beta}{N_{pop}}IS \\ \dot{I} = \frac{\beta}{N_{pop}}IS - rI \\ \dot{R} = rI, \end{cases} \quad (1)$$

which describes the time-evolution of three compartments: individuals (S)usceptible to the disease, individuals (I)nfected with the disease, and finally individuals (R)emoved from the disease dynamics (either because they recovered, assuming immunity after having contracted the disease, or died). The total number of individuals in the population $N_{pop} = S + I + R$ is supposed constant, and individuals transition from one compartment to the next one with certain transition rates β, r . Besides the ODE, the Kermack and McKendrick integro-differential equations can also be rewritten as a stochastic differential equation whose limit is the ODE equation; see [4].

Of course, a simple SIR model is insufficient to capture the dynamics of the COVID-19 disease, due to its biological peculiarities, such as the incubation time and the presence of asymptomatic carriers of the disease, as well as human interventions such as individuals in quarantine (hence with limited transmissivity) and hospitalized. Therefore, many works in the COVID-19 literature consider more complex variations of the

¹On July 14th: 1600+ preprints on arxiv.org, 5100+ preprints on medrxiv.org, 1400+ preprints on biorxiv.org

simple SIR model (1) with, for example, more compartments, time-dependent coefficients, or by introducing network models, in an attempt to better describe the dynamics of the pandemic and provide reliable forecasts of its evolution. Of course, one should always keep in mind that while more complex models have potentially a greater predictive power, they are also more complex to analyze and tune, so that one should ideally look for the model with the optimal trade-off between these two aspects.

Bibliography and further reading.

- For a survey of SIR-like models “pre-COVID-19” for diseases such as Zika, Dengue, Ebola, H1N1, see e.g. [5, 4, 6, 7, 8, 9, 10, 11].
- For some examples of SIR-like models for COVID-19, see e.g. [12, 13, 14, 15, 16, 17, 18].
- A somewhat different approach is proposed in [19], where the underlying model is a simple SIR, with a more complex model for the probability distribution of the delays between infection and the observed events (hospitalization, recovery, death).
- Control strategies for SIR-like systems are also an important topic, see e.g. [20].

3. Forward Uncertainty Quantification (UQ): tools to make predictions under uncertainty

In general, SIR-like models can be written as ODE systems for a state vector X with N_{states} components. The evolution of the system depends on N_{coef} coefficients $\mathbf{p} = [p_1, \dots, p_{N_{coef}}]$ and on the N_{states} initial conditions $\mathbf{q} = [q_1, \dots, q_{N_{states}}]$. Moreover, we might be interested in monitoring not only the states of the system but also some related quantities Y (*Quantities of Interest*, say we have N_{qoi} of them), which can be derived from X by an observation operator G , that in turn might depend on N_{hyp} hyper-parameters $\mathbf{h} = [h_1, \dots, h_{N_{hyp}}]$:

$$\begin{cases} \dot{X} = f(X, \mathbf{p}) \\ X(t_0) = \mathbf{q} \\ Y(t) = G(X(t), \mathbf{h}), \end{cases} \quad (2)$$

where $\forall t \in [0, T]$ we have $X \in \mathbb{R}^{N_{states}}$, $Y \in \mathbb{R}^{N_{qoi}}$, and $f(\cdot, \mathbf{p}) : \mathbb{R}^{N_{states}} \rightarrow \mathbb{R}^{N_{states}}$, $G(\cdot, \mathbf{h}) : \mathbb{R}^{N_{states}} \rightarrow \mathbb{R}^{N_{qoi}}$, $\mathbf{p} \in \mathbb{R}^{N_{coef}}$, $\mathbf{q} \in \mathbb{R}^{N_{states}}$, $\mathbf{h} \in \mathbb{R}^{N_{hyp}}$.

We collect coefficients and initial conditions in a vector $\boldsymbol{\vartheta} = [\mathbf{p}, \mathbf{q}]$ with $N_{\boldsymbol{\vartheta}} = N_{coef} + N_{states}$ components. Throughout the manuscript, we refer to $\boldsymbol{\vartheta}$ as *parameters*, and we will write $X(\boldsymbol{\vartheta}), Y(\boldsymbol{\vartheta}, \mathbf{h})$ to emphasize the dependence of the states and quantities of interest on parameters and hyper-parameters. For SIR, $\mathbf{p} = [\beta, r]$, and Y might be for instance:

- the *prevalence data*, i.e., the number of infected individuals at a specific time, $G(X(t)) = I(t)$;
- the *incidence data*, i.e., the number of new cases over the reference time period, $G(X(t)) = \frac{\beta}{N_{pop}} I(t)S(t)$;
- the *peak-time of number of infected persons*: $G(X) = \arg \max_{t \in [0, T]} I(t)$;
- the *peak-time of the cumulative incidence data* i.e., the new infected persons in a time-window of length Δ : $G(X, \Delta) = \arg \max_t \int_t^{t+\Delta} \frac{\beta}{N_{pop}} S(s)I(s)ds$.

Another important scenario is under-reporting, where we assume that due to insufficient measurements, we observe only a fraction K of the total number of infected, $G(X(t), K) = \frac{1}{K} I(t)$ (K being possibly unknown); similar under-reporting scenarios could be of course conceived also for other quantities. Observe that at this point of the manuscript we are not concerned whether these quantities are reasonably easy to measure and obtain: for instance, for most infections data available are typically incidence rather than prevalence. At this level, we are just giving examples of the mathematical setup of the problem, in an ideal scenario where we have access to all sort of measurements.

Typically, most of the parameters (and possibly the hyper-parameters as well) are not known exactly and they are either taken from literature or calibrated from data. We can then assume that these parameters are random variables with a certain probability density function (pdf): for instance, uniform random variables




Figure 1: Left: SIR dynamics. The colored lines represent the SIR dynamics obtained by Monte Carlo samples of β, r , and the thick black lines represent the average computed by sparse grids. The other panels represent pdfs of quantities of interest: SIR states at $T = 30, 100$ (by Monte Carlo and sparse grids), peak-time for I and peak-value for I (sparse grids only).

over a variability range, or Gaussian random variables centered around a most likely value.² Then, a natural question is: how does the variability of the parameters impact the quantities of interest Y of the SIR-like model at hand? Or otherwise, what is the variability range of Y as the parameters range over their values?

This kind of analysis is known as *forward UQ* in computational science and engineering. The most straightforward way to accomplish this task is by sampling methods, i.e., by generating M samples of the parameters $\vartheta_1, \vartheta_2, \dots, \vartheta_M$ according to their probability distribution, solving the SIR-like system for each ϑ_i , and estimating statistics such as mean, standard deviation, confidence bands, and the probability density function of Y from the corresponding values $Y(\vartheta_1, \mathbf{h}), Y(\vartheta_2, \mathbf{h}), \dots, Y(\vartheta_M, \mathbf{h})$. The easiest sampling scheme is Monte Carlo, but more advanced sampling techniques can be used (Latin Hypercube Sampling, Stratified Sampling, Quasi-Monte-Carlo, Sparse grids, among others; see the bibliography at the end of this section). Within sampling schemes, quantities such as mean, standard deviations and higher moments with respect to the parameters can be computed by averaging over the M samples of the model results:

$$\mathbb{E}_{\vartheta}[Y(\vartheta, \mathbf{h}, t)] \approx \sum_{i=1}^M \omega_i Y(\vartheta_i, \mathbf{h}, t), \quad (3)$$

where the parameter samples ϑ_i and the weights ω_i depend on the specific sampling method used. For instance, Monte Carlo employs random ϑ_i and $\omega_i = \frac{1}{M}$. The probability density function of the quantity of interest can be approximated by, for example, histograms or kernel density estimates, see e.g. [21, 22].

Example 1 (Forward UQ of a SIR model). *Consider a SIR model with initial conditions $S(0) = 0.95, I(0) = 0.05, R(0) = 0$. The survey on the literature performed by [16] suggests these ranges for the parameters: $\beta \in [0.25, 0.35], r \in [0.06, 0.18]$. We assume that a-priori we have no knowledge that any value of β, r is more plausible than others; therefore, we assume that β, r are uniform independent random variables. We solve the SIR system with Matlab's `ode45` up to final time $T = 150$.*

Figure 1-left shows the SIR dynamics obtained by 100 Monte Carlo samples. The black lines are the average trajectories of SIR obtained by sampling values of β and r (65 samples with sparse grids sampling). The remaining panels show pdfs of quantities of interest of SIR: SIR states at $T = 30$ (after the average peak position) and $T = 100$ (when the dynamics is over), again computed both with sparse grids (solid line) and Monte Carlo (circle markers); peak time and peak intensity (we only show the pdf obtained by sparse grids). Figure 2 shows the so-called response surface, i.e., a plot showing how the quantity of interest changes as β and r vary in their range (we report only those obtained by sparse grids). Response surfaces are useful to derive information on the general trends of the system, and to quickly approximate the value of a quantity of interest without evaluating the full model, i.e., they act as a surrogate model for the dynamical system; the pdfs of the quantities of interest obtained by sparse grids have actually been obtained by querying these response surfaces rather than the full model.

Bibliography and further reading.

²Technically, most parameters of compartmental models must be positive, so a Gaussian random variable is not suitable and one should consider other random variables, e.g. Beta or log-normal. However, here we are keeping things on a simple/introductory level, and in any case a Gaussian random variable can be truncated to ensure positivity if needed.




Figure 2: Sparse grids response surfaces for the SIR quantities of interest. From left to right: S at $T = 30$, I at $T = 30$, R at $T = 30$, peak time, and peak value.

- For random sampling methods (Monte Carlo, Latin Hypercube Sampling, Stratified Sampling), see [23, 24, 25]. Random sampling methods are robust and easy to implement, but have a poor accuracy (typically proportional to $M^{-1/2}$, with M denoting the number of samples).
- For sparse grids sampling methods, see [26, 27, 28, 29]. These are deterministic (i.e., non-random) sampling schemes that generalize tensor (cartesian) grid sampling when the parameter space is high-dimensional, in which case a cartesian grid sampling scheme would be too expensive. They are less straightforward than random sampling methods, but guarantee greater accuracy, at least for problems up to a few tens of parameters. These tools have been developed in the context of UQ for models that are expensive to evaluate, whereas evaluating a SIR-like model is typically very fast. Therefore, their use is not as crucial in the context of COVID-19, and random sampling methods might be favored for their straightforwardness. Sparse grids have still an advantage over random sampling for sensitivity analysis, see the next section.
- A somewhat intermediate possibility are Quasi Monte Carlo sampling methods, such as Sobol or Halton sequences. These are also deterministic sampling schemes as well, that aim at covering the space of parameters in the “most uniform way” (space filling) [23, 24, 30]. They typically have accuracy proportional to M^{-1} , with M denoting the number of samples.

4. Sensitivity analysis: pinpointing what parameters we need to get right, and preliminary assessment of feasibility of inversion

Sensitivity analysis aims at assessing which parameters have the largest impact on the quantities of interest. This information is crucial to determine which parameters should be subjected to further investigations to reduce their variability. The sensitivity analysis can be *local* or *global*:

- local sensitivity analysis is usually based on the derivatives of the quantities of interest with respect to ϑ , upon fixing each ϑ_i at some representative value (average, median, mode).³
- global sensitivity analysis considers the total variability of a quantity of interest and decomposes such variability into elementary components, each due to ϑ_i individually or to mixed effects such as $\vartheta_i\vartheta_j$, $\vartheta_i\vartheta_j\vartheta_k, \dots$: the larger the component, the more sensitive the quantity of interest to ϑ_i is.

In these short notes we focus on the Sobol indices for global sensitivity analysis, which are variance-based indices, i.e. the variability of the quantity of interest is measured as its variance [31]. The total variance is decomposed as follows: one term due to each ϑ_i ; one term due to each mixed effect composed of two parameters, $\vartheta_i\vartheta_j$; one term for each mixed effect composed of three parameters $\vartheta_i\vartheta_j\vartheta_k$, and so forth. The

³Of course here we are assuming that the quantity of interest depends smoothly on ϑ , so that it is possible to compute the derivatives. This is not always obvious and should be checked.




Figure 3: Time-evolution of Sobol indices for the SIR compartments, and SIR trajectories obtained by reducing the range of variability of r . The principal indices are represented by the solid line, while the total indices are represented by the dashed line.

sum is then normalized to one, and each quantity thus obtained is called Sobol index:

$$1 = \sum_{i=1}^{N_{\vartheta}} s_i + \sum_{i,j=1, i \neq j}^{N_{\vartheta}} s_{ij} + \sum_{i,j,k=1, i \neq j \neq k}^{N_{\vartheta}} s_{ijk} + \dots \quad (4)$$

The Sobol index of each parameter ϑ_i per se, s_i , is usually reported as an indicator of the importance of each parameter, and is called the *principal Sobol index*. Another relevant quantity is the *total Sobol index* of a parameter, s_i^T , which is obtained by adding to the principal Sobol index of ϑ_i all the Sobol indices of mixed effects of which ϑ_i is part, e.g. for ϑ_1 :

$$s_1^T = s_1 + \sum_{j=1, j \neq 1}^{N_{\vartheta}} s_{1j} + \sum_{j,k=1, j \neq 1, k \neq 1}^{N_{\vartheta}} s_{1jk} + \dots$$

Note that in general $\sum_{i=1}^{N_{\vartheta}} s_i < 1$ and $\sum_{i=1}^{N_{\vartheta}} s_i^T > 1$. This approach bears many similarities with the ANOVA decomposition in statistics. An important observation is that, with an eye to parameter identification, we can expect that if the Sobol index of a parameter is small, it will be hard to recover its value from measurements of the quantity of interest. We also remark that as a general rule of thumb, the larger the range of values of a parameter, the larger its corresponding Sobol indices.

Example 2 (Computing the Sobol indices for a SIR model). *Consider again the SIR example of the previous section. The Sobol decomposition of any quantity of interest Y reads $1 = s_{\beta} + s_r + s_{\beta r}$ and the total Sobol indices can be computed as $s_{\beta}^T = s_{\beta} + s_{\beta r}$, $s_r^T = s_r + s_{\beta r}$. Figure 3 shows the time-evolution of the Sobol indices (principal and total) for the SIR states: the principal indices are represented by the solid line, while the total indices are represented by the dashed line. The principal and total indices behave very similarly, indicating that the interaction between the two parameters is quite limited. Note that the Sobol indices are not constant in time and behave differently for the different compartments. More specifically, the asymptotic regime is mostly dictated by r for all the compartments, while β impacts more in the transient regime, especially in the case of the compartment R . This has an impact on the inversion procedure. In particular, severe difficulties in the estimation of β can be encountered if the data of R are missing or too noisy. Moreover, note that the influence of r is larger, in general. Further evidence of this is shown in the right-most panel, where we show the variability of the trajectories if the range of r is reduced to $[0.06, 0.1]$. The overall variability is greatly reduced, as expected.*

Bibliography and further reading.

- Classical books on sensitivity are e.g. [32, 33].
- Sobol indices can be computed either by Monte Carlo sampling, see e.g. [33], or perhaps more conveniently by polynomial expansions or sparse grids sampling, see e.g. [34].
- Sobol indices analyses for SIR-like problems can be found in [35, 5].
- An alternative to Sobol indices are the Morris indices [36].
- Sobol indices can also be computed with respect to measures of variability other than variance, see e.g. [37].

5. Inverse UQ (data fitting) as a preliminary step to tune the pdf of the parameters to the data

The sections above discussed some general elementary tools to perform predictions under uncertainty once the pdfs for the parameters have been chosen. This section discusses the preliminary step to the UQ process, i.e. how to construct pdfs for the parameters, and in particular how to do so by merging prior information on the parameters and the available data. Upon deriving these data-informed pdfs, we will use them to carry out the UQ analysis. In literature, data-informed pdfs are often called *posterior pdfs*, ρ_{post} , as opposed to *prior pdfs*, ρ_{prior} , before the data are available.

Computing the posterior pdfs of the parameters can be done by means of the Bayes theorem on conditional probabilities. This procedure is quite general and includes as a special case the least-squares approach for data fitting (this connection will be made clearer later). For ease of exposition, we exemplify the procedure over a specific example - extension to other problems is relatively straightforward (see e.g. Examples 4, 5 later on). For now we assume that:

- We have at our disposal N_{meas} measurements of the I state and N_{meas} measurements of the R state, at equispaced times $t_i = i\Delta t$, $i = 1, 2, 3, \dots, N_{meas}$. In total we have $2N_{meas}$ data, $\mathcal{D} = \{\hat{I}_1, \hat{I}_2, \dots, \hat{R}_1, \hat{R}_2, \dots\}$. As stated previously, please note that we are not hinting in any way that having measurements of both $I(t)$ (i.e., prevalence data) and $R(t)$ is what happens in real scenarios: we make this (quite restrictive in practice) assumption to put ourselves in the easiest scenario to illustrate the mathematical procedure. Using more realistic data would add mathematical technicalities without giving any further insight;
- These data correspond to some values $\boldsymbol{\vartheta}_{true}$ of coefficients and initial conditions of the system (2);
- The prior pdfs for $\boldsymbol{\vartheta}$ are uniform (see Example 4 for an example with Gaussian priors);
- Our measurements are under-reported by a factor K , i.e., we are able to measure only a fraction of the actual compartments;
- Data are noisy, i.e. affected by some random errors $\epsilon_{I,i}$, $\epsilon_{R,i}$, that are modeled by independent random variables with zero mean and standard deviation σ ;
- the standard deviation σ is identical for I and R (see Example 5 for the generalization to the case where the two compartments have different σ);
- $\epsilon_{I,i}$, $\epsilon_{R,i}$ are Gaussian random variables $\mathcal{N}(0, \sigma^2)$. See discussion at the end of the section for bibliography on more general models;
- K, σ are hyper-parameters constant in time. We assume for the moment that K is known and σ is unknown (see Examples 6, 7, 8 for a discussion on how to determine K in case it is assumed unknown as well).

