

RESEARCH ARTICLE

OutbreakFlow: Model-based Bayesian inference of disease outbreak dynamics with invertible neural networks and its application to the COVID-19 pandemics in Germany

Stefan T. Radev^{1*}, Frederik Graw², Simiao Chen^{3,4}, Nico T. Mutters⁵, Vanessa M. Eichel⁶, Till Bärnighausen^{3,7,8}, Ullrich Köthe⁹

1 Institute of Psychology, Heidelberg University, Heidelberg, Germany, **2** BioQuant - Center for Quantitative Biology, Heidelberg University, Heidelberg, Germany, **3** Heidelberg Institute of Global Health, Heidelberg, Germany, **4** Chinese Academy of Medical Sciences and Peking Union Medical College, Beijing, China, **5** Institute for Hygiene and Public Health, University Hospital Bonn, Bonn, Germany, **6** Center of Infectious Diseases, University Hospital Heidelberg, Heidelberg, Germany, **7** Harvard T.H. Chan School of Public Health, Boston, Massachusetts, United States of America, **8** Africa Health Research Institute, Durban, South Africa, **9** Computer Vision and Learning Lab, Heidelberg University, Heidelberg, Germany

* stefan.radev93@gmail.com



OPEN ACCESS

Citation: Radev ST, Graw F, Chen S, Mutters NT, Eichel VM, Bärnighausen T, et al. (2021) OutbreakFlow: Model-based Bayesian inference of disease outbreak dynamics with invertible neural networks and its application to the COVID-19 pandemics in Germany. *PLoS Comput Biol* 17(10): e1009472. <https://doi.org/10.1371/journal.pcbi.1009472>

Editor: Mark M. Tanaka, University of New South Wales, AUSTRALIA

Received: January 13, 2021

Accepted: September 23, 2021

Published: October 25, 2021

Copyright: © 2021 Radev et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: All data are available at GitHub: <https://github.com/stefanradev93/AlAgainstCorona>.

Funding: SR was supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation; grant number GRK 2277 "Statistical Modeling in Psychology") and the Covid-19 Research Grants by Google Cloud. SC was supported by Bill & Melinda Gates Foundation

Abstract

Mathematical models in epidemiology are an indispensable tool to determine the dynamics and important characteristics of infectious diseases. Apart from their scientific merit, these models are often used to inform political decisions and interventional measures during an ongoing outbreak. However, reliably inferring the epidemical dynamics by connecting complex models to real data is still hard and requires either laborious manual parameter fitting or expensive optimization methods which have to be repeated from scratch for every application of a given model. In this work, we address this problem with a novel combination of epidemiological modeling with specialized neural networks. Our approach entails two computational phases: In an initial training phase, a mathematical model describing the epidemic is used as a coach for a neural network, which acquires global knowledge about the full range of possible disease dynamics. In the subsequent inference phase, the trained neural network processes the observed data of an actual outbreak and infers the parameters of the model in order to realistically reproduce the observed dynamics and reliably predict future progression. With its flexible framework, our simulation-based approach is applicable to a variety of epidemiological models. Moreover, since our method is fully Bayesian, it is designed to incorporate all available prior knowledge about plausible parameter values and returns complete joint posterior distributions over these parameters. Application of our method to the early Covid-19 outbreak phase in Germany demonstrates that we are able to obtain reliable probabilistic estimates for important disease characteristics, such as generation time, fraction of undetected infections, likelihood of transmission before symptom onset, and reporting delays using a very moderate amount of real-world observations.

(Project INV-006261) and Alexander von Humboldt-Foundation. FG was supported by a Fellowship from the Chica and Heinz Schaller Foundation. UK was supported by Informatics for Life funded by the Klaus Tschira Foundation and by the Cluster of Excellence STRUCTURES funded by DFG. TB was supported by the Alexander von Humboldt Foundation through the Alexander von Humboldt Professor award, funded by the German Federal Ministry of Education and Research. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

Author summary

Emerging infections and epidemic outbreaks are associated with large uncertainties concerning data integrity that challenge the timely detection of disease characteristics and dynamics. Robust parameter inference for mathematical models aiming to describe these dynamics is essential to predict the progression of an epidemic and inform on appropriate public health interventions. In this study, we present a novel method based on invertible neural networks that allows inference of important epidemiological characteristics in case of limited data, thereby allowing for reliable uncertainty quantification. The method circumvents common challenges associated with sparse data by using simulation-based training of an expressive generative neural network. Applying our method to data of the early Covid-19 epidemic in Germany, we are able to obtain reliable estimates on important disease characteristics, such as the proportion of infected individuals remaining undetected, despite limited observations during early outbreak dynamics.

This is a *PLOS Computational Biology Methods* paper.

Introduction

Assessing important disease characteristics and transmission dynamics is of utmost importance in the case of new epidemic outbreaks in order to forecast their progression and guide effective public health measurements. Mathematical models that provide a reliable representation of the processes driving the dynamics of an epidemic are an essential tool for this task (see for example [1]). In the case of communicable diseases, these models typically take the form of systems of ordinary differential equations governing the transitions between different population compartments, such as, “Susceptible”, “Infected”, and “Recovered” (SIR) [1]. Provided that intrinsic properties of the disease (e.g., transmission rates and recovery periods) are already known, SIR models and their extensions are successfully used to simulate outcomes of possible public health interventions.

However, for newly emerging infectious diseases, such as Covid-19, most of these properties are initially unknown and must be estimated before realistic predictions can be made. The task of determining these properties is additionally hampered by limited data availability and integrity within the early outbreak phase. During the initial outbreak of the Covid-19 pandemic, model-based inference was used to provide rapid estimates of key epidemiological parameters, which otherwise can be difficult to infer directly from primary clinical tracing data. For instance, an earlier study [2] incorporated domestic and international travel from and to Wuhan city in a SEIR model and used reported cases outside of Wuhan to infer the reproduction number R_0 and epidemic doubling time. Similar approaches were used to estimate the reproduction number of Covid-19 in various other settings [3–5]. Other studies aimed at identifying critical epidemiological characteristics, such as age-specific mortality rates [6], the impact of implemented control measures on disease transmission [7, 8], or the fraction of undocumented infections [9].

Importantly, since such SIR-type models are employed to forecast the dynamics of an epidemic dependent on public interventions or seasonal effects, reliable inference of such key epidemiological parameters and trustworthy uncertainty quantification is paramount to support

decision making. The estimation of hidden model parameters from observations of model outcomes (*inverse inference*) is also referred to as *model calibration* in the medical decision and health policy modeling literature [10]. Traditionally, model calibration has been considered as an optimization problem seeking the best possible parameter configuration explaining the data (e.g., by performing non-linear least squares minimization or relying on maximum likelihood estimation [11]). However, the resulting optimization and maximum likelihood methods focus on point estimates for the individual parameters and usually lack appropriate approaches to assess their accuracy. This is a severe disadvantage, because reliable uncertainty quantification is crucial when the estimates shall be used to model and predict future outcomes.

In contrast to optimization and maximum likelihood approaches, Bayesian methods provide a principled way to quantify the uncertainty of inferred model parameters, as they return full posterior distributions for the unknown parameters rather than single point estimates. Markov chain Monte Carlo (MCMC) sampling represents one of the classical approaches to Bayesian model calibration [12], and it has been extensively used in Covid-19 studies to infer model parameters describing the dynamics of the disease [3, 4, 8, 13]. However, Bayesian model calibration is computationally expensive and depends explicitly on the evaluation of the likelihood function of model parameters. When the likelihood function is intractable or unknown, approximate Bayesian computation (ABC) can be used to approximate the posterior distribution of parameters [14, 15]. However, standard ABC methods notoriously suffer from poor scalability (i.e., the efficient application to large data sets and complex models), which confines their utility to relatively simple models containing only a few parameters [16].

In order to overcome the limitations of individual parameter estimation methods, our approach aims to combine the advantages of optimization-based and Bayesian methods by using specialized invertible neural networks. In particular, we develop a novel methodological framework based on a neural network architecture called BayesFlow [17] to facilitate model-based inference with a primary, but not exclusive, focus on epidemiology. Our method can incorporate an arbitrary number of epidemiological time series (or other type of temporal information) and can, in principle, be applied to any dynamic model described by (stochastic) ordinary differential equations. Moreover, since the length of the time series is not pre-determined during training and inference, we can consider additional information in our analyzes without having to re-train the networks. By using extensive simulation-based training, our method circumvents the general necessity for large training sets that are lacking during emerging epidemics. Additionally, our method returns posterior distributions which are fully compatible with a Bayesian interpretation and can thus be used to assess the uncertainty associated with any estimation and prediction quantity.

We demonstrate the feasibility of our method by analyzing public Covid-19 data for Germany and the individual German federal states based on the reported daily number of infected, recovered and deceased cases during the first months of the pandemic. Our neural network is trained using simulations from a customized SEIR-model variant [18], in combination with an observation model accounting for the differences between true and reported case numbers, and an intervention model describing the intervention measures for prevention and control imposed by German authorities [8]. Despite the limited number of measurements and the considerable complexity of the model with 34 unknown parameters in total, our network manages to extract information for more than half of the parameters. Credibility intervals of our parameter estimates are well in line with independently published results, and re-simulations starting at our estimated parameters reproduce the observed time series very well. In particular, our inference suggests that approximately four fifths of all infectious individuals remained undetected across all German federal states.

Materials and methods

Data

The model was applied to epidemiological data on the number of reported Covid-19 cases (infected, recovered and deceased) in Germany and the individual federal states (only infected and diseased) from March 01, 2020 until June 11, 2020. Data were collected from publicly available sources during the same time period and were not subsequently cleaned or corrected in the aftermath. Therefore, all sources of uncertainty remain in the data as they would have in the early days of an ongoing pandemic. Code and scripts for reproducing all results and figures as well as the general framework of OutbreakFlow for training new networks on new models are available at <https://github.com/stefanradev93/AIAgainstCorona>.

Neural Bayesian parameter estimation

Following a Bayesian approach for parameter estimation requires prior knowledge about reasonable parameter ranges [12]. Combining this prior knowledge with information extracted from the observed data leads to a posterior distribution which is generally narrower than the prior and thus expresses our updated state of knowledge and associated uncertainty for the individual parameters. More formally, let θ be the vector of all unknown model parameters and $X := x_{1:T} = [x_1, \dots, x_T]$ a multivariate epidemiological time series with T individual time points indicating, for instance, the number of infected, recovered and diseased individuals. Then the well-known analytical formula for the posterior according to Bayes' rule is

$$p(\theta | X) = \frac{p(X | \theta) p(\theta)}{\int p(X | \theta) p(\theta) d\theta} \quad (1)$$

where $p(X|\theta)$ represents the likelihood of observing data X when the true parameters are θ , $p(\theta)$ is the prior distribution encoding our knowledge about plausible parameter configurations for θ , and the denominator represents a normalizing constant (usually called the evidence).