In formulas our data model is the following:

$$\begin{cases} \hat{I}_i = G(I(\boldsymbol{\vartheta}_{true}, t_i), K) + \epsilon_{I,i} = \frac{1}{K}I(\boldsymbol{\vartheta}_{true}, t_i) + \epsilon_{I,i}, & i = 1, 2, \dots, N_{meas} \\ \hat{R}_i = G(R(\boldsymbol{\vartheta}_{true}, t_i), K) + \epsilon_{R,i} = \frac{1}{K}R(\boldsymbol{\vartheta}_{true}, t_i) + \epsilon_{R,i}, & i = 1, 2, \dots, N_{meas}. \end{cases} \quad (5)$$

We also introduce the $2N_{meas}$ misfits $\mathcal{M} = \{M_{I,1}, M_{I,2}, \dots, M_{R,1}, M_{R,2}, \dots\}$ between the data and the model predictions, obtained upon fixing the parameters at some estimate $\boldsymbol{\vartheta}_{guess}$ of $\boldsymbol{\vartheta}_{true}$:

$$\begin{cases} M_{I,i}(\boldsymbol{\vartheta}_{guess}) = \hat{I}_i - \frac{1}{K}I(\boldsymbol{\vartheta}_{guess}, t_i), & i = 1, 2, \dots, N_{meas} \\ M_{R,i}(\boldsymbol{\vartheta}_{guess}) = \hat{R}_i - \frac{1}{K}R(\boldsymbol{\vartheta}_{guess}, t_i), & i = 1, 2, \dots, N_{meas}. \end{cases} \quad (6)$$

5.1. Bayes Theorem and posterior distributions

The Bayes theorem provides us with a practical formula to compute the posterior pdf of the parameters $\boldsymbol{\vartheta}$, i.e., with a means of adjusting the prior pdf to the data at hand. An informal writing of the Bayes formula is

$$\text{pdf}(\boldsymbol{\vartheta} \text{ given } \mathcal{M}) = \text{pdf}(\mathcal{M} \text{ given } \boldsymbol{\vartheta}) \times \text{pdf}(\boldsymbol{\vartheta}) \times \frac{1}{\text{pdf}(\mathcal{M})} \quad (7)$$

where “pdf($\boldsymbol{\vartheta}$ given \mathcal{M})” is the pdf of the parameters when given the misfits, hence given the data (i.e., the posterior pdf that we aim at computing), while “pdf($\boldsymbol{\vartheta}$)” is the pdf of the parameters based only on a-priori information. The “pdf(\mathcal{M})” can be simply considered to be the normalization constant such that the posterior pdf is actually a pdf (i.e., its integral is equal to 1). Therefore, to make the computation of the posterior pdf practical we only need to know the expression of the “pdf(\mathcal{M} given $\boldsymbol{\vartheta}$)”, which is the so-called *likelihood function*; we will denote this quantity as $\mathcal{L}(\boldsymbol{\vartheta})$.

Deriving an expression for $\mathcal{L}(\boldsymbol{\vartheta})$ is quite straightforward. If $\boldsymbol{\vartheta}_{guess}$ were the true values, then the probability that the misfits \mathcal{M} have certain values is the probability that the measurement errors $\epsilon_{I,i}, \epsilon_{R,i}$ have those values (cf. Equations (5) and (6)). By assumption, we know that $\epsilon_{I,i}, \epsilon_{R,i}$ are independent Gaussian random variables with zero mean and standard deviation σ , therefore,

$$\mathcal{L}(\boldsymbol{\vartheta}) = \prod_{i=1}^{N_{meas}} \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2\sigma^2}(\frac{1}{K}I(\boldsymbol{\vartheta},t_i) - \hat{I}_i)^2} \prod_{i=1}^{N_{meas}} \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2\sigma^2}(\frac{1}{K}R(\boldsymbol{\vartheta},t_i) - \hat{R}_i)^2} \quad (8)$$

so that the posterior pdf of the parameters reads

$$\rho_{post}(\boldsymbol{\vartheta}|\mathcal{D}) \propto \mathcal{L}(\boldsymbol{\vartheta})\rho_{prior}(\boldsymbol{\vartheta}) = \prod_{i=1}^{N_{meas}} \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2\sigma^2}(\frac{1}{K}I(\boldsymbol{\vartheta},t_i) - \hat{I}_i)^2} \prod_{i=1}^{N_{meas}} \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2\sigma^2}(\frac{1}{K}R(\boldsymbol{\vartheta},t_i) - \hat{R}_i)^2} \rho_{prior}(\boldsymbol{\vartheta}) \quad (9)$$

where the \propto symbol is used to signify that we have omitted the normalization constant.

5.2. Computational challenges of working with the posterior pdf

Equipped with (9), we would then only need to proceed as in Section 3 and perform the UQ analysis. Although conceptually straightforward, this approach can be practically challenging, because it is not easy to obtain samples of the random parameters $\boldsymbol{\vartheta}$ distributed according to the posterior pdf (9). The classical computational tool to this end is the so-called Markov-Chain Monte Carlo - MCMC [38], which generates a sequence of proposed values of $\boldsymbol{\vartheta}$ that are *asymptotically* distributed according to ρ_{post} . A nice feature of MCMC algorithms is that they do not require knowledge of the normalization constant. The use of MCMC for forward UQ has however some drawbacks:

1. the likelihood function has to be evaluated at every proposed $\boldsymbol{\vartheta}$, which requires evaluating the SIR-like model. Even if evaluating SIR-like models for a single choice of $\boldsymbol{\vartheta}$ is quite cheap, this procedure can be overall expensive, bearing in mind that until the sequence of generated $\boldsymbol{\vartheta}$ enters in the asymptotic regime, the values generated have to be discarded because they are not distributed according to ρ_{post} .
2. most MCMC algorithms proceed by acceptance-rejection criteria, where a new value of $\boldsymbol{\vartheta}$ is generated and then rejected if doesn't agree with certain criteria; this leads to a further increase in the number of model evaluations;
3. the forward UQ analysis based on the MCMC samples is a Monte Carlo analysis, which needs many samples of $\boldsymbol{\vartheta}$ to provide an accurate estimate (the accuracy being proportional to the inverse of the square root of the number of samples as already discussed – or more precisely, the inverse of the square root of the number of accepted samples upon having entered the asymptotic regime).
4. the design of an efficient MCMC algorithm (effective proposal strategies with low rejection rate, quick to enter the asymptotic regime) might be non-trivial.

5.3. Gaussian approximation of the posterior: Maximum Likelihood Estimate (MLE) and Fisher approximation

Instead of using an MCMC approach, the strategy we employ here is to approximate ρ_{post} with a multivariate Gaussian distribution with mean $\boldsymbol{\mu}_G$ and covariance matrix Σ_G ; this is also called *Fisher approximation*. This approximation, provided that the available data are sufficient to determine the parameters, is in general more and more accurate as more data become available, i.e., as $N_{meas} \rightarrow \infty$, and it has the advantage that upon doing so, it is much easier to perform the UQ analysis, because obtaining samples from Gaussian random variables is a standard task. It has, however, some disadvantages that will be made clearer in the later sections, when discussing identifiability of the system: in a nutshell, we can already reveal that the problem is that the Fisher approximation *assumes identifiability of the system*, but this is not always true in practice and whether the system is identifiable or not should be checked beforehand. MCMC instead does not assume identifiability, and can in principle be used even when the system is not identifiable: dealing with a non-identifiable system is, however, intrinsically difficult and care needs to be taken also when tackling it using MCMC methods, as will be made clearer later on.

The Gaussian approximation is centered at the point of maximum of the posterior pdf (maximum a posteriori estimate, MAP). Since we have further made the assumption that the prior pdf for $\boldsymbol{\vartheta}$ is uniform (see Example 4 for the extension to the case of non-uniform prior), this is equivalent to computing the value $\boldsymbol{\vartheta}$ where the likelihood function is maximized; this point is generally known as the *Maximum Likelihood Estimate (MLE)* for $\boldsymbol{\vartheta}$:

$$\boldsymbol{\mu}_G = \boldsymbol{\vartheta}_{MLE} = \arg \max_{\boldsymbol{\vartheta}} \mathcal{L}_{\mathcal{D}}(\boldsymbol{\vartheta}).$$

In practice, it is numerically more convenient to work with the logarithm of the likelihood, and to recast the problem as a minimization problem, i.e., to compute $\boldsymbol{\mu}_G$ as

$$\boldsymbol{\vartheta}_{MLE} = \arg \min_{\boldsymbol{\vartheta}} NLL(\boldsymbol{\vartheta}), \quad NLL(\boldsymbol{\vartheta}) = -2 \log(\mathcal{L}_{\mathcal{D}}(\boldsymbol{\vartheta})). \quad (10)$$

The function $NLL(\boldsymbol{\vartheta})$ is called *negative log-likelihood*, and in the particular case where the noise affecting the data is assumed to be Gaussian random variables (such as in our case), this problem is equivalent to the least-squares estimate of $\boldsymbol{\vartheta}$. Indeed, it is straightforward to combine (8) and (10) to obtain⁴

$$\boldsymbol{\vartheta}_{MLE} = \arg \min_{\boldsymbol{\vartheta}} \frac{1}{\sigma^2} \left[\sum_{i=1}^{N_{meas}} \left(\frac{1}{K} I(\boldsymbol{\vartheta}, t_i) - \hat{I}_i \right)^2 + \sum_{i=1}^{N_{meas}} \left(\frac{1}{K} R(\boldsymbol{\vartheta}, t_i) - \hat{R}_i \right)^2 \right]. \quad (11)$$

Observe that this formulation does not require prior information on the value of the noise variance σ^2 : if σ^2 is unknown, it can be recovered as the sample variance of the misfits at $\boldsymbol{\vartheta}_{MLE}$

$$\sigma^2 \approx \sigma_{MLE}^2 = \frac{1}{2N_{meas}} \left[\sum_{i=1}^{N_{meas}} \left(\frac{1}{K} I(\boldsymbol{\vartheta}_{MLE}, t_i) - \hat{I}_i \right)^2 + \sum_{i=1}^{N_{meas}} \left(\frac{1}{K} R(\boldsymbol{\vartheta}_{MLE}, t_i) - \hat{R}_i \right)^2 \right]. \quad (12)$$

The covariance matrix Σ_G of the Gaussian approximation can be chosen as the inverse of the Hessian of the NLL at the MLE of the parameters $\boldsymbol{\vartheta}_{MLE}$, see e.g. [39]:

$$\Sigma_G = H^{-1}, \quad H_{i,j} = \frac{\partial^2}{\partial \vartheta_i \partial \vartheta_j} NLL(\boldsymbol{\vartheta}) \Big|_{\boldsymbol{\vartheta}=\boldsymbol{\vartheta}_{MLE}}. \quad (13)$$

The matrix H is also called the *Fisher Information Matrix*. The diagonal entries of Σ_G are the variances of the posterior pdfs of the Gaussian approximations of the parameters, i.e.,

$$\vartheta_i \sim \mathcal{N}([\boldsymbol{\vartheta}_{MLE}]_i, [\Sigma_G]_{i,i}). \quad (14)$$

This formula quantifies the intuitive fact that the precision of the MLE is related to how narrow the minimum of the NLL at $\boldsymbol{\vartheta}_{MLE}$ is. A deep, narrow minimum means that moving even slightly from $\boldsymbol{\vartheta}_{MLE}$ will change

⁴the full NLL includes additional terms in $\log(\sigma)$ and $\log(2\pi)$ that we can however drop in the minimization process, since they do not depend on $\boldsymbol{\vartheta}$.

consistently the value of NLL; therefore, we have a significant evidence that the estimate is precise. Conversely, a shallow minimum means that the MLE is not very reliable. At $\boldsymbol{\vartheta}_{MLE}$ the Hessian is positive definite, with large eigenvalues if the minimum is narrow; therefore, its inverse has small eigenvalues, and in general small diagonal entries, that can be used as variances of the parameters (the opposite is true for a shallow minimum: the Hessian has small eigenvalues, which means that the diagonal entries of its inverse will be large, and consequently the variances of the parameters will be large). As already mentioned, approximating the true posterior with equation (14) is in general more and more valid as more data become available, provided that the system is identifiable, as we will make clear below.⁵ Given the expression of the likelihood in equation (8), we can derive an expression for H as follows:

$$H_{i,j} = \sum_{m=1}^{N_{meas}} \frac{1}{K\sigma^2} \left[\frac{1}{K} \frac{\partial}{\partial y_i} I_m(\boldsymbol{\vartheta}_{MLE}) \frac{\partial}{\partial y_j} I_m(\boldsymbol{\vartheta}_{MLE}) + \left(\frac{1}{K} I_m(\boldsymbol{\vartheta}_{MLE}) - \hat{I}_m \right) \frac{\partial^2}{\partial \vartheta_i \partial \vartheta_j} I_m(\boldsymbol{\vartheta}_{MLE}) \right] + \sum_{m=1}^{N_{meas}} \frac{1}{K\sigma^2} \left[\frac{1}{K} \frac{\partial}{\partial y_i} R_m(\boldsymbol{\vartheta}_{MLE}) \frac{\partial}{\partial y_j} R_m(\boldsymbol{\vartheta}_{MLE}) + \left(\frac{1}{K} R_m(\boldsymbol{\vartheta}_{MLE}) - \hat{R}_m \right) \frac{\partial^2}{\partial \vartheta_i \partial \vartheta_j} R_m(\boldsymbol{\vartheta}_{MLE}) \right], \quad (15)$$

with $I_m(\boldsymbol{\vartheta}_{MLE}) := I(\boldsymbol{\vartheta}_{MLE}, t_m)$ and $R_m(\boldsymbol{\vartheta}_{MLE}) := R(\boldsymbol{\vartheta}_{MLE}, t_m)$. Usually, the terms involving the second derivatives of I and R are dropped because they are smaller than the other terms: this is because either the misfits at $\boldsymbol{\vartheta} = \boldsymbol{\vartheta}_{MLE}$ are small or because of near-linearity of the models $I_m(\boldsymbol{\vartheta})$, $R_m(\boldsymbol{\vartheta})$ close to the solution, i.e., $\partial_{\vartheta_i, \vartheta_j} I_m(\boldsymbol{\vartheta}_{MLE})$ and $\partial_{\vartheta_i, \vartheta_j} R_m(\boldsymbol{\vartheta}_{MLE})$ are small [41, Chap. 10]. Collecting all the derivatives of the model predictions with respect to the parameters in the Jacobian matrix J_{IR} , we can write in compact form⁶

$$H_{i,j} \approx \frac{1}{K^2 \sigma^2} J_{IR}(\boldsymbol{\vartheta}_{MLE})^T J_{IR}(\boldsymbol{\vartheta}_{MLE}), \quad (16)$$

$$[J_{IR}]_{m,i} = \frac{\partial}{\partial \vartheta_i} I_m(\boldsymbol{\vartheta}_{MLE}), \quad m = 1, 2, \dots, N_{meas},$$

$$[J_{IR}]_{m+N_{meas},i} = \frac{\partial}{\partial \vartheta_i} R_m(\boldsymbol{\vartheta}_{MLE}), \quad m = 1, 2, \dots, N_{meas}.$$

Finally, we make an important remark: minimizing the NLL to compute $\boldsymbol{\vartheta}_{MLE}$ requires repeatedly evaluating the SIR-like model for the various parameters $\boldsymbol{\vartheta}$ proposed by the optimizer. The minimization procedure should be repeated several times with different starting guesses, to avoid local minima.

Example 3 (Inverse and posterior-based forward UQ of a SIR model). *In this example we show the results of the inversion procedure using artificial/synthetic data, by fixing the values of the parameters to $\boldsymbol{\vartheta}_{true} = [0.29, 0.09]$, adding numerical Gaussian noise with $\sigma = 0.025$, discount factor to $K = 3$, considering data collected at $t = 1, 2, \dots, 30$, and verifying the results of the inversion procedure. We then perform the forward UQ based on the posterior pdf.*

Regarding the results of the inversion procedure, we expect to see that ρ_{post} is centered close to the true value of the parameters with a reasonably small variance, i.e. $\boldsymbol{\vartheta}_{MLE} \approx \boldsymbol{\vartheta}_{true}$ and $[\Sigma_G]_{i,i}$ such that the support of $\rho_{post}(\vartheta_i)$ is smaller than the support of $\rho_{prior}(\vartheta_i)$. We also expect σ_{MLE}^2 to be a reasonable approximation of the true σ^2 . Regarding the subsequent forward UQ, we expect to see that the uncertainty in the prediction is smaller than what would be obtained by using the prior information only, and the expected values of the quantities of interest is closer to the true values when using the posterior pdf than when using the prior. The resulting estimates for the parameters obtained from the inverse UQ are

$$\beta = 0.2848, \quad r = 0.0861, \quad \sigma = 0.02791.$$

The estimated covariances computed using the full Hessian (that we can compute directly by centered finite differences in this simple test) and with Equation (16) (where the Jacobian entries are also computed by

⁵The statistical interpretation of this fact is that the MLE is asymptotically Gaussian distributed with covariance matrix equal to the inverse of the Fisher Information Matrix. This means that it is an *efficient* estimator, because it reaches the Cramer-Rao lower bound on the variance of estimators [40].