Despite being conceptually simple, this formula poses two major challenges in the present setting: (i) Efficient and accurate approximation of the intractable posterior $p(\theta|X)$ is challenging; (ii) The likelihood is only implicitly defined via realizations $X \sim p(X|\theta)$ generated by repeatedly running simulations of the underlying epidemiological model with different θ .

We solve both problems with our recently proposed neural Bayesian inference architecture called BayesFlow, which is explained in full mathematical details in the corresponding methodological work [17]. The core component of BayesFlow is an *invertible neural network* which enables a bidirectional flow of information. During the training phase, the invertible network is run in forward direction to learn an accurate model $q(\theta|X) \approx p(\theta|X)$ for the posterior distribution of parameters θ given observations X , using a large number of simulated pairs $(X_i, \theta_i) \sim p(X|\theta) p(\theta)$ as training data. During the inference phase, the network makes use of its invertible architecture and operates in the inverse direction to estimate the posterior $q(\theta|X = x^{\text{obs}}) \approx p(\theta|X = x^{\text{obs}})$ for the *actually* observed data x^{obs} .

Validation experiments reported in [17] have demonstrated for various model systems, that the BayesFlow method (i) can estimate complex stochastic models of widely varying data types and sizes (e.g., population time series, predator-pray population time series, human response-time data); (ii) outperforms variational or dropout methods for uncertainty quantification; (iii) learns data representations which are more informative than manually selected summary statistics; and (iv) outperforms case-based methods such as ABC and MCMC, whose computations must be re-run from scratch for every observed dataset.

The latter advantage is called “amortized inference”: a BayesFlow network learns to generalize the training knowledge and can efficiently apply it to multiple real observations without retraining. The network’s training effort thus quickly *amortizes* over a sequence of inference queries (e.g., time series), in contrast to sampling methods (e.g., MCMC), which cannot leverage experience and require the same large computational effort for every query. In addition, fast amortized inference facilitates model validation by enabling efficient probabilistic calibration and posterior predictive checks on large validation datasets [17].

OutbreakFlow—The BayesFlow approach to epidemiological inference

We propose OutbreakFlow, an instantiation of our BayesFlow architecture that utilizes a novel combination of three jointly trained neural modules to analyze noisy multivariate time series with potentially long-term temporal dependencies, as are typical in the context of epidemiology. It can process both short and long time series and can thus perform efficient Bayesian updating as new data become available (e.g., on a daily basis in case of Covid-19). Moreover, our method can incorporate additional prior knowledge in the formulation of the underlying generative model.

Our neural architecture comprises three sub-networks: (i) a convolutional *filtering* network performing noise reduction and feature extraction on the raw observational data; (ii) a recurrent *summary* network reducing filtered time series of *arbitrary* length to statistical summaries of *fixed* size; and (iii) an invertible *inference* network performing Bayesian parameter inference, given the learned summaries of the observations. Fig 1 depicts the training and inference phase with our inference architecture and its essential elements.

The design of the convolutional network is inspired by the well-known Inception network, which has shown tremendous success in computer vision tasks [19]. In particular, our network is implemented as a deep fully convolutional network which applies adjustable one-dimensional filters of different size at each level (cf. Fig 1). The intuition behind this design is that filters of different size might capture patterns at different temporal scales (e.g., a filter of size one will capture daily fluctuations whereas a filter of size seven will capture weekly dynamics). This, in turn, should ease the task of extracting informative temporal features for parameter estimation.

The output of the convolutional network is a multivariate sequence containing the filtered epidemiological time series. In order to reduce the filtered sequence to a fixed-size vector, we pass it through a long-short term memory (LSTM) recurrent network [20]. In contrast to standard feed-forward neural networks, LSTM networks incorporate feedback connections which make them ideally suited for processing sequences of observations such as time series.

A standard LSTM network consists of a cell and three gates. The cell stores the *internal memory* of the network over arbitrary temporal intervals. The three gates, comprising an input gate, an output gate, and a forget gate, interact in determining the importance of old and new information. Importantly, LSTM networks can deal with sequences of varying length, which enables them to process data whose duration is dictated by data availability and to perform online inference, that is, process new data instantly as they become available. In contrast to predefined pooling operations (e.g., mean, max, or variance), our recurrent networks learn pooling operations that are adapted to the data and can thus extract potentially richer representations. In this way, our inference architecture learns to filter and extract the most informative features from the noisy observations in an end-to-end manner. Thus, the user is freed from the difficult (and usually suboptimal) task to hand-engineer suitable data features (summary statistics). Finally, the inference network has the task of *inverting* the forward model

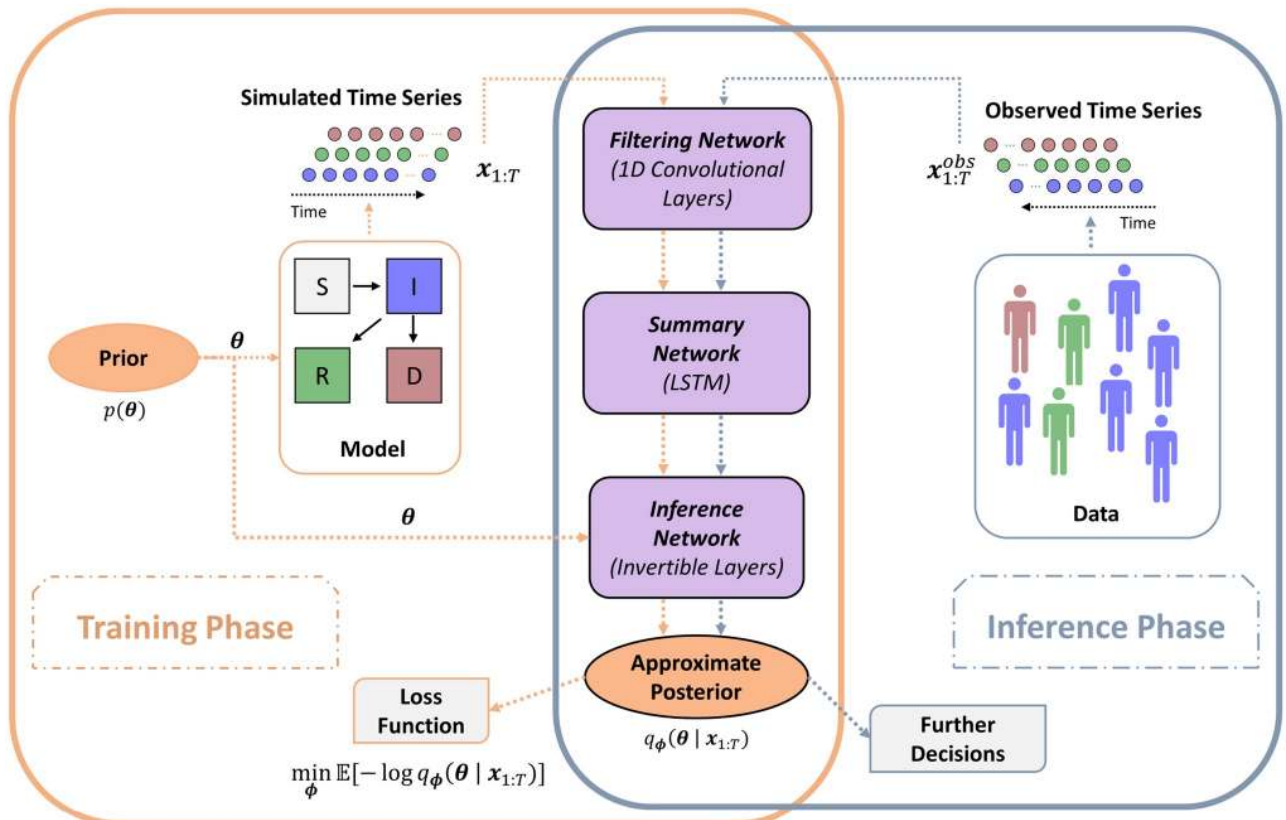


Fig 1. Structure and workflow of OutbreakFlow. During the training phase (orange frame on the left), the assumed epidemiological model is used to simulate time series resembling the observed epidemiological data, based on prior distributions of the unknown parameters. The synthetic time series are used to train the composite neural network consisting of (i) a convolutional filtering network, which extracts relevant features while preserving the temporal structure of the data, (ii) a summary network, which reduces the transformed time series to a fixed-sized vector of maximally informative representations, and (iii) an inference network, which estimates the joint parameter posterior from these data representations. During the inference phase (blue frame on the right), the real epidemiological data $x_{1:T}^{obs}$ are passed to the trained network to infer the posterior distribution of the unknown disease parameters. A full description of the architecture and the methodology is provided within [Materials and methods](#).

<https://doi.org/10.1371/journal.pcbi.1009472.g001>

given the information extracted by the convolutional and recurrent networks (see also [17] for more details on the design of invertible networks for Bayesian inference).

The invertible network has two modes of operation. During training, the network is only evaluated in the *forward direction* and encouraged via a suitable optimization criterion to transform the posterior into a simple base distribution (e.g., Gaussian) from which samples can be easily obtained. Thus, the inference network integrates information from both the prior and the data-generating mechanism (i.e., the implicit likelihood).

During inference, the inference network is only evaluated in the *inverse direction* using conditional information from real observed data passed through the filtering and summary networks. The posterior is approximated by repeatedly sampling from the simple base distribution and applying the inverse of the forward transformation learned during the training phase. Importantly, this method recovers the true posterior under perfect convergence of the optimization method [17].

More formally, let us denote the functions represented by the three networks as g , h , and f . Then the filtering network determines a filtered time series $\tilde{x}_{1:T} = g(x_{1:T})$ from observed data $x_{1:T}$, where the number of time steps T depends on data availability. The summary network

turns the output of the filtering network into a fixed-size representation $y = h(\tilde{x}_{1:T})$ by keeping only the final output vector of the LSTM network, which encodes the accumulated information over all observed time steps. Finally, the inference network generates samples $\hat{\theta} \sim q(\theta | x_{1:T})$ from the parameter posterior by computing $\hat{\theta} = f^{-1}(y, z)$ with normally distributed random vectors $z \sim \mathcal{N}(0, \mathbb{I})$. The parameters of all three networks are optimized jointly during the training phase. Denoting the vector of all trainable network parameters as ϕ , the three networks solve the following optimization criterion

$$\hat{\phi} = \underset{\phi}{\operatorname{argmin}} \mathbb{E}_{X \sim p(X)} [\mathbb{KL}(p(\theta | X) || q_{\phi}(\theta | X))] \quad (2)$$

$$= \underset{\phi}{\operatorname{argmin}} \mathbb{E}_{(X, \theta) \sim p(X, \theta)} [-\log q_{\phi}(\theta | X)] \quad (3)$$

where $\mathbb{KL}(p || q)$ denotes the Kullback-Leibler divergence [21] between probability density functions p and q . We approximate the latter expectation via its empirical mean over realizations $(X, \theta) \sim p(\theta, X)$ obtained via simulations from a dynamic model.

As previously mentioned, one of the most important advantages of our method is *amortized inference*, owing to the fact that we approximate the posterior *globally* via a single set of network parameters $\hat{\phi}$. This is especially advantageous in epidemiological contexts, where the same model is applied in multiple populations (countries, cultures) or at different scales (states, regions), since the same pre-trained network architecture can be repeatedly utilized for different populations and scales. Indeed, in the following real-world application, we demonstrate efficient amortized inference and excellent predictive performance with a single architecture applied simultaneously to epidemiological data from all German federal states.

The epidemiological model

In order to account for the specific nature of the current Covid-19 outbreak, our epidemiological model consists of three submodels: (i) a disease model describing the true dynamics of relevant population compartments; (ii) an intervention model describing the strengthening and relaxation of non-pharmaceutical intervention measures; and (iii) an observation model describing the deviations of publicly reported data from the true values. These models build upon the previous work of Khailaie et al. [18] and Dehning et al. [8], who adapted the general SIR approach to the specifics of the Covid-19 pandemic and the situation in Germany. Parameter priors are based on our current state of knowledge about disease characteristics and government measures, but are chosen very wide to prevent them from dominating the information extracted from the actual observations.

Disease model. The disease model is a system of non-linear ordinary differential equations (ODEs) distinguishing between six compartments: susceptible (S), exposed (E —infected individuals who do not show symptoms and are not yet infectious), infected (I —symptomatic cases that are infectious), carrier (C —infectious individuals who recover without being detected), recovered (R), and dead (D), see Fig 2. Note that direct recovery from the carrier state C covers all reasons why an infection might go undetected, including, among others, truly asymptomatic cases, lack of follow-up on pre-symptomatic cases, limited testing capacity under minor symptoms—our model does not differentiate between these possibilities. Observations with limited accuracy (as described by the observation model below) are available for the compartments I , R , and D . The true time series of all compartments are therefore considered latent and need to be estimated.

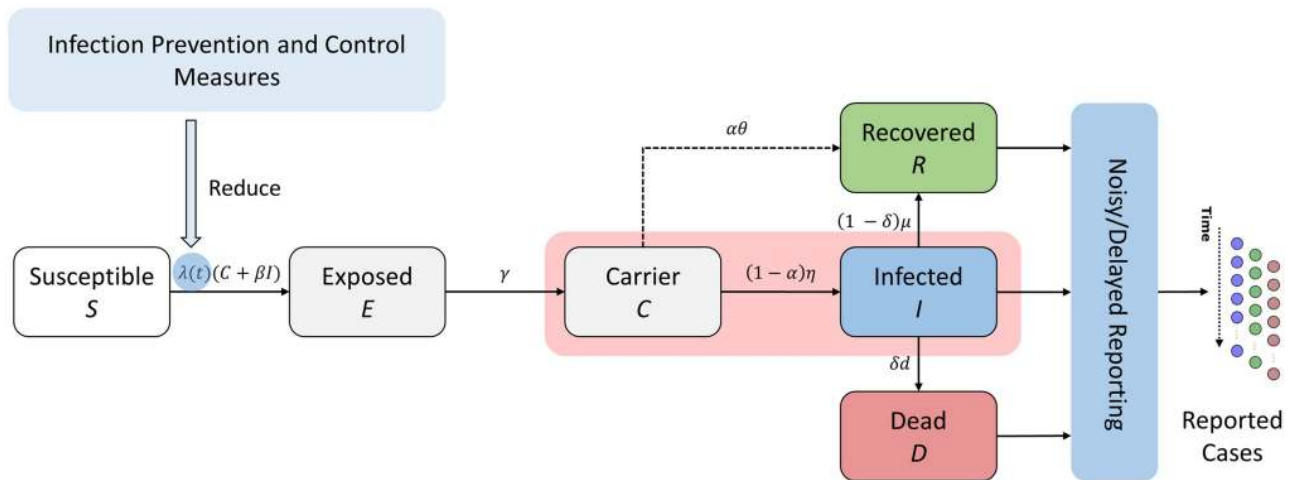


Fig 2. Sketch of the mathematical model describing the epidemiological dynamics of Covid-19. The model is of SEIR-type with six compartments: susceptible (*S*), exposed (*E*, i.e. infected but non-infectious), carrier (*C*, i.e. infectious but undetected), infected (*I*, i.e. infectious and diagnosed), recovered (*R*) and dead (*D*) individuals. In addition, the blue boxes indicate model extensions accounting for external factors, namely intervention measures affecting the transmission rate $\lambda(t)$ and imperfect case reporting due to noise or delays. Note, that data is only reported for the observable compartments *I*, *R*, and *D*. For a detailed description of the model and the different components see [Materials and methods](#).

<https://doi.org/10.1371/journal.pcbi.1009472.g002>

The ODEs for our epidemiological model are defined by:

$$\frac{dS}{dt} = -\lambda(t) \left(\frac{C + \beta I}{N} \right) S \tag{4}$$

$$\frac{dE}{dt} = \lambda(t) \left(\frac{C + \beta I}{N} \right) S - \gamma E \tag{5}$$

$$\frac{dC}{dt} = \gamma E - (1 - \alpha)\eta C - \alpha\theta C \tag{6}$$

$$\frac{dI}{dt} = (1 - \alpha)\eta C - (1 - \delta)\mu I - \delta d I \tag{7}$$

$$\frac{dR}{dt} = \alpha\theta C + (1 - \delta)\mu I \tag{8}$$

$$\frac{dD}{dt} = \delta d I \tag{9}$$

The meaning of the model parameters and their priors are detailed in S1 Table in [S1 Text](#). Prior ranges are based on considerations in [8] and [22]. All rate parameters are considered to be constant, except for the transmission rate $\lambda(t)$ which is considered to be time-dependent as it accounts for possible behavioral changes implied by non-pharmaceutical interventions.

Intervention model. The intervention model accounts for changes in the transmission rate $\lambda(t)$ due to non-pharmaceutical interventions and mitigation measures. Corresponding to the approach followed by [8], we define three change points for $\lambda(t)$ encoding an assumed transmission rate reduction in response to intervention measures imposed by the German authorities. Each change point is represented by a piece-wise linear function with three degrees

of freedom: the effect strength and the boundaries defining the time interval for the effect to fully manifest itself. The chosen priors express the expected effect of each measure to reduce the transmission rate roughly by half after the date when it comes into force, but we allow for wide uncertainty margins to facilitate learning of the actual behavior. In addition to the previous approach in [8], our model also includes a fourth change point expressing the assumption that an eventual withdrawal of effective intervention measures (officially or due to non-compliance) will lead to a slight increase of the transmission rate. Note that we assume that interventions do not affect the risk of infection upon contact with a detected infectious individual (β). Prior distributions and descriptions for the intervention model's parameters are given in S2 Table in [S1 Text](#).

Observation model. The observation model accounts for the fact that officially reported cases might not represent the true case numbers of the epidemics. It represents three error sources: a delay between actual infection and reporting, the weekly modulation of reporting rates (since testing and reporting activities are considerably reduced on weekends), and a noise term describing random fluctuations. Separate parameter sets are learned for each of the three publicly reported time series $I^{(obs)}$, $R^{(obs)}$, and $D^{(obs)}$ —the remaining compartments are considered unobservable. The relationship between the reported counts and their true values is described by the following set of discrete-time difference equations with time steps t measured in days.

$$I_t^{(obs)} = I_{t-1}^{(obs)} + (1 - f_I(t)) (1 - \alpha) \eta C_{t-L_I} + \sqrt{I_{t-1}^{(obs)}} \sigma_I \zeta_t \tag{10}$$

$$R_t^{(obs)} = R_{t-1}^{(obs)} + (1 - f_R(t)) (1 - \delta) \mu I_{t-L_R} + \sqrt{R_{t-1}^{(obs)}} \sigma_R \zeta_t \tag{11}$$

$$D_t^{(obs)} = D_{t-1}^{(obs)} + (1 - f_D(t)) \delta d I_{t-L_D} + \sqrt{D_{t-1}^{(obs)}} \sigma_D \zeta_t \tag{12}$$

where L_I , L_R , and L_D denote the reporting delays (lags), and denote σ_I , σ_R and σ_D the scales of multiplicative reporting noise for the respective compartments. The noise variables ζ_t follow a Student- t distribution with 4 degrees of freedom. The weekly modulation of reporting coverage $f_C(t)$ for each of the compartments $C \in \{I, R, D\}$ is computed as follows:

$$f_C(t) = (1 - A_C) \left(1 - \left| \sin\left(\frac{\pi}{7} t - 0.5 \Phi_C\right) \right| \right) \tag{13}$$

This equation yields three additional unknown parameters for the weekly modulation amplitudes A_I, A_R, A_D , and phases Φ_I, Φ_R, Φ_D , each. The prior distributions and descriptions for the observation model's parameters are listed in S3 Table in [S1 Text](#).

Neural network training

An OutbreakFlow is trained with simulated data by minimizing the negative log posterior according to [Eq 3](#). The training phase can be realized in different ways, depending on the modeling scenario and the modelers' computational budget. First of all, when only a single time series has to be analyzed, non-amortized methods like [23, 24] may outperform OutbreakFlow, because they constrain the simulation scope to the vicinity of the observed data. On the other hand, when the model has to be applied to *multiple* observed time series (e.g., to different federal states in Germany or even countries), our upfront training effort quickly amortizes, since a trained OutbreakFlow executes inference orders of magnitude faster than a case-based (non-amortized) method. We now outline three training modes for OutbreakFlow.

Simulation-based offline learning. Traditional simulation-based approaches utilize a pre-computed reference table $\mathcal{D}^{(S)}$, which is a large data structure containing S pairs (θ, X) of simulation parameters θ and corresponding synthetic observations X [25, 26]. This strategy has also been used in machine learning approaches to simulation-based inference [27, 28], where the problem of inverse inference resembles a supervised learning task. In the context of OutbreakFlow, the resulting *offline learning* method is outlined in S1 Algorithm in [S1 Text](#). It is particularly useful when calls to the simulator are computationally expensive: working with recorded synthetic data is then faster at the expense of higher memory demands during training.