⁶Sometimes the term *Fisher Information Matrix* is used to indicate this approximation rather than the full Hessian.

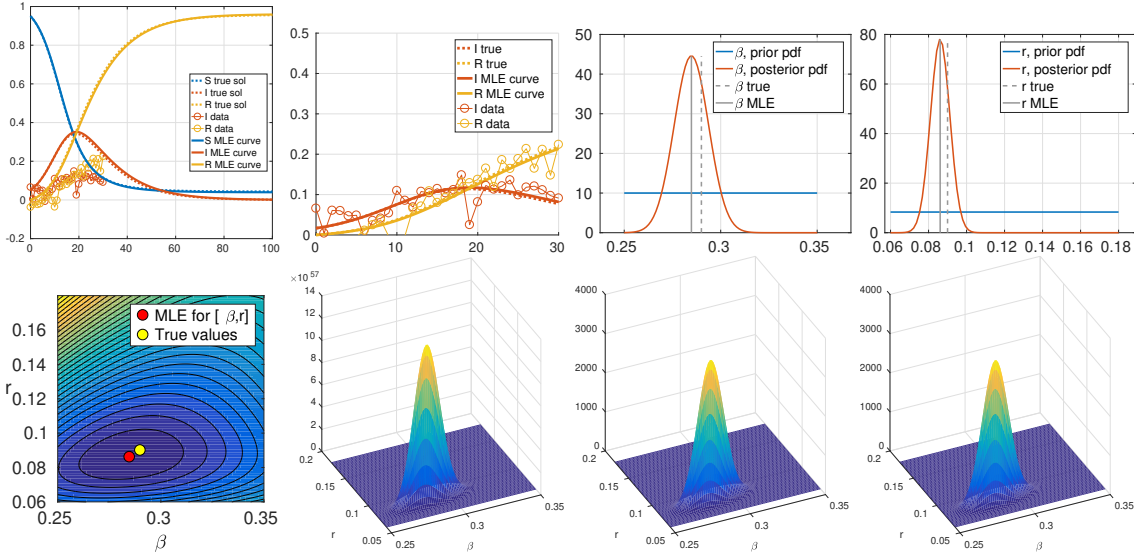


Figure 4: Result of the inverse UQ analysis for SIR. Top row, from left to right: trajectories corresponding to $\vartheta = \vartheta_{true}$ and $\vartheta = \vartheta_{MLE}$, and synthetic data obtained dividing the trajectories for $\vartheta = \vartheta_{true}$ by the under-reporting factor K and adding the Gaussian noise; zoom on the data, and trajectories for $\vartheta = \vartheta_{true}$ and $\vartheta = \vartheta_{MLE}$ rescaled by K ; prior and posterior pdfs for β and r , as well as the true and MLE values of the parameters. Bottom row, from left to right: isolines of NLL ; surface-plot of the full likelihood; surface-plot of the likelihood after the Fisher approximation, cf. equation (15); surface-plot of the likelihood after having further dropped the second derivatives of I, R in the definition of H , cf. equation (16).

centered finite differences) are respectively

$$\Sigma_{G,Jac} = 10^{-4} \times \begin{bmatrix} 0.7995 & 0.1064 \\ 0.1064 & 0.2609 \end{bmatrix}, \quad \Sigma_{G,Hessian} = 10^{-4} \times \begin{bmatrix} 0.8073 & 0.1206 \\ 0.1206 & 0.2616 \end{bmatrix}.$$

Figure 4 provides more details on the results. The top row shows on the left the true trajectories from which the data were generated (dotted thick lines), the noisy data (circles with thin line) and the trajectory obtained by fixing the parameters as $\vartheta = \vartheta_{MLE}$ (solid line). The next panel provides a zoom on the data. We have also rescaled both the true trajectory and the MLE trajectory by K , to emphasize the match with the data. The match between true and MLE trajectories is very good, although not perfect (the distance between the trajectories would further reduce for smaller standard deviations σ of the noise). The last two panels of the row compare the prior and (Gaussian approximation of) the posterior pdfs of the parameters. It can be seen that the posterior are centered close to the true value, and the Gaussian pdf is quite concentrated in comparison to the prior interval. The bottom row provides details about the minimization procedure. More specifically, the leftmost panel shows the contour of the NLL function, the true values of the parameters (yellow dot) and the MLE (red dot). The presence of noise prevents a perfect match between the true values and the MLE, but the match is nonetheless good and the isolines are nicely rounded, which suggest a unique, narrow (hence trustworthy) minimum. The next panel shows the corresponding likelihood function, taken by exponentiating the NLL (remember that since we have assumed uniform prior, the posterior is proportional to the likelihood), which shows a clear Gaussian profile. The next panel shows the Gaussian approximation where the covariance approximation has been computed by inverting the true Hessian of the NLL , while the approximation obtained by using the Jacobian matrix only is shown in the right-most panel, cf. equation (15) and (16). The three latter plots match well (other than by the rescaling factor), indicating that the approximation steps are not introducing significant errors. The MLE has been computed with the `fminsearch` algorithm in Matlab, which implements the derivative-free Nelder–Mead (simplex) algorithm.

Finally, upon calibrating the pdfs of the parameters to the data, the forward UQ analysis can be performed, using the sampling methods discussed in Section 3 to compute the mean of the quantities of interest (S, I, R compartments, location and intensity of the peak), and their pdfs. Results are shown in Figure 5, where we compare the results obtained with prior and posterior pdfs, to appreciate the improvement in the quality of the predictions if data are provided. The top row compares the true and the expected SIR trajectories after the forward UQ analysis based on prior and posterior pdfs. The thick solid lines are the expected values (we

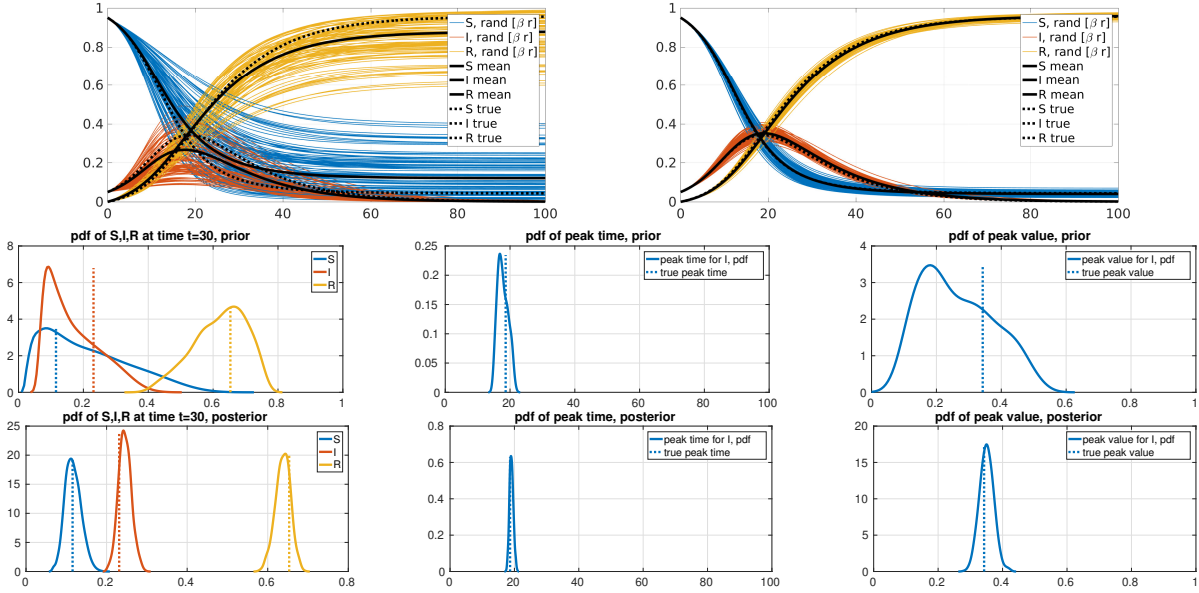


Figure 5: Forward UQ analysis for SIR based on either the prior or the posterior pdfs of the parameters.

compute these with sparse grids sampling), the thick dotted lines are the true trajectories and the colored lines are Monte Carlo trajectories based on the prior/posterior distribution. The results clearly show that the posterior forward UQ is more centered around the true trajectories, and the uncertainty in the prediction is smaller. Similar conclusions can be obtained by looking at the pdfs of the quantities of interest computed based on either the prior or the posterior pdfs, where we have marked with vertical dotted lines the true values (mid-row: prior-based, bottom row: posterior based).

Example 4 (Incorporating prior information on parameters in the inverse UQ). Suppose that we provide prior information about the parameters ϑ_{true} different from the uniform distribution. For instance, we could for simplicity assume that each ϑ_i is a Gaussian random variable with mean $\bar{\vartheta}_i$ and standard deviation s_i , and these variables are all independent (in this way we are allowing ϑ_i to assume negative values with a non-zero probability; we will fix this issue later in the example). Then, the posterior distribution (9) becomes

$$\rho_{post}(\vartheta|\mathcal{D}) \propto \mathcal{L}(\vartheta)\rho_{prior}(\vartheta) = \prod_{i=1}^{N_{meas}} \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2\sigma^2}(\frac{1}{K}I(\vartheta,t_i) - \hat{I}_i)^2} \prod_{i=1}^{N_{meas}} \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2\sigma^2}(\frac{1}{K}R(\vartheta,t_i) - \hat{R}_i)^2} \prod_{i=1}^{N_{\vartheta}} \frac{1}{\sqrt{2\pi s_i^2}} e^{-\frac{1}{2s_i^2}(\vartheta_i - \bar{\vartheta}_i)^2} \quad (17)$$

and the Fisher approximation should now be centered at the maximum of the posterior pdf (we referred to this value as MAP)

$$\boldsymbol{\mu}_{MAP} = \arg \min_{\vartheta} [-2 \log(\mathcal{L}_{\mathcal{D}}(\vartheta)\rho_{prior})],$$

which in practice recovers a form of Tikhonov regularization of the least-squares problem, with penalization parameters $\frac{\sigma^2}{s_i^2}$:

$$\boldsymbol{\mu}_{MAP} = \arg \min_{\vartheta} \left[\frac{1}{\sigma^2} \sum_{i=1}^{N_{meas}} \left(\frac{1}{K}I(\vartheta,t_i) - \hat{I}_i \right)^2 + \frac{1}{\sigma^2} \sum_{i=1}^{N_{meas}} \left(\frac{1}{K}R(\vartheta,t_i) - \hat{R}_i \right)^2 + \sum_{i=1}^{N_{\vartheta}} \frac{1}{s_i^2} (\vartheta_i - \bar{\vartheta}_i)^2 \right].$$

The computation of the sample variance estimator σ_{MLE}^2 and of the Hessian H should be of course updated accordingly. Coming back to the non-positivity issue of the Gaussian prior pdf, a possible workaround is e.g. to assume a log-normal prior for the parameters and then work with the log of the parameters in the model. The expression derived in equation (17) for the posterior pdf would still be valid, the change being “hidden” in the mappings $\vartheta \rightarrow R(\vartheta,t)$, $\vartheta \rightarrow I(\vartheta,t)$.

Example 5 (Inverse UQ when different data types have different noise levels). *Let us consider the data model (5) and assume that the noises $\epsilon_{I,i}$ and $\epsilon_{R,i}$ are independent Gaussian random variables distributed according to $\mathcal{N}(0, \sigma_I^2)$ and $\mathcal{N}(0, \sigma_R^2)$, respectively. The posterior distribution of the parameters can be computed with a slight modification of the procedure explained above to account for the different variances. We detail the new procedure here below, following closely [39, 42].*

We begin by assuming that the ratio λ between the variances of the two sets of data $\lambda := \frac{\sigma_I^2}{\sigma_R^2}$ is known. In this case minimizing the NLL (10) is equivalent to minimizing the following quantity

$$T(\boldsymbol{\vartheta}) = \sum_{i=1}^{N_{\text{meas}}} \left(\frac{1}{K} I(\boldsymbol{\vartheta}, t_i) - \hat{I}_i \right)^2 + \lambda \sum_{i=1}^{N_{\text{meas}}} \left(\frac{1}{K} R(\boldsymbol{\vartheta}, t_i) - \hat{R}_i \right)^2,$$

which is obtained combining (8), opportunely modified to incorporate the different standard deviations, and (10). The parameter λ weighs the sum of squared residuals of I and R , highlighting a different level of trust in the first or second set of data. If λ is small, the residuals of I condition more the quantity T , whereas the data for R are more influential if λ is large. Finally, if $\lambda = 1$ the data sets are equally weighted and we recover the least-squares formula (11). Hence, minimizing the quantity T can be seen as a weighted least-squares criterion.

As λ is however unknown in general, we vary iteratively λ within an appropriate range of values, and minimize T for each value of λ to find the corresponding $\boldsymbol{\vartheta}_{MLE}$. Among the considered λ , we then select the one that realizes the minimum of NLL. The corresponding value of T is also needed, and we denote it by T_{min} . Indeed, the values of the empirical variances of the two data sets are then recovered in the following way

$$\sigma_I = \sqrt{T_{min}/2N_{\text{data}}}, \quad \sigma_R = \sqrt{\sigma_I^2/\lambda}.$$

Note that value of λ can be selected also by means of criteria other than the minimum of the NLL. In particular, in [42] the Kayshap Information Criteria has been employed, that penalizes values of λ that result in shallow NLL.

Finally, we illustrate this minimization procedure with the help of an example. We consider the SIR model, and generate 41 equispaced synthetic data in the time interval $[0, 20]$ with $\boldsymbol{\vartheta}_{true} = [0.29, 0.09]$, adding numerical Gaussian noises with $\sigma_I = 0.2$ and $\sigma_R = 0.05$, where $\lambda_{true} = 16$. We consider integer values of λ in the range $[1, 90]$, and for each of them we minimize T . In Figure 6 on the left we plot the minimum value of the NLL for each considered λ . We observe that the value of λ resulting in the minimum NLL can be approximately correctly identified ($\lambda_{min} \approx 21$), resulting in the following estimates of the parameters

$$\boldsymbol{\vartheta}_{MLE} = [0.2917, 0.0950], \quad \sigma_I = 0.2330 \quad \sigma_R = 0.0508.$$

The figure in the middle displays the posterior distribution of the parameters corresponding to λ_{min} and to the two extremes of the considered range of λ . The figures on the right show 200 Monte Carlo trajectories based on the posterior distribution of the parameters. The value $\lambda = \lambda_{min}$ (left-most panel) gives a narrow bundle of trajectories matching the true ones. The other bundles ($\lambda = 1, 90$ i.e., the smallest and largest values of λ tested) are more spread and in the case of $\lambda = 90$ they are also far from being centered around the true trajectories.

Bibliography and further reading.

- MCMC methods for deriving the posterior pdfs of parameters of SIR-like models for COVID-19 have been used in e.g. [18, 17].
- Several variations of Gaussian approximation of the posterior pdfs can be conceived, depending on the choice of the mean (e.g., MLE, MAP, expected value of the posterior) and of the covariance matrix, see e.g. [43, Result 8, p.224].
- For further approaches to the Bayesian inversion problem, see Approximate Bayesian Computation (ABC, see e.g. [44]) and the Integrated Nested Laplace Approximation (INLA, [45]). ABC is used in cases when it is difficult to evaluate (or even define) the likelihood function. INLA is instead useful when certain conditional posteriors can be reasonably approximated by Gaussian distributions.

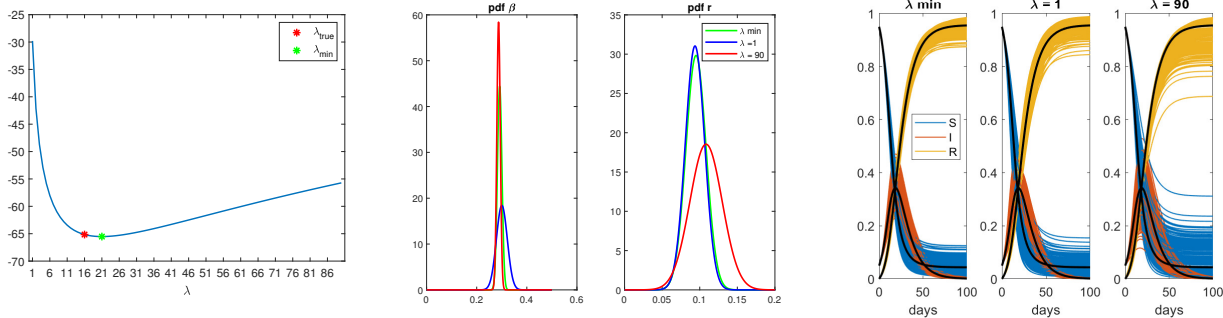


Figure 6: Left: values of NLL across the range of tested values of λ for the test in which the data for I and R have different standard deviations. The smallest NLL is reached for λ_{min} close to the exact one λ_{true} ; middle: posterior distribution of β and r for λ_{min} , $\lambda = 1$, and $\lambda = 90$; right: Monte Carlo trajectories based on the posterior distribution of the parameters, for λ_{min} , $\lambda = 1$, and $\lambda = 90$, respectively. The black lines are the true trajectories.