Simulation-based online learning. Instead of pre-computing synthetic data, we can generate a potentially limitless number of training pairs (θ, X) on-the-fly. Since each simulation result is discarded after the corresponding backpropagation update, the network never encounters the same inputs twice and overfitting is impossible. Moreover, the training phase can continue as long as the network keeps improving, as measured by continuous performance monitoring. Online training is outlined in S2 Algorithm in [S1 Text](#) and is used for all experiments in this work. The present application lends itself to this approach, because the computational cost of running our epidemiological model is negligible, whereas more expensive simulations might become a bottleneck for this strategy.

Simulation-based hybrid learning. Offline and online learning represent two extremes on a continuum of training strategies. Hybrid learning methods combine these two strategies and allow for a more fine-grained allocation of the available simulation budget. For instance, [24] propose a round-based strategy, where each round incorporates its own simulation phase. Thus, the reference table is filled incrementally, and each round can reuse simulations from all previous rounds. Such a round-based training strategy is outlined in S3 Algorithm in [S1 Text](#).

Uncertainty calibration and computational faithfulness

Computational faithfulness refers to the ability of a Bayesian method to recover the correct target posterior in a given modeling scenario. It is an essential precondition for carrying out model-based predictions and interpreting the parameters of a model within a reference theoretical framework. We can estimate the computational faithfulness of any BayesFlow application using simulation-based calibration [29, SBC]. SBC is a diagnostic method which considers the performance of a sampling algorithm over the entire Bayesian joint model $p(\theta, X)$, regardless of the structure of the particular model. It leverages the fact that most Bayesian models are generative by construction as well as the self-consistency of the Bayesian joint model. Accordingly, the *average posterior distribution* over random draws from the Bayesian joint distribution $(\theta, X) \sim p(\theta, X)$ should always recover the prior distribution $p(\theta)$. In other words, for any given parameter combination θ^* , the following should hold:

$$p(\theta^*) = \int \int p(\theta^* | X) p(\theta, X) d\theta dX \quad (14)$$

Random draws from $p(\theta, X)$ are generated by first sampling a configuration θ from the prior $p(\theta)$ and then running the (stochastic) simulator with the sampled parameter configuration to obtain a synthetic outbreak trajectory. This process can be repeated multiple times and does not require an analytically tractable likelihood function. Importantly, if a Bayesian sampling method generates samples from the exact posterior, the equality implied by [Eq 14](#) should hold regardless of the particular form of the posterior. Thus, any violation of this equality indicates some error incurred by the sampling method. The reasons for such an error can be either inaccurate computation of the posterior or an erroneous implementation of the model itself [29].

In practice, we approximate this integral by an ensemble of samples from many posterior distributions estimated from simulated time series with known generating parameters. SBC uses a rank statistic (i.e., the number of posterior draws larger than the prior draw for each simulated time series) to compare the average posterior with the prior. If Eq 14 holds, then the rank statistic of each parameter will be uniformly distributed, allowing us to visually examine the equivalence using univariate histograms. An inspection of the rank histograms thus provides a way to validate the computational faithfulness of OutbreakFlow within the scope of the modeling assumptions [29].

A major disadvantage of SBC is that it can be extremely time-consuming, since it requires inverse inference on potentially thousands of simulated time series. In addition, the obtained posterior draws should exhibit no autocorrelation for SBC to yield reliable results. The latter requirement makes it even more expensive when using MCMC or other non-amortized Bayesian methods yielding dependent samples. Fortunately, amortized inference with OutbreakFlow alleviates these issues, since inference on multiple time series is extremely efficient and posterior draws are independent given perfectly converged networks [17]. Thus, validating the computational faithfulness of any OutbreakFlow application using SBC becomes a matter of seconds.

Outbreak prediction on the basis of estimated posteriors

The posteriors estimated by a trained OutbreakFlow network can be further used to make forecasts about the future dynamics of the pandemic, provided the parameters remain stationary. Due to changing intervention measures, population behavior, testing policies, and possible treatment advances, this is only true for a relatively short period beyond the observed time series, limiting the prediction horizon to a few weeks. Given observed time series $X := x_{1:T}$, the posterior predictive distribution for upcoming data $X' := x_{T+1:T'}$ is given by:

$$p(X' | X) = \int p(X' | \theta, X) p(\theta | X) d\theta \quad (15)$$

Although this quantity is hard to compute exactly, we can approximate it by running simulations with parameters sampled from the posterior: $\{\theta^{(m)} \sim p(\theta | X)\}_{m=1}^M$. Since the $\theta^{(m)}$ are drawn from the joint posterior, statistical dependencies and correlations between parameters are properly taken into account. The resulting ensemble of M simulated time series $\{\tilde{X}^{(m)}\}_{m=1}^M$ can now be used to obtain point predictions (e.g., by computing their mean or median at each future time point) and to quantify the uncertainty of future scenarios (e.g., by computing quantiles or standard deviations at each time point).

In addition to future predictions, posterior predictive re-simulations are also crucial for further model validation. If the estimated posteriors describe the real situation well, re-simulations should replicate the data that served to fit the model in the first place. Mismatches between original and re-simulated time series indicate *model misspecification* or a *simulation gap*, that is, errors arising because important aspects of the disease dynamics or unknown corruption of the observed data are not properly represented by the model. These errors are invisible to simulation-based calibration, because it solely assesses whether the posteriors conform to model specifications. Importantly, our OutbreakFlow experiments demonstrate very good agreement between original and re-simulated time series.

Results

OutbreakFlow as a research tool for inferring dynamics of emerging epidemics

In contrast to mainstream neural network applications, such as image or text analysis, analyzing the dynamics of emerging epidemics poses two major challenges: (i) data are usually sparse, with no large sets of training data available; and (ii) a reliable quantification of the uncertainty associated with neural network outputs, such as estimated parameters, is mandatory to allow for reliable subsequent evaluation of possible scenarios.

Standard neural network architectures do not live up to these challenges. Therefore, we developed a novel method on the basis of a neural network architecture called BayesFlow [17] that addresses the aforementioned challenges in two ways: (i) We leverage the epidemiological insight represented by SIR-type models by means of an alternative training procedure using simulated data—*simulation-based training*; and (ii) we use networks that are specifically designed to perform Bayesian uncertainty quantification over their outputs.

In our framework, a large number of plausible hypothetical scenarios simulating the assumed epidemiological dynamics is processed by the neural network until it becomes an expert in the interpretation of epidemiological observations. After completion of the training phase, the available real-world observations are passed to the network, which then estimates full Bayesian posterior distributions for the real-world parameters of interest. The design of our network architecture is depicted in Fig 1. The details of the method, as well as the model architecture are given in Materials and methods.

The ultimate goal of our approach is comparable to that of traditional simulation-based Bayesian inference methods, such as ABC. However, our method operates much faster and generalizes, without retraining, to any real-world dataset within the scope of its training expertise [17].

Testing and validation of OutbreakFlow

To validate our approach and test its performance in inferring parameter values in epidemiological models, we applied our architecture to a standard SIR-model describing the dynamics of an epidemic. This greatly simplified model is suitable for the initial two weeks of the pandemic and highlights essential properties of our approach. It distinguishes between susceptible, S , infected, I , and recovered, R , individuals with infection and recovery occurring at a constant transmission rate λ and recovery rate μ , respectively. The model is defined by the following system of ODEs:

$$\frac{dS}{dt} = -\lambda \left(\frac{SI}{N} \right) \quad (16)$$

$$\frac{dI}{dt} = \lambda \left(\frac{SI}{N} \right) - \mu I \quad (17)$$

$$\frac{dR}{dt} = \mu I, \quad (18)$$

with $N = S + I + R$ denoting the total population number. In addition to the ODE parameters λ and μ , we consider a reporting delay parameter L and a dispersion parameter ψ , which affect

the number of reported infected individuals via

$$I_t^{(obs)} \sim \text{NegBinomial}(I_{t-L}^{(new)}, \psi), \quad (19)$$

where $I_t^{(new)} = \lambda (S_t I_t / N)$ and we assume that the number of newly observed cases arises from a negative binomial distribution [30]. In addition, we estimate the initial number of infected individuals I_0 , so the full parameter vector of interest becomes $\theta = (\lambda, \mu, L, \psi, I_0)$. Priors over the five parameters are given in S4 Table in S1 Text.

As a first step, we trained our network on simulations from the simple SIR-type model formulated above and then applied the network to the number of reported cases from the first 14 days of the Covid-19 pandemic in Germany. The results from this initial study are depicted in Fig 3. First, we observe that our posterior estimates are in line with those reported in a previous modeling study [8], which utilized the same data and a similar model. Second, we note that the SBC plots indicate no systematic biases in the approximate posteriors and thus suggest that the posterior samples are trustworthy (assuming no simulation gap). Finally, the posterior predictive check indicates that the model can accurately reproduce the observed data (Fig 3).

Since Outbreak-Flow is especially designed to tackle complex stochastic models whose estimation usually necessitates simulation-based approaches, we compared its performance to an

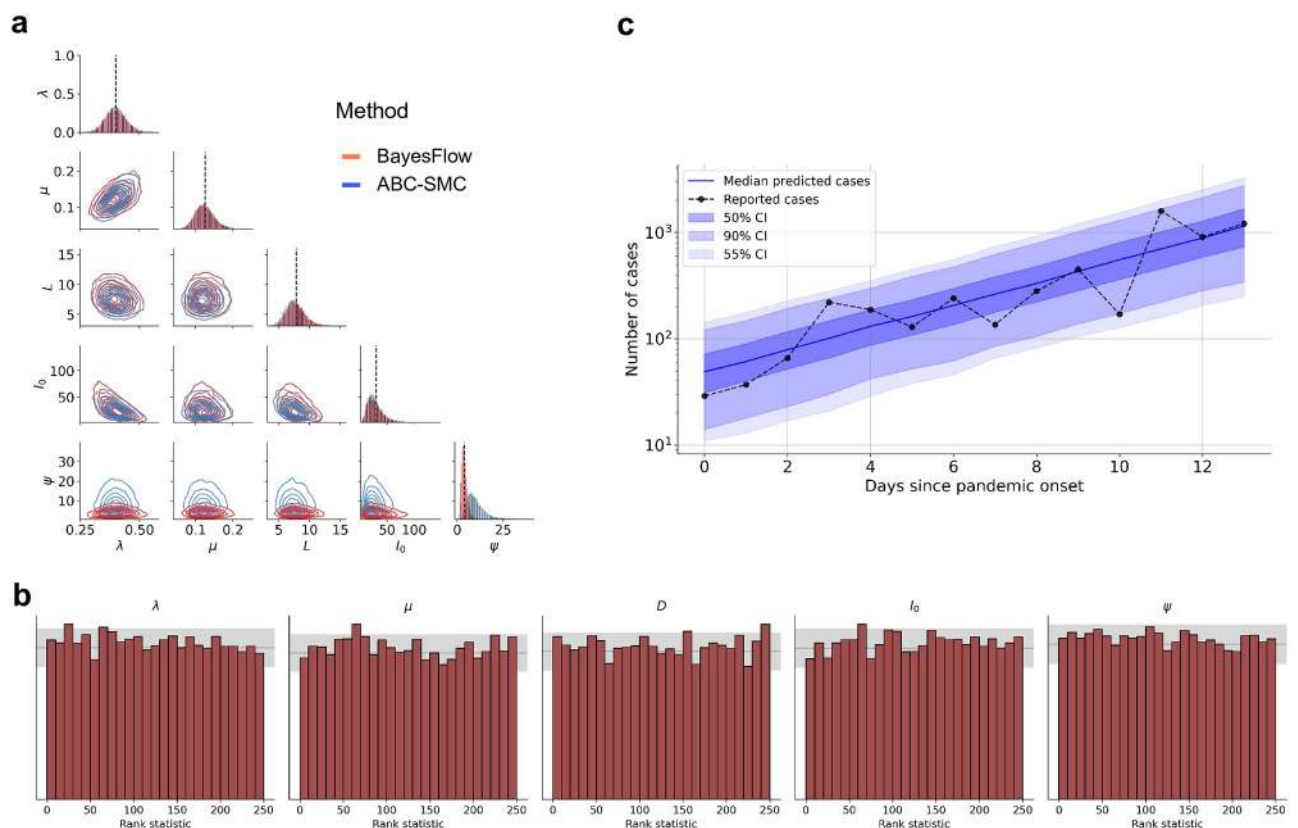


Fig 3. (a) Bivariate posteriors over the five model parameters obtained by OutbreakFlow (red) and ABC-SMC (blue) from cases reported during the first 14 days of the Covid-19 pandemic in Germany. Results of the two methods are very similar except for the dispersion parameter ψ , where OutbreakFlow achieves superior uncertainty reduction (see text for a brief discussion); (b) Model-based predictions based on the posteriors obtained by OutbreakFlow; (c) Simulation-based calibration (SBC) computed from 10000 simulated pandemic scenarios. The uniformity of the rank statistics indicates that there are no systematic biases in the approximate posteriors and implies trustworthy approximation.

<https://doi.org/10.1371/journal.pcbi.1009472.g003>

analysis performed by ABC-SMC, a popular approximate Bayesian computation algorithm based on sequential Monte Carlo sampling [31]. Our benchmark comparison reveals converging results. However, OutbreakFlow implies a much sharper marginal posterior of ψ than ABC-SMC, indicating more uncertainty reduction with regard to the dispersion parameter. Since the SBC plots indicate no overconfidence (i.e., no overdispersion) of the OutbreakFlow posteriors, it is likely that the ABC-SMC algorithm yields an underconfident (i.e., underdispersed) marginal posterior with respect to this one parameter. As for wall-clock running times, the ABC-SMC algorithm converged in 4.1 hours, whereas OutbreakFlow trained with 50,000 iterations using online learning required 4.3 hours on the same laptop machine without parallel simulations. This similar performance speaks in favor of amortized inference, as the training effort already amortizes after as few as two data sets.

To further test if OutbreakFlow is able to provide a reliable quantification of model parameters, we analyzed if the additional consideration of non-identifiable parameters within the model would affect the calibration or predictive performance of our method. To this end, we included 5 dummy variables $u_j \sim \text{Uniform}(0, 1)$ within our model that were not used for the generation of the simulated data, but were later included in the unknown parameter vector θ during training and inference.

Performing the same training and inference phase with these 5 additional dummy parameters neither hurts the calibration of OutbreakFlow, nor does it impact inference on the observed time series in a noticeable way, since the posterior estimates for the relevant parameters appear to be unaffected by the dummy parameters (see [S1 Text](#) for full results). The same is true for model-based posterior predictions, which underlines the ability of OutbreakFlow to reliably characterize parameter identifiability in case of insufficient data or over-parameterized models.

Inferring epidemiological characteristics from the early Covid-19 pandemic in Germany

After validation of the general applicability of our novel approach, we applied OutbreakFlow to data from the Covid-19 pandemic in Germany, analyzing reported cases (infected, recovered and deceased) in the time period from March 01, 2020 until June 11, 2020. These data captured the early dynamics of the emerging epidemic associated with considerable uncertainty and stochasticity with regard to the number of detected cases, as well as the effect of subsequent public interventions. For our analysis, we used an extended SEIR-type model that had been developed recently and distinguishes between detected and undetected carriers of the disease comprising a total of 34 unknown model parameters (see [Fig 2](#) and [Materials and methods](#) for a detailed description of the model) [18]. The model was trained on a time-period from March 01 until May 21 using wide prior distributions across plausible parameter ranges from previous literature [8, 22]. The remaining data (three weeks from May 22 until June 11) are then used to assess the predictive value of the model.

The observed and predicted dynamics, as well as the marginal posterior distributions of the individual model parameters are depicted in [Fig 4](#) (see [S6 Fig](#) in [S1 Text](#) for simulation-based calibration). Our model was able to recover the observed dynamics and yields good predictions for the future period, with its forecasts having well-calibrated uncertainty bounds for the newly infected, recovered, and diseased cases (see [Fig 4](#)). Despite the large number of unknown model parameters and limited data, our analysis indicated considerable reduction in uncertainty in relation to the prior knowledge for most of the model parameters ([Fig 4](#)). Standard point estimates (median, mean, MAP mode) and credibility intervals (95%-CI between

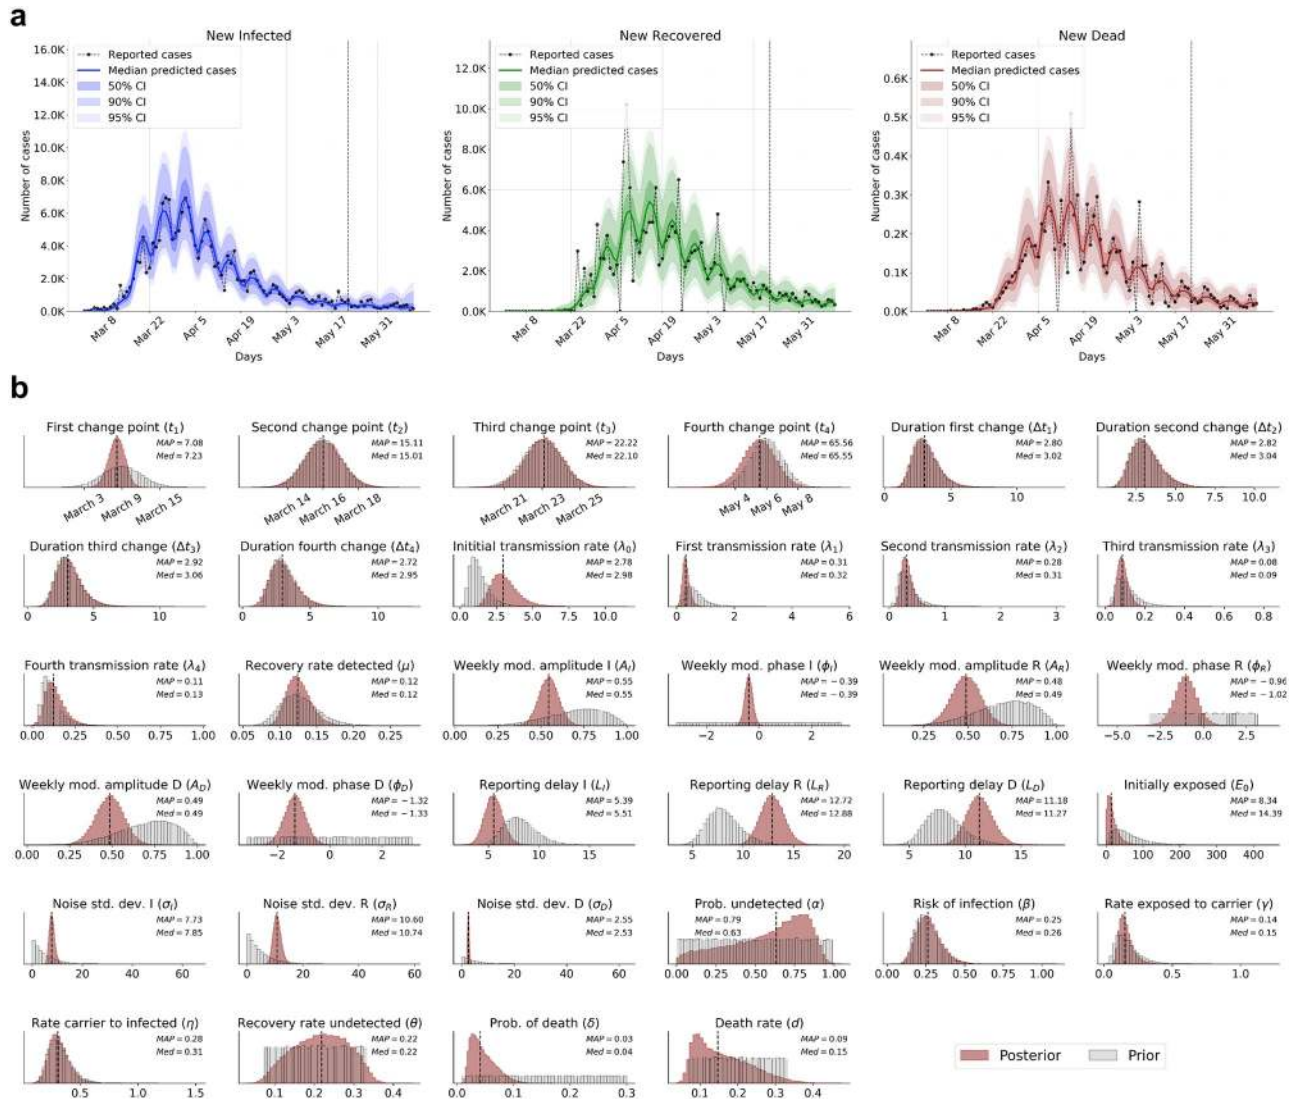


Fig 4. (a) Posterior predictions and forecasts of new cases obtained by inferring model parameters from epidemiological data available for reported infected, assumed recovered and deaths by Covid-19 for the entire Germany. Cases to the left of the vertical dashed line were used for posterior checking (model training) and cases to the right for posterior forecasts (predictions) on unseen data; (c) Marginal posteriors of all 34 model parameters inferred from data for the entire Germany alongside median and MAP summary statistics. Gray lines depict prior distributions for comparison with the posteriors. Vertical dashed lines indicate posterior medians.

<https://doi.org/10.1371/journal.pcbi.1009472.g004>

the 2.5% and 97.5% quantiles of the corresponding posterior) for all 34 model parameters are given in [Table 1](#).