- In this section we have assumed a simple error model with additive Gaussian errors. See e.g. [6, 5] for more general error models. Furthermore, one can model under-reporting and additional features of the measurements, such as the fact that these kind of data are intrinsically integer numbers, by choosing more suitable likelihood functions than in equation (8), such as negative binomial or quasi-Poisson likelihoods, see e.g. [46, 47, 19].
- In this section we have assumed that ρ_{prior} are either uniform or Gaussian random variables; other choices with good properties are possible, e.g. the Jeffreys non-informative priors [48].
- Using sparse grids for UQ with exact posterior pdf in equation (9) is not straightforward, because the functional shape of ρ_{post} does not fall into classical families of pdfs (e.g. uniform, Gaussian, gamma) for which these methods are developed. A possible remedy would be to compute ad-hoc polynomials, in the spirit of [49, 50]. Another possibility is to keep sampling according to the prior pdf, see e.g. [51], but this is possibly suboptimal if the prior and the posterior are significantly different (e.g., a uniform prior and a very peaked posterior).
- Conversely, sparse grids sampling for Gaussian random variables has been discussed multiple times in literature, see e.g. [52]; therefore, sparse grids computation are easy to use upon having performed the Fisher approximation.
- Replacing the evaluation of the full-model with a response surface either in the MCMC sampling of posterior pdf or in the minimization of NLL is possible (this operation is routine in computationally-heavy inverse problems, see e.g. [53]), but in this context this operation does not dramatically speed-up the computational time because SIR-like models are rather cheap to evaluate.
- We could extend the error model for ϵ_I, ϵ_R to account for the various sources of error (model error, response surface error if used, numerical discretization errors), see e.g. [54].
- Minimizing the NLL function is an example of non-linear least-squares optimization problems, for which ad-hoc algorithms exist, like Gauss-Newton, Levenberg-Marquardt, VarPro, Trust-region reflective, see e.g. [41, 55].
- Inverse UQ / data fitting methodologies can also be applied to stochastic models, where the dynamics of the system are influenced by a “process noise”, and measurements are further corrupted by “observation noise” (i.e., the errors $\epsilon_{I,i}, \epsilon_{R,i}$ in equation (5)). The methodology for data fitting described above of course needs to be adjusted accordingly, see e.g. [56, 57, 58].

6. Summary: an ideal UQ workflow, from data to prediction

Summarizing the discussion so far, the *ideal UQ workflow* for predictions under uncertainty would consist of the steps reported in Algorithm 1. This ideal workflow however is missing one step, i.e., the identifiability

analysis, which is a crucial preliminary analysis to perform. We discuss it in details in the next sections. The adjusted ideal UQ workflow will then be presented in the final Section 9 (Discussion and conclusion), see Algorithm 2.

Algorithm 1: Ideal UQ workflow

- 1 Choose a model and the prior distributions for its parameter (literature, expert opinion);
 - 2 Compute Sobol indices to assess which parameters are more influential and can be inferred from data.
Fix the remaining parameters to some reasonable value;
 - 3 Perform the inverse UQ analysis to adjust the prior distribution to the data evidence;
 - 4 Perform the forward UQ analysis based on the posterior distribution to obtain statistical information about the quantities of interest of the model (e.g. expected value, variance, full pdf of the outputs);
-

7. Structural identifiability

The fundamental assumption underlying the inversion approach proposed in Section 5 is that there exists a certain set of parameters and hyper-parameters $[\boldsymbol{\vartheta}_{true}, \mathbf{h}_{true}]$ that generated the observed data from the system (2), as expressed in Equation (5). We then embrace the fact that the data are noisy, and that this noise might prevent us from correctly determining the values $[\boldsymbol{\vartheta}_{true}, \mathbf{h}_{true}]$; we therefore give up on giving a “one-shot” estimate of $[\boldsymbol{\vartheta}_{true}, \mathbf{h}_{true}]$, and rather content ourselves with computing a posterior pdf, which quantifies our degree of belief on each possible value of $[\boldsymbol{\vartheta}_{true}, \mathbf{h}_{true}]$. The Fisher approximation then further assumes that the NLL has a unique, well-shaped minimum, which means that the posterior pdf of the parameters is sufficiently well-approximated by a Gaussian pdf centered at $[\boldsymbol{\vartheta}_{MLE}, \mathbf{h}_{MLE}]$, whose variance gets smaller as we acquire more data. If instead we believe that there exists a certain set $[\boldsymbol{\vartheta}_{true}, \mathbf{h}_{true}]$ but for some reason we think that our data do not support the assumption that the posterior is Gaussian (for instance, because we have only limited data), we could consider the “full” posterior given by equation (9) instead (using e.g. MCMC as computational tool), and ideally three scenarios might then occur:

1. the posterior pdf is actually close to Gaussian;
2. the posterior pdf is unimodal but it departs from Gaussian in that it might show “heavy tails” and/or some degree of skewness. This would indicate that we might be introducing a bias that leads to over/underestimates;
3. the posterior pdf is multi-modal. This would mean that the inversion procedure is suggesting a few “likely” combinations of parameters $\boldsymbol{\vartheta}$, each corresponding to one peak of the posterior pdf: in this case the heights of the peaks represent our belief on the plausibility that such $\boldsymbol{\vartheta}$ is the “true one”.

In any case, the crucial point that one has to address is: can we guarantee that *there is* a unique set $[\boldsymbol{\vartheta}_{true}, \mathbf{h}_{true}]$ that generates the observed outputs? Or, equivalently, is the inverse problem well-posed? If not, the Fisher approximation is bound to fail (for instance, item 3 in the list above) and the MCMC approach also needs to be handled with care. In the field of mathematical epidemiology (and more generally of dynamical systems/systems control), this question falls into the study of the so-called *system identifiability*, which can be divided in two consecutive steps:

Structural identifiability: studying from a theoretical point of view the well-posedness of the identifiability (inverse) problem, assuming that perfect information is available, i.e., that infinitely many noise-free observations of the outputs are available. It is an intrinsic property of the system and is the topic of this section. This topic is well-studied in the epidemiological literature: a list of references is available at the end of the section.

Practical identifiability: addressing the identifiability of the system given limited and noisy observations of the outputs. It depends not only on the properties of the system but also on the quality of the data. In other words, the structural identifiability is a necessary but not sufficient condition for the practical identifiability of the system. The practical identifiability is the topic of the next Section 8.

Mathematically, a system is structurally identifiable if the model map, i.e. the function $[\boldsymbol{\vartheta}, \mathbf{h}] \mapsto Y(\boldsymbol{\vartheta}, \mathbf{h})$ mapping each realization of $\boldsymbol{\vartheta}, \mathbf{h}$ to the corresponding values of the outputs / quantities of interest is injective. Of course, numerical estimates of parameters obtained by UQ techniques for structurally non-identifiable systems are not reliable and might lead to very wrong predictions. We will show some results on this in Example 8. In the following we discuss two approaches to structural identifiability, one based on differential algebra and a mapping approach. Other methods are available, see the bibliography at the end of the section.

7.1. Differential algebra

The differential algebra approach to the structural identifiability problem is based on deriving a set of differential equations for the model outputs Y of the form $\mathcal{P}(Y, \dot{Y}, \ddot{Y}, \dots, \boldsymbol{\vartheta}, \mathbf{h}) = 0$, where \mathcal{P} is a monic differential polynomial including only Y , their derivatives and the model parameters/hyperparameters; see e.g. Example 6. These equations are known as input-output equations and are an implicit form of the model map, as they generate the same output as the original model. The coefficients of the input-output equations give indication on the identifiability of the system: if the map from $[\boldsymbol{\vartheta}, \mathbf{h}]$ to the coefficients of the input-output equation is injective, the system is structurally identifiable; conversely, if there are multiple values of $[\boldsymbol{\vartheta}, \mathbf{h}]$ that generate the same input-output equations, the system is structurally non-identifiable.

One approach to obtain such equations is by ad-hoc substitution and differentiation to eliminate the unwanted quantities, starting from the original system (2) – of course, one must be careful not to remove/introduce additional solutions e.g. by canceling/multiplying every term by Y . For systems with many compartments, the manual ad-hoc substitution method might be impractical, and a more algorithmic approach is needed: one possibility consists in generating the input-output equations as part of the so-called characteristic set of the algebraic ideal generated by the polynomials defining the model, see [59]. Finally, note that the calculations required to derive the input-output equations can be done also using symbolic calculus software, e.g. Mathematica and Maple.

We refer to [60] for the theoretical background underlying this procedure: in particular, in that work it is shown that the identifiability result does not depend on the particular method employed to derive the input-output equations, as long as a certain property, called *mutual reduction*, holds true – see Example 6. This property is always true in the case of a single output quantity, whereas it needs to be enforced in the case of multiple output quantities. The fact that the identifiability result does not depend on the specific form of the input-output equation stems from the fact that by definition all the input-output equations generate the same output trajectory as the original model. From this, it follows that all forms of such equations contain the same identifiability information of the original system.

Another (quite technical) preliminary condition for the differential algebra approach to be valid is the so-called *solvability condition* [61], which can be safely be assumed to hold for a wide class of practical problems, and we do this in this work as well. We refer again to [60] for a discussion on why this condition can be “safely assumed”, and to [62] for an ad-hoc example where this condition is not valid and the differential algebra approach fails.

Example 6 (Structural identifiability of a SIR model by differential algebra). *In this example, we focus on the case of the SIR model (1) with output $Y = \frac{1}{K}I$. Our argument here is similar to those in [6, 9]; as already mentioned, having prevalence data is somehow unrealistic, yet it is a scenario already rich enough for our “didactical” purposes. We consider the following system*

$$\begin{cases} \dot{S} = -\frac{\beta}{N_{pop}}IS \\ \dot{I} = \frac{\beta}{N_{pop}}IS - rI \\ Y = \frac{1}{K}I. \end{cases}$$

Note that we have neglected the equation for R in (1) as it does not influence the dynamics of the system. Combining the differential equation for I and Y yields

$$\dot{Y} = \frac{\beta}{N_{pop}}YS - rY,$$

which can be solved for S ; from the latter, an expression for \dot{S} can then be derived. Then, replacing these expressions for S and \dot{S} in the differential equation for S of the SIR model, we obtain the following equation

$$N_{pop}\ddot{Y}Y - N_{pop}\dot{Y}^2 + K\beta\dot{Y}Y^2 + Kr\beta Y^3 = 0.$$

We then divide all the terms by N_{pop} to obtain the following monic polynomial, i.e. a polynomial with the coefficient of the highest order term equal to 1:

$$\ddot{Y}Y - \dot{Y}^2 + K\frac{\beta}{N_{pop}}\dot{Y}Y^2 + Kr\frac{\beta}{N_{pop}}Y^3 = 0.$$

Assuming N_{pop} and K known, the map $[\beta, r] \mapsto [K\frac{\beta}{N_{pop}}, Kr\frac{\beta}{N_{pop}}]$ is injective, i.e., the system

$$\begin{cases} C_1 = K\frac{\beta}{N_{pop}}, \\ C_2 = Kr\frac{\beta}{N_{pop}}, \end{cases}$$

can be solved for β, r . This means that it is possible to uniquely identify β and r , and the model is structurally identifiable. Conversely, if K is unknown, only r and the combination $K\beta$ can be uniquely estimated, i.e., the model is structurally non-identifiable.

If instead we consider the case of having also a second quantity of interest $Z = \frac{1}{K}R$, all the parameters of the SIR model result to be structurally identifiable, as we show in the following. Since we have data of I and R , we drop the differential equation for the compartment S of the SIR model and consider the following equivalent system

$$\begin{cases} \dot{I} = \frac{\beta}{N_{pop}}I(N_{pop} - I - R) - rI \\ \dot{R} = rI \\ Y = \frac{1}{K}I \\ Z = \frac{1}{K}R. \end{cases}$$

The practice of eliminating a compartment is possible as we assume that $N_{pop} = S + I + R$ for all times, which means that the dynamics of the compartment that we eliminate is completely determined by the remaining ones. In particular, if we let $S = N_{pop} - I - R$, we can immediately see that

$$\dot{S} = -\dot{I} - \dot{R} = -\frac{\beta}{N_{pop}}I(N_{pop} - I - R) + rI - rI = -\frac{\beta}{N_{pop}}IS,$$

i.e., we recover the initial equation for S . In this second example we have two outputs, therefore we have to derive two input-output equations. We rewrite the differential equations for I and R in terms of Y and Z and obtain

$$K\dot{Y} - \frac{\beta}{N_{pop}}(N_{pop} - KY - KZ)KY + rKY = 0 \quad \text{and} \quad K\dot{Z} - KrY = 0.$$

To make the polynomials monic we have to introduce a ranking among the variables.⁷ A common choice for the ranking is $Y < Z < \dot{Y} < \dot{Z}$ (any other ranking would lead to the same results as already mentioned, of course after different computations), from which it follows that the leading monomials of the input-output equations are \dot{Y} and \dot{Z} , respectively. We then divide all the terms of the equations above by the coefficient of the corresponding leading term and obtain the following monic input-output equations

$$\dot{Y} - \beta Y + \frac{\beta K}{N_{pop}}Y^2 + \frac{\beta K}{N_{pop}}YZ + rY = 0 \quad \text{and} \quad \dot{Z} - rY = 0.$$

⁷i.e., a total ordering of the variables and of their derivatives. For a formal definition, see [63].




Figure 7: The mapping approach for structural identifiability.

Since we have more than one output, before concluding on the identifiability of the system we have to check that the two equations are mutually reduced; if not, the analysis could lead to spurious results, as pointed out in [60]. The concept of reduction is again based on the chosen ranking of the variables: a polynomial \mathcal{P}_i is reduced with respect to the polynomial \mathcal{P}_j if it does not contain neither the leading monomial of \mathcal{P}_j with equal or greater degree nor its derivatives. In our case it can be easily seen that the input-output equations are mutually reduced, and by looking at their coefficients we conclude that β , r and K can be simultaneously identified. Hence, the SIR model is structurally identifiable from prevalence data of I and R .

We close this example mentioning that [6] discusses structural identifiability of SIR (with no under-reporting factor) in the more realistic scenario where cumulative incidence data rather than prevalence are available. The result obtained is that SIR is structurally identifiable even from such kind of data (but other problems that we will discuss in Example 8 occur, so that one should restrain from attempting parameter identification of SIR from cumulative incidence data).

7.2. Mapping approach

This approach is discussed in [64, 65]; we refer the reader interested to the theoretical background to these two references and only sketch the main idea and the “practical recipe” here. Example 7 gives an example of the application of this method to the SIR model. For ease of notation, in this discussion the vector $\mathbf{p} \in \mathbb{R}^p$ collects all the uncertain elements of (2), i.e., not just the coefficients, but also the initial conditions and the hyper-parameters, i.e. $p = N_{coef} + N_{states} + N_{hyp}$. Moreover, we change the notation in (2) to a more compact form and write $f^{\mathbf{p}}(X)$ instead of $f(X, \mathbf{p})$ and similarly for G , and let $N_{states} = n, N_{qoi} = m$. Summarizing the new notation, (2) becomes

$$\begin{cases} \dot{X}(t, \mathbf{p}) = f^{\mathbf{p}}(X(t, \mathbf{p})) \\ X(0, \mathbf{p}) = X_0(\mathbf{p}) \\ Y(t, \mathbf{p}) = G^{\mathbf{p}}(X(t, \mathbf{p})), \end{cases} \quad (18)$$

with $X \in \mathbb{R}^n, Y \in \mathbb{R}^m, f^{\mathbf{p}} : \mathbb{R}^n \rightarrow \mathbb{R}^n, G^{\mathbf{p}} : \mathbb{R}^n \rightarrow \mathbb{R}^m, \forall t \in [0, t_{max}]$, and $X_0 : \mathbb{R}^p \rightarrow \mathbb{R}^n$.

We explain the method with the support of Figure 7. If the system (18) is not identifiable, then there are two different sets of parameters, say \mathbf{p} and $\bar{\mathbf{p}}$, such that the trajectories $X(t, \mathbf{p})$ and $X(t, \bar{\mathbf{p}})$ are different but the corresponding outputs are identical, $Y(t, \mathbf{p}) = Y(t, \bar{\mathbf{p}}) \forall t \in [0, t_{max}]$. If these conditions hold true, then it is possible to rework the equality $Y(t, \mathbf{p}) = Y(t, \bar{\mathbf{p}})$ to explicitly construct a map $\lambda : \mathbb{R}^n \rightarrow \mathbb{R}^n$ that maps the trajectories $X(t, \mathbf{p})$ and $X(t, \bar{\mathbf{p}})$ to one another: $\lambda(X(\bar{\mathbf{p}})) = X(\mathbf{p})$ (we will come back to this point with more details later on).