Interestingly, our parameter estimates (cf. [Fig 4](#)) are consistent with previous findings about central disease parameters [8]. With the number of undetected cases being one of the most important estimates to assess epidemic-related dynamics, our network estimates a median probability of remaining undetected (parameter α) of 0.63 with the maximum a posteriori (MAP) estimate at 0.79 (95%-CI [0.07–0.91]). Notwithstanding the large uncertainty surrounding the number of undetected cases, the posterior of α is clearly far from uniform (our prior assumption), and peaks well beyond 0.5 (see [Fig 4](#)). This estimate is consistent with the results of representative seroprevalence studies in Germany [32, 33], which find that 75% of

Table 1. Posterior summaries and 95%-CIs for each model parameter inferred from data for entire Germany.

Parameter	Symbol	Median	Mean	MAP	95%-CI
Onsets of changes	t_1	March 8	March 8	March 8	[March 6–11]
		(Day 7.23)	(Day 7.20)	(Day 7.08)	[Day 4.64–9.59]
	t_2	March 16	March 16	March 16	[March 14–18]
		(Day 15.01)	(Day 15.01)	(Day 15.11)	[Day 12.99–17.05]
	t_3	March 23	March 23	March 23	[March 21–25]
		(Day 22.10)	(Day 22.10)	(Day 22.22)	[Day 20.23–24.00]
	t_4	May 6	May 6	May 6	[May 4–May 8]
		(Day 65.55)	(Day 65.53)	(Day 65.56)	[Day 63.48–67.54]
Duration of changes [days]	Δt_1	3.02	3.14	2.80	[1.61–5.34]
	Δt_2	3.04	3.16	2.82	[1.65–5.34]
	Δt_3	3.06	3.18	2.92	[1.63–5.41]
	Δt_4	2.95	3.08	2.72	[1.45–5.43]
Transmission rates	λ_0	2.98	3.12	2.78	[1.70–5.31]
	λ_1	0.32	0.34	0.31	[0.13–0.63]
	λ_2	0.31	0.33	0.28	[0.16–0.58]
	λ_3	0.09	0.09	0.08	[0.05–0.15]
	λ_4	0.13	0.14	0.11	[0.05–0.28]
Reporting delays	L_I	5.51	5.54	5.39	[3.87–7.35]
	L_R	12.88	12.91	12.72	[10.73–15.21]
	L_D	11.27	11.31	11.18	[9.15–13.65]
Weekly modulation amplitudes	A_I	0.55	0.55	0.55	[0.43–0.66]
	A_R	0.49	0.49	0.48	[0.33–0.65]
	A_D	0.49	0.49	0.49	[0.32–0.64]
Weekly modulation phases	φ_I	-0.39	-0.39	-0.39	[-0.69–-0.09]
	φ_R	-1.02	-1.02	-0.96	[-2.36–-0.33]
	φ_D	-1.33	-1.33	-1.32	[-2.13–-0.55]
Reporting noise scales	σ_I	7.85	7.92	7.73	[5.96–10.31]
	σ_R	10.74	10.85	10.60	[8.36–13.88]
	σ_D	2.55	2.54	2.53	[1.95–3.23]
Number of initially exposed	E_0	14.39	18.72	8.34	[1.20–61.37]
Risk of infection from I	β	0.26	0.27	0.25	[0.14–0.45]
Rate $E \rightarrow C$ [1/days]	γ	0.15	0.17	0.14	[0.09–0.30]
Rate $C \rightarrow I$ [1/days]	η	0.31	0.32	0.28	[0.18–0.52]
Rate $I \rightarrow R$ [1/days]	μ	0.12	0.13	0.12	[0.10–0.16]
Rate $C \rightarrow R$ [1/days]	θ	0.22	0.22	0.22	[0.09–0.33]
Rate $I \rightarrow D$ [1/days]	d	0.15	0.16	0.09	[0.07–0.32]
Probability of $C \rightarrow R$	α	0.63	0.58	0.79	[0.07–0.91]
Probability of $I \rightarrow D$	δ	0.04	0.05	0.03	[0.02–0.10]

<https://doi.org/10.1371/journal.pcbi.1009472.t001>

the sero-positive individuals at the end of the first wave had not been diagnosed with the disease before. An even higher fraction of undetected cases around 80% was reported by seroprevalence studies focusing on the hotspot regions of Gangelte, Kupferzell, and Tischenreuth [34–36]. Together with our estimated case fatality rate (parameter δ) of 4.1% (median) resp. 3% (MAP), this results in an infection fatality rate of about 0.63% (MAP estimates) or 1.5% (median estimates).

Additionally, our informative priors for the parameters η and γ are not updated by the observed data. Accordingly, around 3.2 days will typically pass before the infection is detected

(95%-CI [1.92–5.55], $1/\eta$). This estimate is in line with results from [37] (around 4 days) and the World Health Organization [38] (5–6 days). Combined with the latent period (i.e., the time in compartment E : median 6.7 days, 95%-CI [3.33–11.1], $1/\gamma$), this leads to a median incubation period until symptom onset or a positive test outcome of around 9.9 days, as originally implied by our priors.

If the disease remains undiagnosed throughout (asymptomatic, weakly symptomatic etc.), recovery takes a median number of 4.6 days (95%-CI [2.99–11.11], $1/\theta$), a time period in which pre-symptomatic and undiagnosed individuals can be responsible for a considerable fraction of infections. Furthermore, if we assume that most infections occur near the end of the carrier stage, that is, after 2.5 days in compartment C , we arrive at a generation time of around 9 days. In conjunction with the delay for reporting infections I of 5.5 days (parameter L_I), this is consistent with the generally acknowledged fact that the success of intervention measures only becomes apparent after around two weeks.

For diagnosed individuals, the median recovery period is estimated at 8.1 days (95%-CI [6.13–10.20], $1/\mu$). Thus, manifestly ill cases have a more severe disease progression than undiagnosed individuals and typically require $1/\eta + 1/\mu = 3.2 + 8.1 = 11.3$ days until recovery [39]. The time between diagnosis and death (6.7 days, 95%-CI [3.12–14.3], parameter $1/d$) is shorter than in clinical reports, with parameter identification possibly impaired by the estimated reporting delay for disease-associated deaths D of 11.3 days (parameter L_D), which is probably much longer than in reality. From the available time series alone, the model is not able to distinguish a long critical phase with short reporting delay from rapid death with long reporting delay. Nevertheless, it is remarkable how much information about 34 free parameters our networks can extract from seeing only about 70 time steps of real data (see Fig 4).

Finally, our results corroborate the timing of intervention measures and the gradual reduction in transmission rate as observed in [8]. According to our estimates, the lifting of measures around May 6 would have led to an approximately 40% increase in the transmission rate, as assumed by our prior. However, since the spreading rate at t_4 has already decreased to a median of 0.09 (95%-CI: [0.05–0.15]), the increase to a median of 0.13 (95%-CI [0.05–0.28]) does not lead to an exponential growth of infections.

Testing the robustness of model analysis and inferred dynamics

To test the validity of our results, we performed a series of *ablation studies* which independently reduce either the architecture of our neural network or the considered epidemiological model. Thereby, we are able to test (i) the importance of specific technical components for parameter inference, as well as (ii) the importance of specific aspects in the model structure when trying to explain the observed dynamics. Accordingly, the implemented changes comprise: 1) removing the convolutional filtering network; 2) removing the recurrent summary network; 3) removing the observation model; 4) removing the intervention model; 5) removing the carrier (C) compartment from the latent disease model. For each of these ablation studies, we used the same number of simulations and training settings as before. We then compared the predictive performance, calibration, and reliability in parameter estimation of each of the modified analyzes to our previous analysis.

The results from all ablation studies are available in [S1 Text](#). Indeed, our experiments indicate that all network and model components are crucial for the performance of OutbreakFlow and for the fidelity of model predictions. Interestingly, no ablation scenario leads to dramatic miscalibration across all marginal posteriors. However, predictive performance markedly deteriorates in all ablation studies. For instance, removing either the intervention or the observation model leads to the rejection of all time series drawn from the posterior predictive due to

divergences. Further, removing the latent carrier compartment leads to poor fits and notable underestimation of the reported cases. On the other hand, removing either the filtering or the summary network does not lead to dramatic misfits but prevents the architecture from fully capturing the structure of the observed time series (e.g., weekly modulation is poorly captured, see [S1 Text](#)).

Predicting the epidemiological dynamics within the individual German federal states

In the previous section, we demonstrated the ability of our method to recover observed epidemiological dynamics and to infer reliable parameter estimates that determine disease characteristics based on data covering the early Covid-19 epidemic in whole Germany. Arguably, the importance of uncertainty concerning the number of reported cases, as well as stochastic effects increase with lower case numbers. To this end, we also applied our method to each German federal state separately, as individual states are characterized by different population sizes, as well as different onsets and progression of the epidemic. Because data for recovered cases per federal state was lacking (in contrast to recovered cases for the entire Germany), the network was trained solely on the simulated reported infected cases and deaths. Note that we only trained a *single* OutbreakFlow, which we then applied unchanged to each German federal state. In this way, the training effort amortized over the repeated applications of the same neural estimator (see S13 Fig in [S1 Text](#) for simulation-based calibration of the trained network).

Posterior predictions and forecasts for cumulative infections (derived from estimated new cases, see S14 Fig in [S1 Text](#)) in each federal state are depicted in [Fig 5](#) (see also S15 Fig in [S1 Text](#) for predictions of cumulative deaths). As for Germany as a whole, we observe that median predictions follow very closely the reported cumulative number of cases across all federal states. Furthermore, the reported cases are very well represented by the uncertainty bounds derived from the parameter posteriors, with prediction uncertainty increasing towards the future (i.e., predictions after the dotted vertical lines in [Fig 5](#)). However, median predictions can become unreliable when only a few cases are available for model training (see predictions for cumulative deaths for the state Mecklenburg-Western Pomerania, S15 Fig in [S1 Text](#)). Therefore, well-calibrated uncertainty estimates are particularly important and need to be taken into account when reporting point predictions.

The posterior distributions of individual model parameters for each of the 16 German federal states are depicted in [S1 Text](#). For comparison of individual estimates between different states, we focused on four latent parameters that are essential for assessing early epidemical dynamics: (i) the probability of infecteds to remain undetected (α), (ii) the recovery rate of undetected (θ), (iii) the number of initially exposed individuals (E_0), and (iv) the initial transmission rate (λ_0).

We observe that posterior estimates of α across states tend to peak well above 0.5 (see [Fig 6a](#)), suggesting once again that many infections have remained undetected/undiagnosed during the initial months of the Covid-19 pandemic in Germany. Furthermore, we observe that some states have smaller probabilities of undiagnosed infections, especially Bavaria, Berlin, and North Rhine-Westphalia (possibly indicating a more successful testing policy), although these estimates are associated with substantial uncertainty. Interestingly, the α posteriors of each state are considerably sharper than the posterior of α for the entire Germany (see [S1 Text](#) for full marginal posterior plots for each state). This may be a consequence of the high variability of α across states, but could also be an estimation artifact of the simplified model (only two epidemiological time series, new infections and deaths, are available for each state).