Since the map λ has been derived by enforcing equality between the outputs $Y(t, \mathbf{p})$ and $Y(t, \bar{\mathbf{p}})$, without taking into account the dynamics of the system, one has to further check that $\lambda(X(\bar{\mathbf{p}}))$ still solves the

dynamical system. Taking time derivatives of both sides of $\lambda(X(\bar{\mathbf{p}})) = X(\mathbf{p})$ and using (18) one gets (removing the dependence on t for ease of notation):

$$\begin{aligned}
\lambda(X(\bar{\mathbf{p}})) &= X(\mathbf{p}) \\
\nabla\lambda|_{X(\bar{\mathbf{p}})}\dot{X}(\bar{\mathbf{p}}) &= \dot{X}(\mathbf{p}) \\
\nabla\lambda|_{X(\bar{\mathbf{p}})}f^{\bar{\mathbf{p}}}(X(\bar{\mathbf{p}})) &= f^{\mathbf{p}}(X(\mathbf{p})) \\
\nabla\lambda|_{X(\bar{\mathbf{p}})}f^{\bar{\mathbf{p}}}(X(\bar{\mathbf{p}})) &= f^{\mathbf{p}}(\lambda(X(\bar{\mathbf{p}}))). \tag{19}
\end{aligned}$$

Therefore, one has to check that the last equation, i.e. equation (19) is valid for the proposed λ . This will result in a set of conditions for the components of \mathbf{p} and $\bar{\mathbf{p}}$: if the resulting conditions are $\mathbf{p} = \bar{\mathbf{p}}$ the system is identifiable; otherwise, there will be non-trivial conditions between some of the components of \mathbf{p} and some of the components of $\bar{\mathbf{p}}$, (see e.g. Example 7), which means that the system is not identifiable.

We now come back to the issue of constructing the map λ . As already mentioned, the idea is to construct λ from the condition $Y(\bar{\mathbf{p}}) = Y(\mathbf{p})$, i.e. from

$$G^{\bar{\mathbf{p}}}(X(\bar{\mathbf{p}})) = G^{\mathbf{p}}(X(\mathbf{p})) = G^{\mathbf{p}}(\lambda(X(\bar{\mathbf{p}}))).$$

In principle, it would be enough to solve for λ in the latter, i.e. $\lambda(X) = (G^{\mathbf{p}})^{-1}[G^{\bar{\mathbf{p}}}(X)]$, but the system might be underdetermined if we have $m < n$ observables. Then, the idea is to complement the observables with additional equations and to create an ‘‘augmented’’ observables vector $H^{\mathbf{p}} : \mathbb{R}^n \rightarrow \mathbb{R}^n$, as follows, by taking into account the directional derivatives of the outputs Y_j along the direction $f^{\mathbf{p}}(X)$:

$$\begin{aligned}
H_1^{\mathbf{p}} &= Y_1 = G_1^{\mathbf{p}}(X(\mathbf{p})) \\
H_2^{\mathbf{p}} &= Y_2 = G_2^{\mathbf{p}}(X(\mathbf{p})) \\
&\dots \\
H_m^{\mathbf{p}} &= Y_m = G_m^{\mathbf{p}}(X(\mathbf{p})) \\
H_{m+1}^{\mathbf{p}} &= \nabla_X Y_1 \cdot f^{\mathbf{p}}(X) \\
H_{m+2}^{\mathbf{p}} &= \nabla_X Y_2 \cdot f^{\mathbf{p}}(X) \\
&\dots \\
H_n^{\mathbf{p}} &= \nabla_X Y_{n-m} \cdot f^{\mathbf{p}}(X).
\end{aligned}$$

The next step is to verify that $H^{\mathbf{p}}$ is a bijective map (this condition is named Observability Rank Condition, ORC), i.e. that the Jacobian of $H^{\mathbf{p}}$ with respect to X is non-singular in the domain of definition of the trajectories X for every admissible \mathbf{p} . Finally, the map λ can be computed by solving for λ the system of equations:

$$H^{\bar{\mathbf{p}}}(X(\bar{\mathbf{p}}), \bar{\mathbf{p}}) = H^{\mathbf{p}}(\lambda(X(\bar{\mathbf{p}})), \mathbf{p}),$$

leading to (cf. Figure 7)

$$\lambda(X) = (H^{\mathbf{p}})^{-1}H^{\bar{\mathbf{p}}}(X),$$

where the well-posedness of $(H^{\mathbf{p}})^{-1}$ is guaranteed by the ORC. We close this discussion with a couple of remarks on the construction of $H^{\mathbf{p}}$:

- In case the number of observables $m < n/2$, adding the directional derivatives $\nabla_X Y_j \cdot f^{\mathbf{p}}(X)$, $j = 1, \dots, m$ will not be enough to reach n observables. In this case, one should also add derivatives of higher order.
- In the opposite case, in which $2m > n$, one can choose among multiple directional derivatives $\nabla_X Y_j \cdot f^{\mathbf{p}}(X)$: after having dropped the choices that make the mapping $H^{\mathbf{p}}$ singular in the domain of definition of the trajectories X , one should check all remaining combinations (the fact that one choice results in identifiability in this case does not rule out the possibility that another choice might result in non-identifiability). If $m > n$, then there are more observables than states, and the same principle applies.

Example 7 (Structural identifiability of a SIR model by mapping approach). *Let us consider the SIR model with output $y = \frac{1}{K}I$ as in Example 6, and rewrite it replacing S, I as $X = [x_1, x_2]$:*

$$\begin{cases} \dot{x}_1 = -\frac{\beta}{N_{pop}}x_1x_2 \\ \dot{x}_2 = \frac{\beta}{N_{pop}}x_1x_2 - rx_2 \\ y = \frac{1}{K}x_2. \end{cases} \Rightarrow f^{\mathbf{p}}(x_1, x_2) = \begin{pmatrix} -\frac{\beta}{N_{pop}}x_1x_2 \\ \frac{\beta}{N_{pop}}x_1x_2 - rx_2 \end{pmatrix}.$$

Our goal is to recover the structural identifiability results already obtained in Example 6 with the differential algebra approach. The first step is to build the two maps $H^{\mathbf{p}}$ and λ , where $\mathbf{p} = [\beta, r, K, N_{pop}]$. As for $H^{\mathbf{p}} = (H_1^{\mathbf{p}}(x_1, x_2), H_2^{\mathbf{p}}(x_1, x_2))^T$, we need to augment the observable y with a directional derivative:

$$\begin{cases} H_1^{\mathbf{p}} = y = \frac{1}{K}x_2, \\ H_2^{\mathbf{p}} = \nabla_X y \cdot f^{\mathbf{p}}(X) = \frac{1}{K} \frac{\beta}{N_{pop}}x_1x_2 - \frac{1}{K}rx_2. \end{cases}$$

The Jacobian of this mapping has non-zero determinant whenever $x_2 \neq 0$, which is a value never attained by the trajectories X unless the initial condition is $x_2(0) = 0$ (in which case the trajectory is the uninteresting case $X = [0, 0]^T$). Therefore, we can apply the methodology. The mapping $\lambda = (\lambda_1(x_1, x_2), \lambda_2(x_1, x_2))^T$ is obtained by solving the equation $H^{\mathbf{p}}(\lambda(X)) = H^{\bar{\mathbf{p}}}(X)$, i.e.

$$\begin{cases} \frac{1}{K}\lambda_2 = \frac{1}{K}x_2, \\ \frac{1}{K} \frac{\beta}{N_{pop}}\lambda_1\lambda_2 - \frac{1}{K}r\lambda_2 = \frac{1}{K} \frac{\bar{\beta}}{N_{pop}}x_1x_2 - \frac{1}{K}\bar{r}x_2, \end{cases} \Rightarrow \begin{cases} \lambda_1 = \frac{N_{pop}}{\beta} \left(\frac{\bar{\beta}}{N_{pop}}x_1 - \bar{r} + r \right), \\ \lambda_2 = \frac{K}{K}x_2, \end{cases}$$

whose Jacobian is

$$\nabla_x \lambda = \begin{bmatrix} \frac{N_{pop}\bar{\beta}}{N_{pop}\beta} & 0 \\ 0 & \frac{K}{K} \end{bmatrix}.$$

Then, enforcing condition (19) results in the following equations:

$$\begin{cases} -\frac{N_{pop}\bar{\beta}}{N_{pop}\beta} \frac{\bar{\beta}}{N_{pop}}x_1(\bar{\mathbf{p}})x_2(\bar{\mathbf{p}}) = -\frac{\beta}{N_{pop}} \frac{N_{pop}}{\beta} \left(\frac{\bar{\beta}}{N_{pop}}x_1(\bar{\mathbf{p}}) - \bar{r} + r \right) \frac{K}{K}x_2(\bar{\mathbf{p}}), \\ \frac{K}{K} \left(\frac{\beta}{N_{pop}}x_1(\bar{\mathbf{p}})x_2(\bar{\mathbf{p}}) - \bar{r}x_2(\bar{\mathbf{p}}) \right) = \frac{\beta}{N_{pop}} \frac{N_{pop}}{\beta} \left(\frac{\bar{\beta}}{N_{pop}}x_1(\bar{\mathbf{p}}) - \bar{r} + r \right) \frac{K}{K}x_2(\bar{\mathbf{p}}) - r \frac{K}{K}x_2(\bar{\mathbf{p}}). \end{cases}$$

The second equation is identically verified. Conversely, the first one holds true for every x_1, x_2 if

$$r = \bar{r}, \quad \frac{K\beta}{N_{pop}} = \frac{K\bar{\beta}}{N_{pop}},$$

i.e. we obtain the same non-trivial condition previously obtained in Example 6: at most one parameter out of K, β, N_{pop} can be identified. If both observations of I and R are available instead, then we have two observables

$$\begin{cases} y_1 = \frac{1}{K}x_2, \\ y_2 = \frac{1}{K}(N_{pop} - x_1 - x_2), \end{cases}$$

and the analysis must be repeated. In particular, the map $H^{\mathbf{p}}$ can now be constructed without using directional derivatives, as

$$\begin{cases} H_1^{\mathbf{p}} = \frac{1}{K}x_2, \\ H_2^{\mathbf{p}} = \frac{1}{K}(N_{pop} - x_1 - x_2). \end{cases}$$

This map is linear, therefore it is bijective for every X and we can carry on with the analysis. The mapping λ is obtained by solving the equations $H^{\mathbf{p}}(\lambda(X)) = H^{\bar{\mathbf{p}}}(X)$, resulting in

$$\begin{cases} \lambda_1 = \frac{K}{\bar{K}}x_1 + N_{pop} \left(1 - \frac{K}{\bar{K}}\right), \\ \lambda_2 = \frac{K}{\bar{K}}x_2, \end{cases}$$

whose Jacobian is $\frac{K}{\bar{K}}$ times the identity matrix. Then, enforcing condition (19) results in the following equations:

$$\begin{cases} -\frac{K}{\bar{K}} \frac{\bar{\beta}}{N_{pop}} x_1(\bar{\mathbf{p}})x_2(\bar{\mathbf{p}}) = -\frac{\beta}{N_{pop}} \left[\frac{K}{\bar{K}}x_1(\mathbf{p}) + N_{pop} \left(1 - \frac{K}{\bar{K}}\right) \right] \frac{K}{\bar{K}}x_2(\bar{\mathbf{p}}), \\ -\frac{K}{\bar{K}} \left(\frac{\bar{\beta}}{N_{pop}} x_1(\bar{\mathbf{p}})x_2(\bar{\mathbf{p}}) - \bar{r}x_2(\bar{\mathbf{p}}) \right) = \frac{\beta}{N_{pop}} \left[\frac{K}{\bar{K}}x_1(\mathbf{p}) + N_{pop} \left(1 - \frac{K}{\bar{K}}\right) \right] \frac{K}{\bar{K}}x_2(\bar{\mathbf{p}}) - r \frac{K}{\bar{K}}x_2(\mathbf{p}), \end{cases}$$

which result in the conditions $r = \bar{r}$, $K = \bar{K}$, $\frac{\beta}{N_{pop}} = \frac{\bar{\beta}}{N_{pop}}$, i.e., the system is now structurally identifiable if we assume N_{pop} to be known, as already discussed in Example 7.

7.3. Sensitivity analysis

As already mentioned in Section 4, if an output / quantity of interest Y is only weakly influenced by a parameter, such parameter might be non-identifiable from Y : therefore, computing the sensitivity of a quantity to a parameter gives an indication about the structural identifiability of the parameter. As explained in Section 4, sensitivity analysis can be local or global. In the first case, a small gradient of the quantity of interest value indicates that the quantity of interest is not very sensible to small variations in the parameter, which might then be non-identifiable. The global sensitivity analysis can be performed by means of Sobol indices: parameters with low Sobol index do not strongly influence the quantity of interest Y , and might be non-identifiable from observations of Y .

Bibliography and further reading.

- The concept of structural identifiability was first introduced in [66].
- Some theoretical background on the differential algebra approach is provided in [63, 60].
- Many methods we have not directly mentioned here (e.g. Taylor series approach, generating series approach, and methods based on the implicit function theorem) are explained in [63].
- See [6] for the detailed discussion of structural identifiability of a SEIR model by differential algebra. As already mentioned, [6] also shows that the SIR model is structurally identifiable also in the case of cumulative incidence data.
- Further discussion on the solvability condition for the differential algebra approach can be found in [67, 68].
- Sensitivity analysis is discussed in, e.g., [5] and [63], where several methods based on the Jacobian matrix J and on the Fisher matrix H are discussed, see equations (13) and (16).

8. Practical identifiability

Upon having assessed the structural identifiability of a system, it is still not obvious that the system can be identified from limited, noisy data, that possibly cover only a fraction of the time-span of the dynamics (e.g., when one measures the initial part of a trajectory and wants to assess the parameters for long-term forecast). With reference to Figure 7, if the output trajectories $Y(\mathbf{p}), Y(\bar{\mathbf{p}})$ are distinct but close, they might become indistinguishable from one another if we only have at our disposal a noisy cloud of points around them rather than the entire exact trajectories. The study of this setting is called *practical identifiability*

analysis, and is typically performed on synthetic data to see under which conditions the inversion procedure obtains results “close enough” to the true values of the parameters. In the following, we give a short outlook on the main tools; see the bibliography at the end of the section for further readings on each method.

Monte Carlo simulations/bootstrap: generate M sets of synthetic data and for each data set, compute the MLE $\vartheta^{(k)}$, $k = 1, \dots, M$. Then, compute dispersion indices for the M MLE of each parameter, such as their sample variance or their average relative error, which is defined as

$$\text{ARE}(\vartheta_i) = \frac{1}{M} \sum_{k=1}^M \frac{|\vartheta_i^{(k)} - \vartheta_{\text{true},i}|}{\vartheta_{\text{true},i}}.$$

Repeat the procedure for synthetic data with increasing levels of noise and observe the trend of the dispersion indices as the noise increases. If the ARE of the estimates is e.g. higher than the noise level, the parameters are not-identifiable (other criteria in the same spirit might be used as well).

The same procedure should be repeated by considering each time a different number of parameters to be jointly estimated, until all the parameters are included. In this way, it is possible to detect the influence of each parameter on the quality of the joint estimate.

Finally, note that the amount and the time-span of the available data might be relevant for practical identifiability, and the analysis should ideally be repeated while varying these settings. For instance, [6] discusses the bootstraps analysis for the SIR model computing the ARE of parameter estimates using data sets of increasing time-span: the results indicate that the model with unknown β, r is practically identifiable from data on the I compartment only after that the epidemic peak is reached (despite the fact the model is structurally identifiable regardless of the time-span of the data, see Example 6).

Fisher Information Matrix: as already explained in Section 5.3, the Fisher Information Matrix is the Hessian of the NLL at the maximum log-likelihood estimate (or its approximation in Equation (16)). If the Fisher Information Matrix is positive definite with large eigenvalues, the NLL has a narrow minimum and we can conclude *local* practical identifiability of the system, i.e. identifiability in a neighborhood of the maximum likelihood estimate; conversely, small eigenvalues are symptom of practical non-identifiability, since the minimum of the NLL occurs in a shallow region. Numerically, the optimization procedure might even end up in a critical point where the Fisher Information Matrix is not even positive definite, e.g. it could be non-definite or it could have rank smaller than the number of parameters. In that case, the rank of Fisher Information Matrix gives an indication on the maximum number of parameters that can be simultaneously inferred, see e.g. [69].

Correlation matrix C : The inverse of the Fisher Information Matrix H gives an approximation of the covariance matrix of the parameters, Σ_G , see Equation (13). From Σ_G , it is possible to compute the *correlation matrix* of the parameters, rescaling each entry as

$$C_{ij} = \frac{\Sigma_{G,ij}}{\Sigma_{G,ii}\Sigma_{G,jj}}.$$

If two parameters have correlation close to 1, they are linearly dependent, and cannot be estimated separately. Hence, they are practically non-identifiable.

Optimization with multiple restart: As already mentioned in Section 5.3, one should repeat several times the optimization to determine the MLE for different initial guesses and check where the minimization ends. If the algorithm ends always at the same point, there is empirical evidence that the NLL has a unique minimum and the system is *globally* practically identifiable (the minimum might be different from the nominal value of the parameter, i.e. the noise might introduce a bias in the estimate, see e.g. Example 3). If several local minima or a manifold of minima are detected we conclude that the system is not practically identifiable: in particular, in the case of a manifold of minima there are many sets of parameters that fit equally well the data, denoting a possible case of structural non-identifiability. For instance, in Example 6 when only data of I are considered, we expect that the region $K\beta = \text{const}$ will be a manifold of minimum points for the NLL.

Profile log-likelihood for each parameter: The previous discussion, pointing to the possibility that the NLL might have a manifold of minima, allows us to introduce the last tool to assess practical identifiability, i.e., the profile log-likelihood; this tool is actually “in between”, and could also be used for structural identifiability, as it will be made clearer later on.