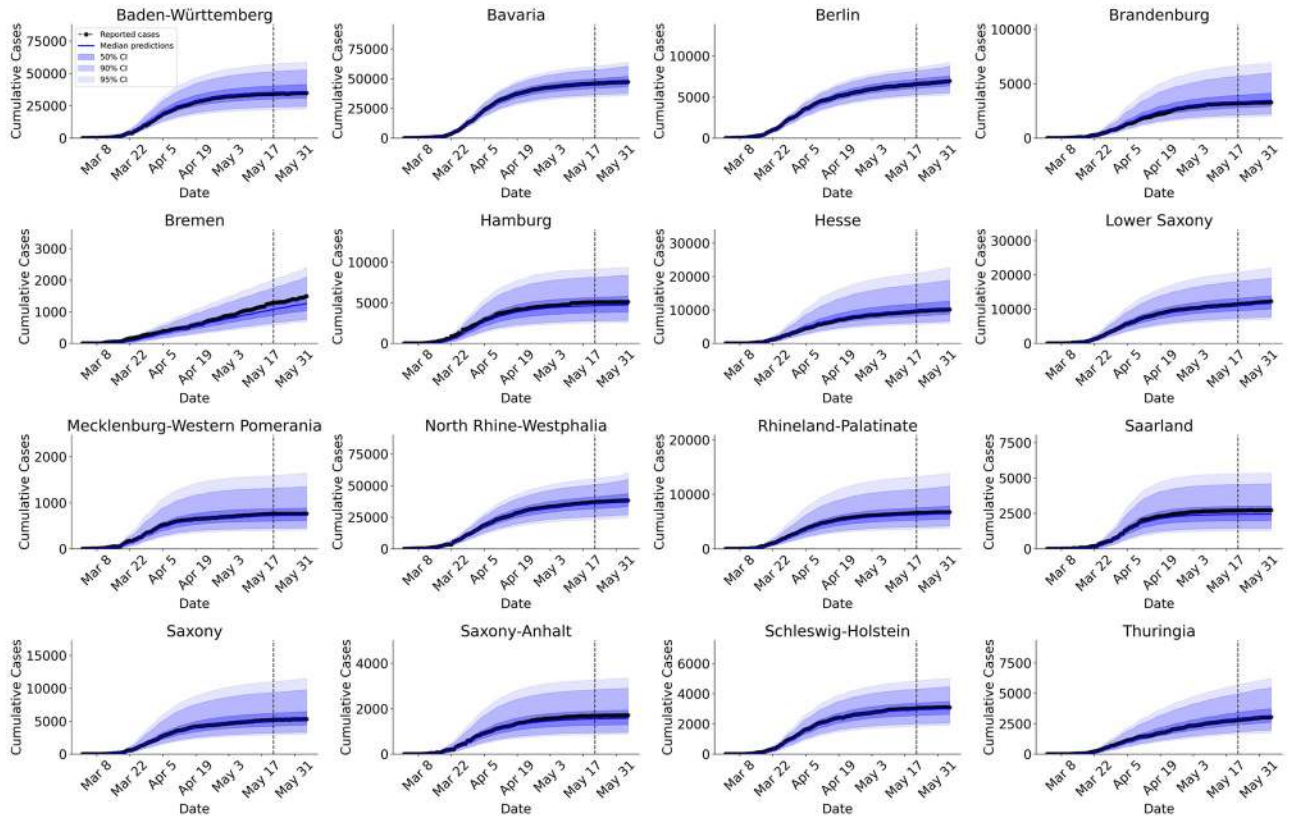


Fig 5. Model predictions of cumulative Covid-19 cases (derived from estimated new cases) for each German federal state. Cases to the left of the vertical dashed line (8 weeks) were used for model fitting and posterior checking and cases to the right (3 weeks) for forecasts on new data. We observe that median model predictions closely match both past and future reported cases for each German federal state. Most importantly, the reported cases always lie within the estimated CIs, which vary across the federal states.

<https://doi.org/10.1371/journal.pcbi.1009472.g005>

Less variability between states is observed in the estimates for the recovery rate of undetected, θ , suggesting an overall fast recovery of undiagnosed individuals (Fig 6b). However, some differences between states are evident in the estimates of initially exposed individuals, E_0 (cf. Fig 6c), with the states of Bavaria, Berlin, and North Rhine-Westphalia having significantly more exposed individuals at onset than other states, which may correspond to observed outbreaks in these states after skiing holidays and festivals. Finally, Fig 6c depicts a comparison of initial transmission rates λ_0 between states. We observe that estimates vary around a median value of 2.27 across states, with the state Mecklenburg-Western Pomerania having the lowest and the state Baden-Württemberg having the highest median transmission rate at onset.

Thus, our analysis is able to reveal differences in the epidemiological dynamics between individual German federal states and proves reliable for different population sizes.

Discussion

In this work, we presented a novel simulation-based Bayesian inference framework for complex epidemiological models. We directly demonstrated the utility of our method by applying it to publicly available data on the early reported infected, recovered, and deceased individuals during the first phase of the Covid-19 pandemic in Germany.

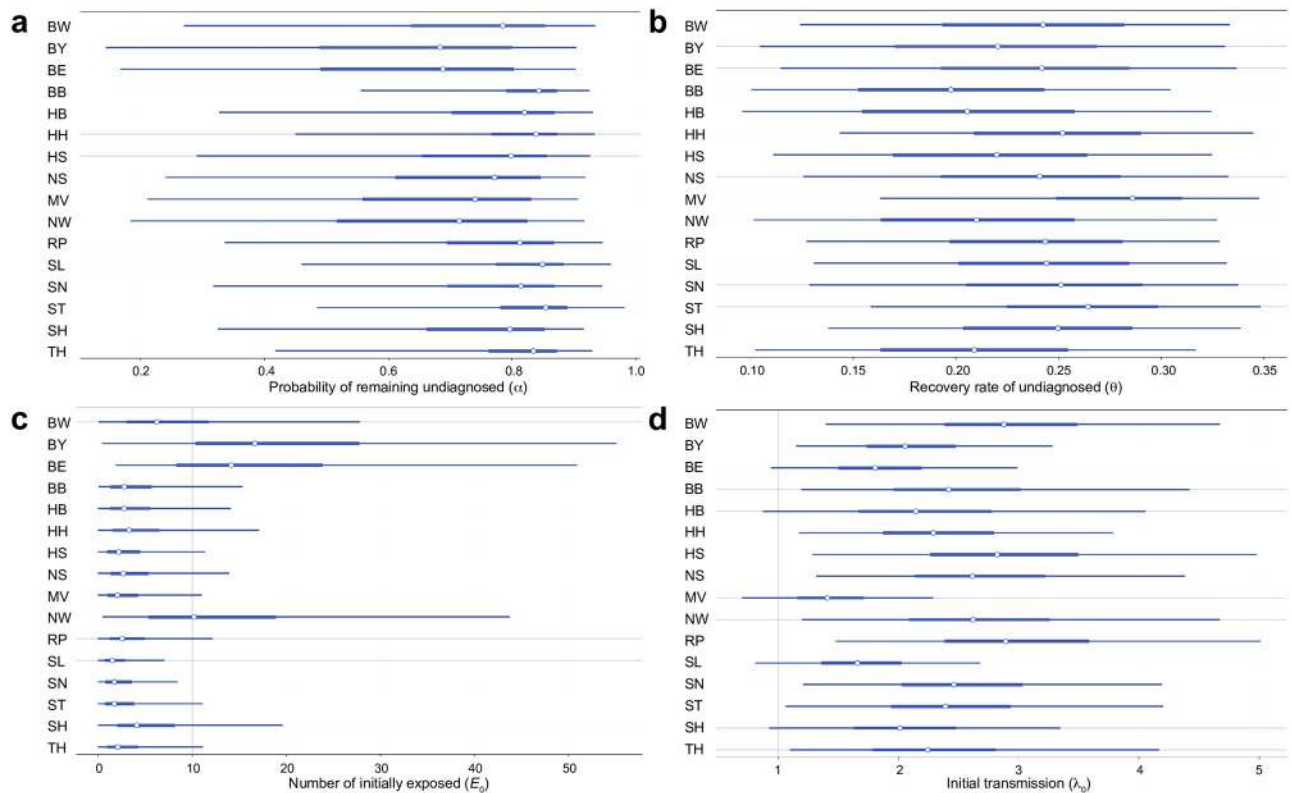


Fig 6. (a) Forest plot depicting 95% credibility intervals for the probability of remaining undetected (α) obtained by amortized inference on data from all German federal states. Thin lines depict highest density intervals (HDI), thick lines depict posterior quartiles, and white points depict the corresponding medians of the estimated posteriors; (b-d) Corresponding plots for (b) the recovery rate of undetected (θ), (c) the number of initially exposed (E_0), and (d) the initial transmission rate (λ_0).

<https://doi.org/10.1371/journal.pcbi.1009472.g006>

Given the general uncertainty in reported numbers for emerging infectious diseases, estimation methods need to account for this uncertainty when providing parameter estimates and be able to efficiently incorporate incoming data. Our method works reliably well, providing well calibrated uncertainty bounds for individual parameter estimates even in case of small sample sizes and a limited amount of observations. Given the limited amount of data available for model calibration (i.e., number of reported infections, defined recoveries and deaths) and the uncertainty concerning these values, our analysis is able to reduce the uncertainty about quantities that could not be obtained during the time course of the epidemics.

Our estimates suggest that a large fraction of the infected individuals (60–80%) might have gone undetected through the course of the early Covid-19 outbreak in Germany. This finding has been confirmed by subsequent seroprevalence studies from Germany [32–36] and other countries [40–43] and is also in line with estimates on reporting rates [44]. However, our posteriors also suggest that there is non-negligible uncertainty surrounding this estimate when derived in a purely model-based manner. Moreover, different summary statistics (e.g., means, medians, MAPs) derived from non-symmetric posteriors offer slightly different conclusions. This observation highlights the need to consider the full posteriors and corresponding credibility intervals when aiming to draw substantive conclusions about epidemiological parameters and possible forecasts for the progression of the epidemic or the effect of specific public health interventions. When interpreting the results of parameter estimates, one should also be

aware that mechanistic models like the one used here only describe the average behavior of entire compartments. Accordingly, the given confidence intervals quantify our uncertainty about the inferred parameter averages and can not serve as a measure for the variability between individual cases.