Whenever $\boldsymbol{\vartheta}$ is practically identifiable, the NLL should have a global minimum at $\boldsymbol{\vartheta} = \boldsymbol{\vartheta}_{true}$. Therefore, a visual inspection of the NLL can immediately tell whether the system is identifiable or not. Of course, this is not a viable solution for problems with more than two parameters. In this case, one can resort to the profile likelihood, which is a mono-dimensional slice of the log-likelihood function in the direction of the considered parameter ϑ_i . The profile likelihood can be obtained by changing the parameter ϑ_i iteratively in a certain range of values around its MLE value $\vartheta_{MLE,i}$, while reoptimizing all other parameters. Thus, the profile log-likelihood is defined as

$$PL_i(\boldsymbol{\vartheta}) = \min_{\{\mathbf{y}: y_i = \vartheta_{MLE,i}\}} NLL(\mathbf{y}).$$

If the profile log-likelihood has a unique minimum, the parameter is practically identifiable, whereas more complex shapes, with shallow regions and multiple minima indicate that the parameter is practically non-identifiable. In particular, a flat profile means that the NLL has a manifold of minima that fit the data equally well and we can again conclude that the system is structurally non-identifiable. For a schematic illustration we refer to Figure 8. Crucially, if the NLL has a manifold of minima, the Fisher approximation of the posterior will be completely wrong, since it is based on the assumption that NLL has a unique minimum; the case of finitely many local minima could instead be fixed by acquiring more data, which should hopefully make the “spurious peaks” become smaller and smaller.

In summary, the analysis of the profile log-likelihood is based on checking the “flatness”/“shallowness” of such function. To make this criterion more quantitative and discriminate between different levels of shallowness one can provide confidence thresholds of the profile log-likelihood CI_{PL} using a χ_1^2 test:

$$CI_{PL}(y_i) = \{\mathbf{y} | \mathbf{PL}_i(\mathbf{y}) \leq \mathbf{PL}(\vartheta_{MLE,i}) + \Delta_\alpha \chi_1^2\},$$

where $\Delta_\alpha \chi_1^2$ is the α quantile of the χ_1^2 distribution. If the likelihood profile of a parameter exceeds the confidence threshold $\Delta_\alpha \chi_1^2$ on both sides of $\vartheta_{MLE,i}$ the parameter is practically identifiable (the threshold is indicated by the black dotted lines in Figure 8). This is an alternative approach to determine a posterior pdf for the MLE, consisting of a uniform pdf instead of the Gaussian Fisher approximation discussed in Section 5.3. Such uniform estimate might be more robust for limited datasets [70].

Finally, observe that if the data are generated in a noise-less way by sampling the trajectories of $\boldsymbol{\vartheta} = \boldsymbol{\vartheta}_{true}$, and they are sufficiently many, we end up in the scenario of structural identifiability. The profile log-likelihood can still be constructed and evaluated also in this case, as the NLL is just the sum of square misfits (without the $\frac{1}{\sigma^2}$ factor), and therefore it can be used as a computational tool to verify structural identifiability, i.e., for the existence of a manifold of minima.

In case the system is structurally or practically non-identifiable, a possible workaround is to learn some of the parameters from independent studies, thus reducing the number of parameters to be simultaneously identified. Sometimes even a hierarchical optimization approach might be effective: in this approach, one does a first round of optimization to obtain the values of the entire set of parameters but retains the values obtained for the identifiable parameters only. Upon fixing these parameters to the values just obtained, the optimization of the remaining parameters can be repeated. An alternative is to reparametrize the model, replacing the original parameters with the combination that can be identified, see e.g. [11]. Finally, one could resort to so-called *marginalization techniques*, see e.g. [71, 72, 73].

Example 8 (Practical non-identifiability of a SIR model with unknown under-reporting factor). *In the scenario of COVID-19, it has been often pointed out that data of infected and dead persons have been under-reported (even significantly), but the exact value of the under-reporting factor K is not known. Some discussion on this aspect is provided in [13, 17]. In this example, we consider a SIR model, generate synthetic data and investigate the practical identifiability of β, r, K when measurements of I and R are considered, by computing the profile likelihood for K ; we already know from Example 6 that in this scenario (somehow*




Figure 8: Three cases of profile likelihood (blue line), with nominal value of the parameter (red star), and confidence threshold $\Delta_\alpha\chi_1^2$ (dashed line). Left: the parameter is practically identifiable; center: the parameter is not structurally identifiable; right: the parameter is practically non-identifiable.




Figure 9: Profile Likelihood for K for the three different scenarios, i.e. data collection ending: before peak of I (left), around the peak of I (center), after the peak of I (right).

unrealistic, yet already rich enough for exposition purposes) the parameters are structurally identifiable. We repeat practical identifiability analysis in three settings, that differ by the time-span covered by the data: until past the peak of I ($T = 20$), up to the peak of I ($T = 30$), and before the peak of I ($T = 40$).

We fix $\vartheta = [0.28, 0.11]$, $\sigma = 0.025$, and the discount factor to $K = 3$. The results are reported in Figure 9 and suggest that for data before peak the profile likelihood of K is shallow (look at the scale on the vertical axis) and even has a minimum at the wrong value $K = 2$, denoting practical non-identifiability of K . For longer collection times instead, the profile likelihood shows an increasingly deep minimum around the correct value $K = 3$. The fact that time might be important in determining whether a system is practically identifiable was already discussed in [5, 8, 6]. Fixing K to the wrong value to perform the forward UQ analysis of course leads to predictions that are far from the true behavior of the system, see Figure 10-top-left. An important disclaimer to do here is that it would be hard to say that the set of parameters obtained for $K = 2$ fits the data worse than those obtained for $K = 3$, see Figure 10-top and bottom-second figure. Even using quantitative criteria to evaluate the goodness of fit of the two fittings, such as

$$\text{Root Mean Square Error (RMSE)} : \frac{1}{N_{meas}} \sqrt{\sum_{m=1}^{N_{meas}} \left(F_m(\vartheta_{MLE}) - \hat{F}_m \right)^2}, \text{ for } F = I, R$$

$$\text{Mean Absolute Error (MAE)} : \frac{1}{N_{meas}} \sum_{m=1}^{N_{meas}} |F_m(\vartheta_{MLE}) - \hat{F}_m|, \text{ for } F = I, R$$

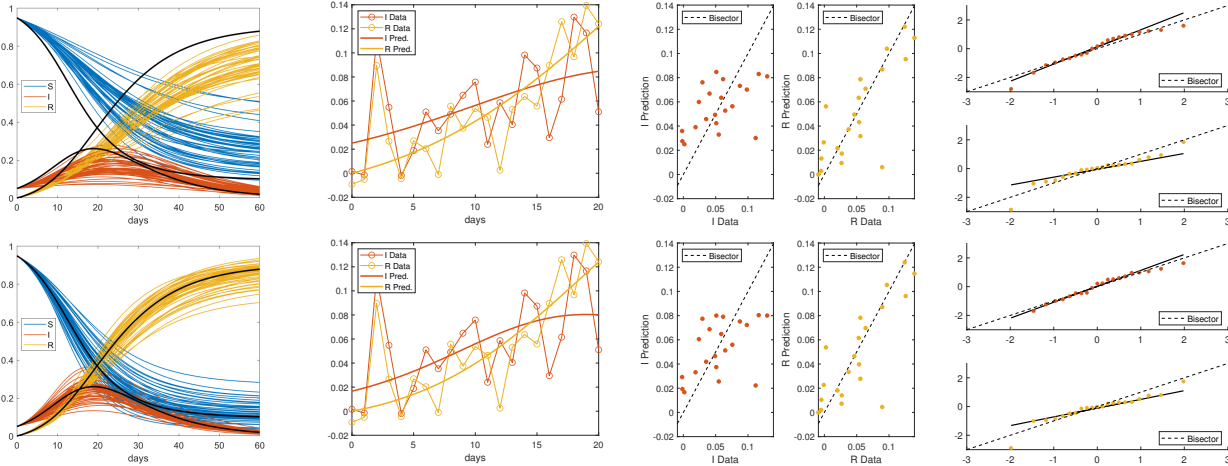


Figure 10: Comparison of quality of predictions and data fitting obtained for $K = 2$ (top) and $K = 3$ (bottom) in the scenario of data available only before the peak ($T = 20$). From left to right: forward UQ based on the posterior pdf; the fitting of the data in the two cases; scatterplot of data vs predictions; qqplot of misfits.

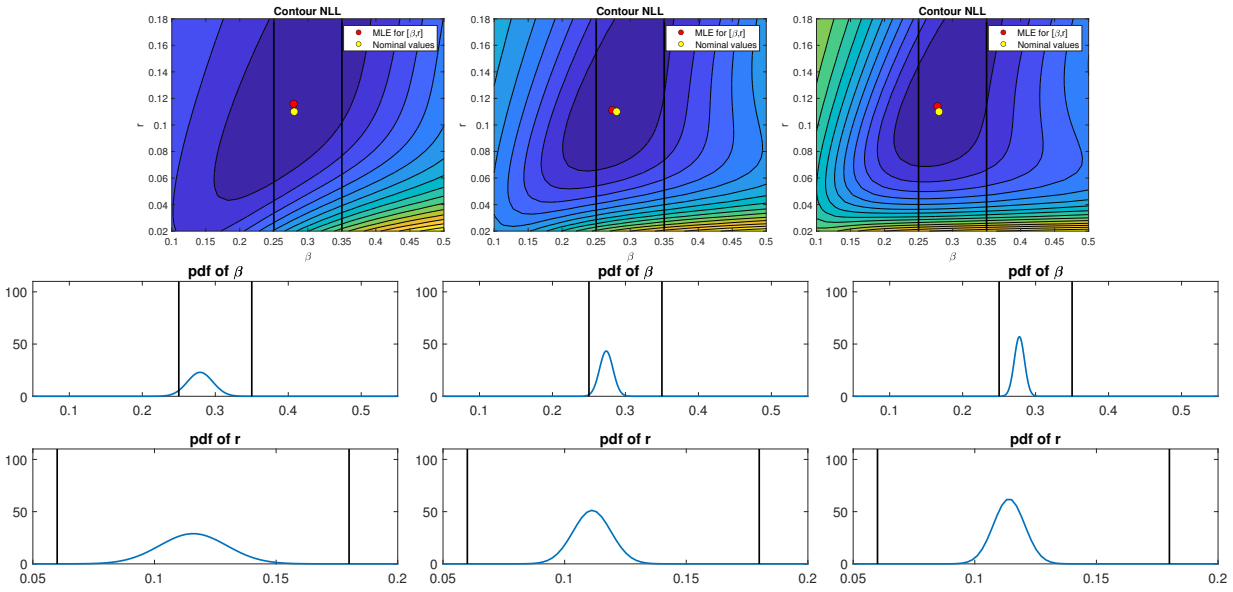


Figure 11: Results of inversion for the three cases $T = 20, 30, 40$ for $K = 3$. Top row: contours of the NLL in the three cases (left to right); the true value of the parameters is the yellow dot, the MLE is the red dot. The NLL contour lines suggest that the NLL has a unique minimum at the MLE, which is always close to the true value of the parameters. Bottom rows: Gaussian approximation of the posterior pdfs of β, r for the three cases $T = 20, 30, 40$ for $K = 3$.

Mean Absolute Percentage Error (MAPE) : $\frac{1}{N_{meas}} \sum_{m=1}^{N_{meas}} (F_m(\boldsymbol{\vartheta}_{MLE}) - \hat{F}_m) / \hat{F}_m$, for $F = I, R$.

would actually tell that the fitting of the case $K = 2$ is slightly better than the case $K = 3$ (numbers not reported for brevity). The scatterplots of predictions vs data are qualitatively identical, and reasonably aligned with the bisector, see Figure 10-top and bottom third and fourth panel. The empirical distribution of the misfits in both cases are qualitatively identical and close to a Gaussian (see the quantile-quantile plots in Figure 10-top and bottom-right panels). In summary, all of these diagnostic tools give little-to-no evidence that $K = 2$ is the wrong choice of the under-reporting parameter. We close this example with some remarks:

- the quality of the fitting of the cases $K = 2, 3$ cannot be assessed by the R^2 coefficient, which is not well-defined for non-linear least-squares problems [74].
- if instead K is known, the SIR system is practically identifiable regardless of time, since the NLL in the $\beta - r$ plan has always a unique minimum close to the true value, in all of the three scenarios, see Figure 11-top. Of course, the time span of data collection still has an impact on the quality of the results. Indeed, if time increases the minimum of the NLL is less and less shallow, which implies that the uncertainty on the parameters is smaller and smaller. This is visible in Figure 11-bottom, where we show the Gaussian approximation of the posterior pdfs of β, r . As the data time-span increases, the posterior pdfs are more and more concentrated and, equivalently, the system is more and more practically identifiable (cf. Fisher Information Matrix criterion for practical identifiability).
- As already mentioned in Example 6, [6] discusses whether the SIR model without under-reporting factor is identifiable from cumulative incidence data, rather than prevalence data; since cumulative incidence data are most typically reported in an outbreak, this question is very relevant for practical purposes. The finding is that while SIR is indeed structurally identifiable from cumulative incidence data, it is not practically identifiable from this kind of data, and advocates for a broader diffusion of prevalence data, from which SIR is practically identifiable, as just discussed in this example.

Example 9 (Structural and practical identifiability of a SEIRD model). In this example, we discuss the identifiability of a slightly more complex model with incubation period (compartment “E”, exposed), where we distinguish between recovered and dead persons (compartments “R” and “D”, respectively). We also assume that at $T = T_{lock}$ the parameter β changes, due to some restriction measure being enforced (lockdown).⁸ We call this model SEIRDz:

$$\begin{cases} \dot{S} = -\frac{\beta(t)}{N_{pop}} IS \\ \dot{E} = \frac{\beta(t)}{N_{pop}} IS - iE \\ \dot{I} = iE - dI - rI \\ \dot{R} = rI \\ \dot{D} = dI \end{cases}$$

where $\beta(t) = \beta_1$ for $t \leq T_{lock}$ and $\beta(t) = \beta_1 - z$ for $t > T_{lock}$.⁹ We consider these ranges for the parameters (that we consider as uniform random variables):

$$\beta_1 \in [0.25, 0.35] \quad r \in [0.06, 0.18] \quad d \in [0.01, 0.02] \quad i \in [0.14, 0.33] \quad z \in [0.1, 0.2].$$

In this setting the rate r is the recovery rate (inverse of the average days of sickness), the rate d is the mortality rate and the rate i is the inverse of the incubation time. With these intervals we are assuming that the average number of days of sickness is roughly between 5 and 16, while the average incubation time is roughly between 3 and 7 days. Moreover, upon introducing the fatality ratio $d/(d+r)$, i.e. the proportion of infected that eventually die, this example considers a very lethal infection, with a fatality ratio roughly

⁸A change in β of 90% was observed on real-data e.g. in [18].

⁹Another option would be to consider $\beta(t) = \beta_1(1-z)$ after T_{lock} , i.e., to use z as the percentage decrease of β_1 instead of the absolute decrease. The treatment of the two models would be identical.




Figure 12: Prior-based forward UQ for the SEIRDz model. Left-most panel: Monte Carlo realizations and expected value of the compartments. Remaining panels: time-evolution of the Sobol indices.

between 5% and 25%. We consider $T_{lock} = 15$ and run the simulation until $T = 100$. We set the initial conditions to $S(0) = 0.95, E(0) = 0.04, I(0) = 0.01, R(0) = 0, D(0) = 0$, and we assume of having at disposal measurements of the compartments I, R, D .

If we assume that β is constant in time, we can show that the system is structurally identifiable, by means of the differential algebra approach. The computations are shown in details in Appendix A. Since we have assumed that we know the time T_{lock} where the change in β happens, we can apply the structural analysis results to both the time intervals $t \leq T_{lock}$ and $t > T_{lock}$ separately, and conclude that the model SEIRDz is structurally identifiable.

Figure 12 shows some results for the preliminary prior-based forward UQ analysis. The left-most panel shows 100 Monte Carlo trajectories, and the expected value of the compartments, computed with a sparse grid (2433 model evaluations¹⁰). It is clearly visible that the trajectories are significantly scattered, and that the asymptotic values of the compartments vary considerably. The remaining panels show the time-evolution of the Sobol indices, from which we can derive some information about the identifiability of the model. We can see that not all parameters impact equally the variability of the solution, and we expect in particular that it will be difficult to recover by the inversion procedure the value of those parameters that have the smallest impact (i, z, β_1). Even more so given that we only measure the compartments I, R, D (while β_1 would be best recovered from the S compartment), and that we will only measure them up to a certain time, and the Sobol indices are not constant in time (for instance, β_1 has a significant impact on R but only at late times, say $T \geq 30$, while we measure essentially early times). Thus, while the system is structurally identifiable, it might be practically non-identifiable. Next, we perform the inversion and check for practical identifiability. We fix the parameters as

$$\beta_{1,true} = 0.28, \quad r_{true} = 0.11, \quad d_{true} = 0.018, \quad i_{true} = 0.18, \quad z_{true} = 0.18, \quad K = 3, \quad \sigma = 0.01;$$

and measure data until $T = 40$ (after peak). We repeat the minimization procedure 20 times with different initial guesses for the parameters, to investigate the presence of local minima of the NLL, and select as MLE the results that led to the smallest NLL. The results of this procedure are shown in Figure 13-top, where we report the initial and final values of the NLL, as well as the initial and final values of the parameters. We can conclude that while the final values of the NLL are all close (yet not identical), the values of the parameters show a significant variability, denoting the fact that the NLL has a multiple minima with similar NLL value. Moreover, most of the parameters are not correctly identified. This further suggests that the model might be practically non-identifiable, at least for the value of σ tested here. Therefore, the trajectories corresponding to the computed values of the parameters are quite far from the true one, see Figure 13-bottom.

To confirm our diagnosis of non-identifiability, we compute the profile likelihood of the problem. Results are shown in Figure 14. The only parameter with a deep, narrow minimum is r (compare the vertical scales) while the other ones are in rather shallow regions. Moreover, the profile likelihoods for i, β_1 and z are very noisy (as expected, due to the fact that their Sobol indices are rather small). This confirms that the only

¹⁰the number of sparse grids points is larger than in Example 1, where we considered a SIR model. This is because now we have to sample a 5-dimensional parameters space, and sparse grids suffer to a certain degree the so-called ‘‘curse of dimensionality’’ [75], i.e. loosely speaking, the number of sampling points grows rapidly (more than linearly) with the number of dimensions. A sampling method that fully suffers from this problem is cartesian sampling, where the number of points grows exponentially with the number of dimensions.

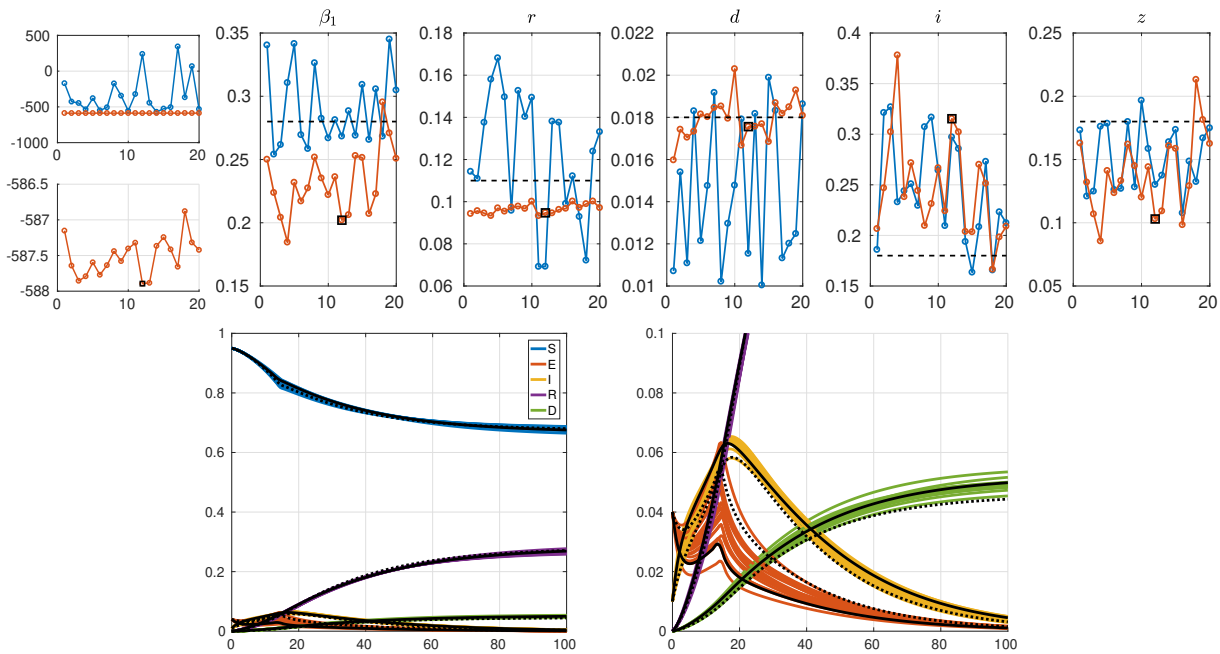


Figure 13: Results for inversion of SEIRDz. Top row, from left to right: the initial and final value of NLL for the 20 trials (blue and red line, respectively) and a zoom on the final values (the value with the smallest NLL is marked with a black square); the initial and final values of the parameters for the 20 trial (blue and red line, respectively), the correct values (black dash lines), and the values of the parameters that yield the smallest NLL (black square marker). Bottom row: the trajectories of the 20 MLEs and a zoom on the E, I, D compartments. The colored trajectories are those obtained by the values of the parameters obtained by the 20 optimization trials. The MLE trajectories are reported in full black line, while the true ones in dashed black lines.

parameter that can be easily identified are r and to a certain extent d (given that the shape of the profile likelihood is not noisy, although shallow), and overall the system is practically non-identifiable.

As mentioned, a possible workaround which might help in this situation is to learn some of the parameters from independent studies, and thus reduce the number of parameters to be simultaneously identified. For instance, medical studies might give us estimates of the incubation time, recovery time and death rate, so that we are left to identify only the contact probabilities β_1 and z . One has to pay attention to the fact that setting the influential parameters to the wrong values can be however detrimental to the procedure. Here, we fix the values of i, r , and d to their exact values and repeat the identification procedure. This is of course an over-optimistic scenario. Another possibility would be to use, for example, r, d obtained from the inversion procedure (whose profile likelihood is “well-shaped”, even though the one for d is quite shallow; we named this procedure as “hierarchical optimization” in the previous discussion). The results obtained by fixing i, r , and d to their exact values are reported in Figure 15, which shows the same information of Figure 13, i.e. initial and final values of NLL and parameters, and trajectories of the model. The presence of local minima in the NLL is greatly reduced, and the identification of β_1 and z is more robust and closer to the true values (this might not always be the case though, depending on the noise level and the quality of the data). As a result, the trajectories corresponding to this new set of parameters are closer to the true ones than the previous results, as shown in the right-most panel.

Bibliography and further reading.

- For a general survey on practical identifiability, we refer the reader to [63, 6, 7].
- Bootstrap approaches in the context of epidemiological models are considered e.g. in [6, 7]. The analysis in [6] considers a criterion based on ARE to conclude on identifiability, whereas in [7] a combination of the size of the confidence interval and mean squared error of the parameters is employed.
- An extensive explanation on the use of the likelihood profiles is given in [70] and further references are [9, 11].




Figure 14: Profile likelihood for the SEIRDz identifiability problem, centered around the MLEs of the parameters (red circle markers). Observe that the red markers do not always coincide with the profile likelihood. This is because the optimization for the remaining parameters might land at different local optima (despite running the optimizer multiple times with random initial points).




Figure 15: Results for inversion of SEIRDz upon blocking the parameters r, d, i . Left panel: the initial and final value of NLL for the 20 trials (blue and red line, respectively) and a zoom on the final values (the value with the smallest NLL is marked with a black square). Mid panel: the initial and final values of β_1, z for the 20 trial (blue and red line, respectively), the correct values (black dash lines), and the values of the parameters that yield the smallest NLL (black square marker). Right panel: the true trajectories (black lines), the MLE when trying to identify all parameters (colored dotted lines) and the MLE when trying to identify β_1, z only (full colored lines).

- [8] provides a quite comprehensive step-by-step guide on fitting compartmental models for epidemiology, with an eye to identifiability.
- In the context of COVID-19, some discussion about identifiability is provided by [76].
- [9] provides an interesting example of the consequences of identifiability in the context of a SEIR-based model for a Dengue outbreak: the system is practically non-identifiable, resulting in two sets of parameters that show an excellent fit of the data but provide dramatically different predictions when used to test a possible non-medical remediation strategy (removal of mosquitos).
- Optimal design of experiments can be used to improve the quality of the data to minimize the impact of practical identifiability issues, see e.g. [77].

9. Discussion and conclusions: a revisited UQ workflow

In these notes we have reviewed some computational tools for prediction under uncertainty and parameter identification for dynamical systems, which we referred to as forward and inverse UQ analyses, respectively. Combined together, these tools provide a powerful framework for reliable predictions of the outputs of a dynamical system, and complement the punctual predictions with confidence estimates with solid ground in probability/statistics theory. However, investigators should always carefully check whether their model is actually identifiable, both structurally and practically. Dynamical systems might indeed be not fully identifiable, and blindly using the prediction tools in this case can be harmful, see e.g. Example 8 or the above-mentioned case study on Dengue reported in [9].

If the system is not identifiable, the Fisher approach to the inversion problem is bound to fail, because it intrinsically assumes identifiability of the system, or in other words because it assumes that the NLL has a unique minimum where the Gaussian approximation of the posterior should be centered, whereas in case of structural non-identifiability the NLL has a manifold of minima. The case of NLL with a finite number of local minima (symptom of practical non-identifiability, see Section 8) could instead hopefully be fixed by acquiring more data of the right kind, that should hopefully rule out the “wrong minima”. Conversely, MCMC methods make no assumptions on the shape of the NLL and therefore might be a partially safer technique, However, MCMC algorithms come with a much larger computational cost and are not entirely safe either, since they typically implement some sort of adaptive sampling, where most of the samples are collected in regions of large likelihood, i.e., they cannot entirely escape the problem of computing the maxima of the likelihood. Therefore, unless they are properly designed and tuned, they could fail to realize that there might be an entire manifold of minima. A compromise solution could be to use the Fisher approximation not plainly as the posterior pdf of the parameters, but only in the context of an importance sampling strategy, where one still generates samples from the Fisher approximation of the posterior but then rescales them suitably to remove any bias [78, 79]. In summary, the ideal UQ workflow sketched in Section 6 (cf. Algorithm 1) can be adjusted as detailed in Algorithm 2.

Of course, the underlying assumption here is that we know what is the exact model that generated the data, and we are dealing with parametric identification only. Discussing how model mis-specification affects identifiability further adds to the complexity of the problem and it is out of the scope of this work.

Our theoretical discussion has been complemented with a number of small examples. None of these consider the initial conditions of the system as unknown, but doing so would not pose any conceptual challenge from a numerical point of view. As already hinted in Section 7, the structural identifiability of parameters connected to the initial conditions can be investigated e.g. by the mapping approach.

An important point that we did not discuss is the issue of *model-selection*. In the case when multiple models are available (quite common in the case of epidemics modeling and, in particular, of COVID-19), is there any way to tell which one has the largest statistical evidence? A large body of work is available on this topic in the statistical literature, where several criteria have been developed to select the “best model”. The underlying principle is that adding more parameters might lead to a better fit of the data, but the more parameters, the larger the chances that the model is overfitted, i.e., that it adjusts to the noise and gets limited predicting power. Thus, one should restrain from blindly adding more parameters.¹¹ Criteria that

¹¹this is a mathematical formulation of the Occam’s razor.

Algorithm 2: Ideal UQ workflow

```
1 Choose a model and the prior distributions for its parameter (literature, expert opinion);
2 Determine whether the system is structurally identifiable (Sobol indices, profile likelihood,
  differential algebra, mapping approach, etc.);
3 if the model is structurally identifiable then
4   | while the model is not practically identifiable (bootstrap, profile likelihood, multiple restart, etc.)
5   |   do
6   |     Acquire more data / get information on some parameters from independent studies / perform a
7   |     hierarchical optimization;
8   |   end
9   |   choose Fisher approximation as inversion method;
10 else
11 | the likelihood has a manifold of minima: choose an appropriate MCMC algorithm as inversion
12 | method;
13 end
14 perform the inverse UQ analysis;
15 perform the forward UQ analysis based on the posterior distribution to obtain statistical information
16 about the quantities of interest of the model (e.g. expected value, variance, full pdf of the outputs);
```

try to identify the optimal model among a pool of possible ones include, for example, the Akaike Information Criterion (AIC), the Bayes Information Criterion (BIC), and the Kayshap Information Criterion (KIC). We refer the interested reader e.g. to [80, 81, 82] for a more thorough discussion, as well as to e.g. [83] for a discussion on a model selection strategy in the case when some models are not identifiable.

Acknowledgments

The authors acknowledge the many fruitful discussions with several colleagues, and in particular the colleagues at CNR-IMATI that participated in the COVID-19 modeling study group.

Declarations of interest

None.

Funding

Lorenzo Tamellini and Chiara Piazzola have been supported by the PRIN 2017 project 201752HKH8 “Numerical Analysis for Full and Reduced Order Methods for the efficient and accurate solution of complex systems governed by Partial Differential Equations (NA-FROM-PDEs)”. Lorenzo Tamellini also acknowledges the support of GNCS-INdAM (Gruppo Nazionale Calcolo Scientifico - Istituto Nazionale di Alta Matematica). This work was supported by the KAUST Office of Sponsored Research (OSR) under Award No. URF/1/2584-01-01 and the Alexander von Humboldt foundation. Raúl Tempone is a member of the KAUST SRI Center for Uncertainty Quantification in Computational Science and Engineering.

Appendix A. Structural identifiability of a SEIRD model by differential algebra

In this section we consider the SEIRD model, which is a simplified version of the SEIRDz model considered in Example 9 with β constant in time. We show by means of the differential algebra technique explained in Section 7 that it is structurally identifiable from prevalence data of I , R and D .

Let us consider the following system

$$\left\{ \begin{array}{l} \dot{S} = -\frac{\beta}{N_{pop}}IS \\ \dot{I} = i(N_{pop} - S - I - R - D) - dI - rI \\ \dot{R} = rI \\ \dot{D} = dI \\ Y = \frac{1}{K}I \\ Z = \frac{1}{K}R \\ W = \frac{1}{K}D, \end{array} \right.$$

where we have removed the equation for the compartment E , as it holds that $N_{pop} = S + E + I + R + D$ (the same argument was used in Example 6, when we removed the equation for S while discussing identifiability of SIR with data of I and R). We then rewrite the differential equations in terms of the observed variables. In particular, we derive the following explicit expression for S from the second equation:

$$S = -\frac{K}{i}\dot{Y} + N_{pop} - K\frac{i+d+r}{i}Y - KW - KZ.$$

From this, we compute \dot{S} , insert both formulas in the first differential equation above, and obtain the first input-output equation:

$$-\frac{K}{i}\ddot{Y} - K\frac{i+d+r}{i}\dot{Y} - K\dot{W} - K\dot{Z} - \frac{\beta K^2}{iN_{pop}}Y\dot{Y} + \beta KY - \frac{\beta K^2}{iN_{pop}}(i+d+r)Y^2 = 0.$$

The other two input-output equations follow from the third and fourth differential equation above and are:

$$K\dot{Z} - rKY = 0 \quad \text{and} \quad K\dot{W} - dKY = 0.$$

The set of these three input-output equations is not mutually reduced with respect to the ranking $Y < Z < W < \dot{Y} < \dot{Z} < \dot{W} < \ddot{Y} < \ddot{Z} < \ddot{W}$ (other ranking would lead to the same results). The third equation is not reduced with respect to the first one, as its leader monomial \dot{W} appears also in the first one. Similarly, the second equation is not reduced with respect to the first one. By doing some further substitutions to eliminate \dot{Z} and \dot{W} in the first equation we finally get a set of mutually reduced equations:

$$\left\{ \begin{array}{l} -\frac{K}{i}\ddot{Y} - K\frac{i+d+r}{i}\dot{Y} - KdY - KrY - \frac{\beta K^2}{iN_{pop}}Y\dot{Y} + \beta KY - \frac{\beta K^2}{iN_{pop}}(i+d+r)Y^2 = 0 \\ K\dot{Z} - rKY = 0 \\ K\dot{W} - dKY = 0. \end{array} \right.$$

The last step is to make the polynomials monic with respect to their leaders, which are \ddot{Y} , \dot{Z} , and \dot{W} , respectively. It then follows that all the coefficients can be uniquely determined; hence, the SEIRD model is structurally identifiable from data of I , R , and D .

References

- [1] A. Alahmadi, S. Belet, A. Black, D. Cromer, J. Flegg, T. House, P. Jayasundara, J. Keith, J. McCaw, M. R., J. Ross, F. Shearer, S. Thein Than Tun, W. J., L. White, J. Whyte, A. Yan, A. Zarebski, Influencing public health policy with data-informed mathematical models of infectious diseases: Recent developments and new challenges, *Epidemics* 32 (2020) 100393.

- [2] J. H. Guillaume, J. D. Jakeman, S. Marsili-Libelli, M. Asher, P. Brunner, B. Croke, M. C. Hill, A. J. Jakeman, K. J. Keesman, S. Razavi, J. D. Stigter, Introductory overview of identifiability analysis: A guide to evaluating whether you have the right type of data for your modeling purpose, *Environmental Modelling & Software* 119 (2019) 418 – 432.
- [3] W. O. Kermack, A. G. McKendrick, A contribution to the mathematical theory of epidemics, *Proceedings of the Royal Society A*. 115 (772) (1927) 700–721.
- [4] M. A. Capistrán, J. A. Christen, J. X. Velasco-Hernández, Towards uncertainty quantification and inference in the stochastic SIR epidemic model, *Mathematical Biosciences* 240 (2) (2012) 250–259.
- [5] A. Capaldi, S. Behrend, B. Berman, J. Smith, J. Wright, A. L. Lloyd, Parameter estimation and uncertainty quantification for an epidemic model, *Mathematical Biosciences and Engineering* 9 (3) (2012) 553–576.
- [6] N. Tuncer, T. T. Le, Structural and practical identifiability analysis of outbreak models, *Mathematical Biosciences* 299 (2018) 1–18.
- [7] K. Roosa, G. Chowell, Assessing parameter identifiability in compartmental dynamic models using a computational approach: application to infectious disease transmission models, *Theoretical Biology and Medical Modelling* 16 (1) (2019).
- [8] G. Chowell, Fitting dynamic models to epidemic outbreaks with quantified uncertainty: A primer for parameter uncertainty, identifiability, and forecasts, *Infectious Disease Modelling* 2 (3) (2017) 379–398.
- [9] Y.-H. Kao, M. C. Eisenberg, Practical unidentifiability of a simple vector-borne disease model: Implications for parameter estimation and intervention assessment, *Epidemics* 25 (2018) 89–100.
- [10] M. G. Roberts, Epidemic models with uncertainty in the reproduction number, *Journal of Mathematical Biology* 66 (7) (2013) 1463–1474.
- [11] C. Tönsing, J. Timmer, C. Kreutz, Profile likelihood-based analyses of infectious disease models, *Statistical Methods in Medical Research* 27 (7) (2018) 1979–1998.
- [12] G. Giordano, F. Blanchini, R. Bruno, P. Colaneri, A. Di Filippo, A. Di Matteo, M. Colaneri, Modelling the COVID-19 epidemic and implementation of population-wide interventions in Italy, *Nature Medicine* 26 (2020) 855–860.
- [13] C. Anastassopoulou, L. Russo, A. Tsakris, C. Siettos, Data-based analysis, modelling and forecasting of the COVID-19 outbreak, *PLOS ONE* 15 (3) (2020) 1–21.
- [14] L. Peng, W. Yang, D. Zhang, C. Zhuge, L. Hong, Epidemic analysis of COVID-19 in China by dynamical modeling, *arXiv preprints (2002.06563)* (2020).
- [15] H. Wang, Z. Wang, Y. Dong, R. Chang, C. Xu, X. Yu, S. Zhang, L. Tsamlag, M. Shang, J. Huang, Y. Wang, G. Xu, T. Shen, X. Zhang, Y. Cai, Phase-adjusted estimation of the number of Coronavirus Disease 2019 cases in Wuhan, China, *Cell Discovery* 6 (2020).
- [16] F. Della Rossa, D. Salzano, A. Di Meglio, F. De Lellis, M. Coraggio, C. Calabrese, A. Guarino, R. Cardona, P. DeLellis, D. Liuzza, F. Lo Iudice, G. Russo, M. Di Bernardo, A network model of Italy shows that intermittent regional strategies can alleviate the COVID-19 epidemic, *Nature Communications* 11 (2020) 5106.
- [17] M. Gatto, E. Bertuzzo, L. Mari, S. Miccoli, L. Carraro, R. Casagrandi, A. Rinaldo, Spread and dynamics of the COVID-19 epidemic in Italy: Effects of emergency containment measures, *Proceedings of the National Academy of Sciences* 117 (19) (2020) 10484–10491.

- [18] E. Lavezzo, E. Franchin, C. Ciavarella, G. Cuomo-Dannenburg, L. Barzon, C. Del Vecchio, L. Rossi, R. Manganelli, A. Loregian, N. Navarin, D. Abate, M. Sciro, S. Merigliano, E. De Canale, M. C. Vanuzzo, V. Besutti, F. Saluzzo, F. Onelia, M. Pacenti, S. G. Parisi, G. Carretta, D. Donato, L. Flor, S. Cocchio, G. Masi, A. Sperduti, L. Cattarino, R. Salvador, M. Nicoletti, F. Caldart, G. Castelli, E. Nieddu, B. Labella, L. Fava, M. Drigo, K. A. M. Gaythorpe, K. E. C. Ainslie, M. Baguelin, S. Bhatt, A. Boonyasiri, O. Boyd, L. Cattarino, C. Ciavarella, H. L. Coupland, Z. Cucunubá, G. Cuomo-Dannenburg, B. A. Djafaara, C. A. Donnelly, I. Dorigatti, S. L. van Elsland, R. FitzJohn, S. Flaxman, K. A. M. Gaythorpe, W. D. Green, T. Hallett, A. Hamlet, D. Haw, N. Imai, B. Jeffrey, E. Knock, D. J. Laydon, T. Mellan, S. Mishra, G. Nedjati-Gilani, P. Nouvellet, L. C. Okell, K. V. Parag, S. Riley, H. A. Thompson, H. J. T. Unwin, R. Verity, M. A. C. Vollmer, P. G. T. Walker, C. E. Walters, H. Wang, Y. Wang, O. J. Watson, C. Whittaker, L. K. Whittles, X. Xi, N. M. Ferguson, A. R. Brazzale, S. Toppo, M. Trevisan, V. Baldo, C. A. Donnelly, N. M. Ferguson, I. Dorigatti, A. Crisanti, Suppression of a SARS-CoV-2 outbreak in the Italian municipality of Vo', *Nature* 584 (2020) 425–429.
- [19] A. Pugliese, S. Sottile, Inferring the COVID-19 infection curve in Italy, arXiv preprints (2004.09404) (2020).
- [20] G. Albi, L. Pareschi, M. Zanella, Control with uncertain data of socially structured compartmental epidemic models, arXiv preprints (2004.13067) (2020).
- [21] M. Rosenblatt, Remarks on some nonparametric estimates of a density function, *The Annals of Mathematical Statistics* 27 (3) (1956) 832–837.
- [22] E. Parzen, On estimation of a probability density function and mode, *The Annals of Mathematical Statistics* 33 (3) (1962) 1065–1076.
- [23] R. E. Caflisch, Monte Carlo and quasi-Monte Carlo methods, in: *Acta numerica*, 1998, Vol. 7 of *Acta Numer.*, Cambridge Univ. Press, Cambridge, 1998, pp. 1–49.
- [24] H. Niederreiter, Random number generation and quasi-Monte Carlo methods, CBMS-NSF regional conference series in applied mathematics, SIAM, 1992.
- [25] M. D. McKay, R. J. Beckman, W. J. Conover, Comparison of three methods for selecting values of input variables in the analysis of output from a computer code, *Technometrics* 21 (2) (1979) 239–245.
- [26] I. Babuška, F. Nobile, R. Tempone, A stochastic collocation method for elliptic partial differential equations with random input data, *SIAM Review* 52 (2) (2010) 317–355.
- [27] D. Xiu, J. Hesthaven, High-order collocation methods for differential equations with random inputs, *SIAM J. Sci. Comput.* 27 (3) (2005) 1118–1139.
- [28] A. Cohen, R. DeVore, Approximation of high-dimensional parametric PDEs, *Acta Numerica* 24 (2015) 1–159.
- [29] M. D. Gunzburger, C. G. Webster, G. Zhang, Stochastic finite element methods for partial differential equations with random input data, *Acta Numerica* 23 (2014) 521–650.
- [30] I. H. Sloan, H. Woźniakowski, When are quasi-Monte Carlo algorithms efficient for high-dimensional integrals?, *J. Complexity* 14 (1) (1998) 1–33.
- [31] I. M. Sobol', Sensitivity estimates for nonlinear mathematical models, *Math. Modeling Comput. Experiment* 1 (4) (1993) 407–414 (1995).
- [32] D. Cacuci, *Sensitivity & Uncertainty Analysis, Volume 1: Theory*, no. v. 1, CRC Press, 2003.
- [33] A. Saltelli, M. Ratto, T. Andres, F. Campolongo, J. Cariboni, D. Gatelli, M. Saisana, S. Tarantola, *Global Sensitivity Analysis: The Primer*, Wiley, 2008.
- [34] L. Formaggia, A. Guadagnini, I. Imperiali, V. Lever, G. Porta, M. Riva, A. Scotti, L. Tamellini, Global sensitivity analysis through polynomial chaos expansion of a basin-scale geochemical compaction model, *Computational Geosciences* 17(1) (2013) 25–42.

- [35] E. Borgonovo, X. Lu, Is Time to Intervention in the COVID-19 Outbreak Really Important? A Global Sensitivity Analysis Approach, arXiv preprints (2005.01833) (2020).
- [36] M. D. Morris, Factorial sampling plans for preliminary computational experiments, *Technometrics* 33 (2) (1991) 161–174.
- [37] A. Dell’Oca, M. Riva, A. Guadagnini, Moment-based metrics for global sensitivity analysis of hydrological systems, *Hydrology and Earth System Sciences* 21 (12) (2017) 6219–6234.
- [38] A. M. Stuart, Inverse problems: A bayesian perspective, *Acta Numerica* 19 (2010) 451–559.
- [39] J. Carrera, S. P. Neuman, Estimation of aquifer parameters under transient and steady state conditions: 1. maximum likelihood method incorporating prior information, *Water Resources Research* 22 (2) (1986) 199–210.
- [40] S. M. Kay, *Fundamentals of statistical signal processing*, Prentice Hall PTR, 1993.
- [41] J. Nocedal, S. Wright, *Numerical Optimization*, 1st Edition, Springer-Verlag New York, 1999.
- [42] G. Porta, L. Tamellini, V. Lever, M. Riva, Inverse modeling of geochemical and mechanical compaction in sedimentary basins through polynomial chaos expansion, *Water Resources Research* 50 (12) (2014).
- [43] J. Berger, *Statistical Decision Theory and Bayesian Analysis*, Springer Series in Statistics, Springer, 1985.
- [44] C. C. Drovandi, *Approximate Bayesian Computation*, American Cancer Society, 2017, pp. 1–9.
- [45] H. Rue, S. Martino, N. Chopin, Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 71 (2) (2009) 319–392.
- [46] J. M. Ver Hoef, P. L. Boveng, Quasi-poisson vs. negative binomial regression: How should we model overdispersed count data?, *Ecology* 88 (11) (2007) 2766–2772.
- [47] A. Lindén, S. Mäntyniemi, Using the negative binomial distribution to model overdispersion in ecological count data, *Ecology* 92 (7) (2011) 1414–1421.
- [48] D. Firth, Bias reduction of maximum likelihood estimates, *Biometrika* 80 (1) (1993) 27–38.
- [49] J. Witteveen, S. Sarkar, H. Bijl, Modeling physical uncertainties in dynamic stall induced fluid–structure interaction of turbine blades using arbitrary polynomial chaos, *Computers and Structures* 85 (11) (2007) 866–878.
- [50] S. Oladyshkin, W. Nowak, Data-driven uncertainty quantification using the arbitrary polynomial chaos expansion, *Reliability Engineering and System Safety* 106 (0) (2012) 179 – 190.
- [51] C. Schillings, C. Schwab, Sparse, adaptive Smolyak quadratures for Bayesian inverse problems, *Inverse Problems* 29 (6) (2013).
- [52] O. G. Ernst, B. Sprungk, L. Tamellini, On Expansions and Nodes for Sparse Grid Collocation of Log-normal Elliptic PDEs, Arxiv e-prints (1906.01252), accepted. Also available as IMATI report 19-02 (2019).
- [53] Y. Marzouk, D. Xiu, A stochastic collocation approach to bayesian inference in inverse problems, in: *Communications in computational physics* 6, 2009, pp. 826–847.
- [54] A. Manzoni, S. Pagani, T. Lassila, Accurate solution of bayesian inverse uncertainty quantification problems combining reduced basis methods and reduction error models, *SIAM/ASA Journal on Uncertainty Quantification* 4 (1) (2016) 380–412.
- [55] G. Golub, V. Pereyra, Separable nonlinear least squares: the variable projection method and its applications, *Inverse Problems* 19 (2) (2003) R1–R26.

- [56] P. De Valpine, A. Hastings, Fitting population models incorporating process noise and observation error, *Ecological Monographs* 72 (1) (2002) 57–76.
- [57] D. He, E. L. Ionides, A. A. King, Plug-and-play inference for disease dynamics: measles in large and small populations as a case study, *Journal of The Royal Society Interface* 7 (43) (2010) 271–283.
- [58] L. Martín-Fernández, G. Gilioli, E. Lanzarone, J. Míguez, S. Pasquali, F. Ruggeri, D. P. Ruiz, A rao-blackwellized particle filter for joint parameter estimation and biomass tracking in a stochastic predator-prey system, *Mathematical Biosciences and Engineering* 11 (1551-0018_2014_3_573) (2014) 573.
- [59] J. Ritt, *Differential Algebra*, Colloquium publications, American Mathematical Society, 1950.
- [60] M. Eisenberg, Input-output equivalence and identifiability: some simple generalizations of the differential algebra approach, *arXiv preprints* (1302.5484v2) (Feb. 2019).
- [61] M. P. Saccomani, S. Audoly, L. D’Angiò, Parameter identifiability of nonlinear systems: the role of initial conditions, *Automatica* 39 (4) (2003) 619–632.
- [62] H. Hong, A. Ovchinnikov, G. Pogudin, C. Yap, Global identifiability of differential models, *Communications on Pure and Applied Mathematics* 73 (9) (2020) 1831–1879. doi:10.1002/cpa.21921.
- [63] H. Miao, X. Xia, A. S. Perelson, H. Wu, On identifiability of nonlinear ODE models and applications in viral dynamics, *SIAM Review* 53 (1) (2011) 3–39.
- [64] N. Evans, M. Chapman, M. Chappell, K. Godfrey, The structural identifiability of a general epidemic (SIR) model with seasonal forcing, *IFAC Proceedings Volumes* 35 (1) (2002) 109–114, 15th IFAC World Congress.
- [65] N. D. Evans, L. J. White, M. J. Chapman, K. R. Godfrey, M. J. Chappell, The structural identifiability of the susceptible infected recovered model with seasonal forcing, *Mathematical Biosciences* 194 (2) (2005) 175–197.
- [66] R. Bellman, K. Åström, On structural identifiability, *Mathematical Biosciences* 7 (3) (1970) 329–339.
- [67] A. Ovchinnikov, G. Pogudin, P. Thompson, Input-output equations and identifiability of linear ode models, *arXiv preprints* (1910.03960) (2020).
- [68] A. Ovchinnikov, A. Pillay, G. Pogudin, T. Scanlon, Computing all identifiable functions for ode models, *arXiv preprints* (2004.07774) (2020).
- [69] M. C. Eisenberg, M. A. Hayashi, Determining identifiable parameter combinations using subset profiling, *Mathematical Biosciences* 256 (2014) 116–126.
- [70] A. Raue, C. Kreutz, T. Maiwald, J. Bachmann, M. Schilling, U. Klingmüller, J. Timmer, Structural and practical identifiability analysis of partially observed dynamical models by exploiting the profile likelihood, *Bioinformatics* 25 (15) (2009) 1923–1929.
- [71] M. Iglesias, Z. Sawlan, M. Scavino, R. Tempone, C. Wood, Bayesian inferences of the thermal properties of a wall using temperature and heat flux measurements, *International Journal of Heat and Mass Transfer* 116 (2018) 417–431.
- [72] V. Kolehmainen, T. Tarvainen, S. Arridge, J. Kaipio, Marginalization of uninteresting distributed parameters in inverse problems-application to diffuse optical tomography, *International Journal for Uncertainty Quantification* 1 (1) (2011) 1–17.
- [73] F. Ruggeri, Z. Sawlan, M. Scavino, R. Tempone, A hierarchical bayesian setting for an inverse problem in linear parabolic pdes with noisy boundary conditions, *Bayesian Anal.* 12 (2) (2017) 407–433.
- [74] A.-N. Spiess, N. Neumeyer, An evaluation of R2 as an inadequate measure for nonlinear models in pharmacological and biochemical research: a Monte Carlo approach, *BMC Pharmacology* 10 (1) (2010).

- [75] G. Wasilkowski, H. Wozniakowski, Explicit cost bounds of algorithms for multivariate tensor product problems, *Journal of Complexity* 11 (1) (1995) 1 – 56.
- [76] W. C. Roda, M. B. Varughese, D. Han, M. Y. Li, Why is it difficult to accurately predict the COVID-19 epidemic?, *Infectious Disease Modelling* 5 (2020) 271–281.
- [77] Q. Long, M. Scavino, R. Tempone, S. Wang, A Laplace method for under-determined bayesian optimal experimental designs, *Computer Methods in Applied Mechanics and Engineering* 285 (2015) 849–876.
- [78] J. Beck, B. M. Dia, L. F. Espath, Q. Long, R. Tempone, Fast bayesian experimental design: Laplace-based importance sampling for the expected information gain, *Computer Methods in Applied Mechanics and Engineering* 334 (2018) 523–553.
- [79] J. Beck, B. Mansour Dia, L. Espath, R. Tempone, Multilevel double loop Monte Carlo and stochastic collocation methods with importance sampling for bayesian optimal experimental design, *International Journal for Numerical Methods in Engineering* 121 (15) (2020) 3482–3503.
- [80] A. Schöniger, T. Wöhling, L. Samaniego, W. Nowak, Model selection on solid ground: Rigorous comparison of nine ways to evaluate bayesian model evidence, *Water Resources Research* 50 (12) (2014) 9484–9513.
- [81] G. Claeskens, N. L. Hjort, *Model Selection and Model Averaging*, no. 9780521852258 in Cambridge Books, Cambridge University Press, 2008.
- [82] K. P. Burnham, D. R. Anderson, *Model selection and inference: a practical information-theoretical approach*, New-York: Springer-Verlag (1998).
- [83] M. Drton, M. Plummer, A bayesian information criterion for singular models, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 79 (2) (2017) 323–380.