Our approach has two key advantages over standard Bayesian and likelihood-based point estimation methods. First, it can flexibly deal with arbitrarily complex models and data structures, requiring no closed-form likelihoods or *ad hoc* distributional restrictions regarding the shape of the joint prior or posterior. As standard SIR-models based on (stochastic) ordinary differential equations generally provide a coarse-grained view on the epidemic dynamics [1], more complex models accounting for heterogeneous social interactions, age-dependent effects, and/or spatial and temporal heterogeneity become more and more important to predict the progression of an epidemic or guide intervention measures [8, 45–49]. Such agent-based and stochastic models can be easily incorporated within our neural Bayesian framework.

As a second advantage, the *amortized inference* property of our method, that is, training the network only *once* on simulated data, allows efficient posterior sampling and simultaneous application to multiple data sets. In addition, it allows for efficient online-learning and validation (i.e., the continuous integration of upcoming data), once the networks have been trained with sufficient amounts of simulated data. These advantages are important, since they enable researchers to concentrate on formulating, testing, and validating complex model systems without worrying about estimation efficiency or analytical tractability.

Future developments will include Bayesian model comparison, multilevel modeling with hierarchical priors and a systematic comparison between different neural inference architectures. Hierarchical modeling allows us to better distinguish between disease-specific and region-specific parameters by representing the natural cluster structure of epidemiological data. In addition, model comparison is an especially important research avenue to compare different possible disease transmission and progression schemes that could explain the observed dynamics. Currently, in order to compare multiple competing dynamic models, separate neural networks would need to be trained and stored—one network corresponding to each model. Even though such a training can be carried out in parallel, it would be much more efficient to embed all competing models within a single neural architecture, which can perform both prior predictive and posterior predictive model comparison. This idea has already been explored in [50] and future research should investigate its utility for comparing complex epidemiological models.

In addition, tackling non-stationary dynamic processes is an important open problem of our framework that we plan to address in future work. This limitation is one reason why we focus on the early outbreak dynamics, where the assumption of stationarity appears plausible (e.g., absence of major virus mutations, largely fixed testing policy and implied fraction of undetected infections, no significant progress on treatment options and subsequent recovery times), with the exception of government interventions and behavioral changes, which we explicitly model. Moreover, our Bayesian treatment ensures that deviations from stationarity and shortcomings of the reported data do not result in catastrophic failure of our method, but are reflected in wider uncertainty regions than would otherwise have been achievable.

When applying OutbreakFlow, one should be aware of three potential error sources that can distort the outcomes and interpretations of amortized Bayesian workflows. First, model misspecification and data contamination can result in a *simulation gap*. A simulation gap occurs when the model cannot represent the actual disease dynamics or when data collection is biased or contaminated in ways not accounted for by the model. We addressed these issues by suitable model extensions, which are motivated by theoretical considerations and ablation studies. Remaining misspecifications can often be detected via standard Bayesian model

checking methods, for instance, insufficient posterior predictive accuracy, divergent re-simulations, or very low posterior probability under the prior. However, theoretical guarantees on upper bounds for the residual errors remain an important open problem of our approach.

The second source is the *Monte Carlo error* introduced by approximating the expectation in Eq 3 with only a finite number of simulations. It is also referred to as approximation error and widely acknowledged throughout all Monte Carlo methods. This error can be mitigated relatively easily in our online learning setting, because we can generate a potentially endless stream of synthetic training data until the continuously monitored prediction accuracy is satisfactory. In this respect, simulation-based inference is better positioned to fully utilize the capacity of deep neural networks than traditional supervised learning methods, which rely on a limited supply of annotated real data.

The third source of error is an *amortization gap*, which refers to potential deficiencies of the network for atypical real dynamics. Atypical dynamics are, by definition, underrepresented in any training set, and an amortized inference scheme may be less accurate in such cases. An amortization gap can be detected via the probabilistic calibration methods advocated in this work, that is, simulation-based calibration. When the amortization gap is too large, resorting to non-amortized methods such as [23, 24] might be a viable option. Alternatively, the accuracy of an OutbreakFlow architecture for more rare situations can be improved by increasing the expressiveness of the filtering, summary, and inference networks (e.g., more layers and more units, along with more training iterations). The choice of optimal neural network architectures and training algorithms is a topic of ongoing research in the deep learning community and an important target for future work on simulation-based inference.

In summary, our OutbreakFlow architecture provides a general inference framework for complex epidemiological scenarios and enables efficient simulation-based inference for key epidemiological parameters. We therefore believe that our proposed architecture can facilitate uncertainty-aware inference with complex and realistic epidemiological models, especially during the early phase of epidemic outbreaks when information is scarce and data reliability low. Rapid trustworthy inference can timely reveal crucial dynamic aspects of a spreading disease and inform effective public health interventions.

Supporting information

S1 Text. Online supporting information. S1 Algorithm. Simulation-based offline learning. **S2 Algorithm.** Simulation-based online learning. **S3 Algorithm.** Simulation-based hybrid learning. **S1 Table.** Disease model priors. **S2 Table.** Intervention model priors. **S3 Table.** Observation model priors. **S4 Table.** Validation SIR model priors. **S5 Table.** Periodic SIR model priors. **S1 Fig.** Model-based predictions with non-identifiable parameters. **S2 Fig.** Calibration of OutbreakFlow trained with additional non-identifiable parameters. **S3 Fig.** Bivariate posteriors from real data. **S4 Fig.** Calibration of OutbreakFlow trained on the periodic SIR model. **S5 Fig.** Model-based predictions on 16 simulated time-series from the periodic SIR model. **S6 Fig.** Calibration of OutbreakFlow trained for inference on entire Germany. **S7 Fig.** Calibration of OutbreakFlow without a convolutional filtering network. **S8 Fig.** Model-based predictions using OutbreakFlow without a convolutional filtering network. **S9 Fig.** Calibration of OutbreakFlow without a recurrent summary network. **S10 Fig.** Model-based predictions using OutbreakFlow without a recurrent summary network. **S11 Fig.** Calibration of OutbreakFlow without a latent carrier compartment in the epidemiological model. **S12 Fig.** Model-based predictions using an epidemiological model without a carrier compartment. **S13 Fig.** Calibration of OutbreakFlow trained for inference on the German federal states. **S14 Fig.** Model predictions of new cases for each German federal state. **S15 Fig.** Model predictions of

cumulative Covid-19 deaths (derived from predicted new cases) for each German federal state. **S16 Fig.** Marginal parameter posteriors from data available for the German federal state Baden-Württemberg. **S17 Fig.** Marginal parameter posteriors from data available for the German federal state Bavaria. **S18 Fig.** Marginal parameter posteriors from data available for the German federal state Berlin. **S19 Fig.** Marginal parameter posteriors from data available for the German federal state Brandenburg. **S20 Fig.** Marginal parameter posteriors from data available for the German federal state Bremen. **S21 Fig.** Marginal parameter posteriors from data available for the German federal state Hamburg. **S22 Fig.** Marginal parameter posteriors from data available for the German federal state Hesse. **S23 Fig.** Marginal parameter posteriors from data available for the German federal state Saxony. **S24 Fig.** Marginal parameter posteriors from data available for the German federal state Mecklenburg-Western Pomerania. **S25 Fig.** Marginal parameter posteriors from data available for the German federal state North Rhine-Westphalia. **S26 Fig.** Marginal parameter posteriors from data available for the German federal state Rhineland-Palatinate. **S27 Fig.** Marginal parameter posteriors from data available for the German federal state Saarland. **S28 Fig.** Marginal parameter posteriors from data available for the German federal state Saxony-Anhalt. **S29 Fig.** Marginal parameter posteriors from data available for the German federal state Saxony. **S30 Fig.** Marginal parameter posteriors from data available for the German federal state Schleswig-Holstein. **S31 Fig.** Marginal parameter posteriors from data available for the German federal state Thuringia. (PDF)

Author Contributions

Conceptualization: Stefan T. Radev, Frederik Graw, Ullrich Köthe.

Data curation: Stefan T. Radev, Simiao Chen, Nico T. Mutters, Vanessa M. Eichel, Ullrich Köthe.

Formal analysis: Stefan T. Radev, Frederik Graw, Ullrich Köthe.

Funding acquisition: Stefan T. Radev.

Investigation: Stefan T. Radev, Simiao Chen, Ullrich Köthe.

Methodology: Stefan T. Radev, Frederik Graw, Ullrich Köthe.

Project administration: Ullrich Köthe.

Resources: Stefan T. Radev, Nico T. Mutters, Vanessa M. Eichel.

Software: Stefan T. Radev.

Supervision: Frederik Graw, Till Bärnighausen, Ullrich Köthe.

Validation: Stefan T. Radev, Frederik Graw, Ullrich Köthe.

Visualization: Stefan T. Radev, Frederik Graw, Ullrich Köthe.

Writing – original draft: Stefan T. Radev, Frederik Graw, Simiao Chen, Nico T. Mutters, Vanessa M. Eichel, Till Bärnighausen, Ullrich Köthe.

Writing – review & editing: Stefan T. Radev, Frederik Graw, Simiao Chen, Nico T. Mutters, Vanessa M. Eichel, Till Bärnighausen, Ullrich Köthe.

References

1. Keeling MJ, Rohani P. Modeling infectious diseases in humans and animals. Princeton University Press; 2011.

2. Wu JT, Leung K, Leung GM. Nowcasting and forecasting the potential domestic and international spread of the 2019-nCoV outbreak originating in Wuhan, China: a modelling study. *The Lancet*. 2020; 395(10225):689–697. [https://doi.org/10.1016/S0140-6736\(20\)30260-9](https://doi.org/10.1016/S0140-6736(20)30260-9) PMID: [32014114](https://pubmed.ncbi.nlm.nih.gov/32014114/)
3. Chinazzi M, Davis JT, Ajelli M, Gioannini C, Litvinova M, Merler S, et al. The effect of travel restrictions on the spread of the 2019 novel coronavirus (COVID-19) outbreak. *Science*. 2020; 368(6489):395–400. <https://doi.org/10.1126/science.aba9757> PMID: [32144116](https://pubmed.ncbi.nlm.nih.gov/32144116/)
4. Kucharski AJ, Russell TW, Diamond C, Liu Y, Edmunds J, Funk S, et al. Early dynamics of transmission and control of COVID-19: a mathematical modelling study. *The lancet infectious diseases*. 2020; 20(5):553–558. [https://doi.org/10.1016/S1473-3099\(20\)30144-4](https://doi.org/10.1016/S1473-3099(20)30144-4)
5. Scire J, Nadeau SA, Vaughan T, Gavin B, Fuchs S, Sommer J, et al. Reproductive number of the COVID-19 epidemic in Switzerland with a focus on the Cantons of Basel-Stadt and Basel-Landschaft. *Swiss medical weekly*. 2020; 150(19-20):w20271. PMID: [32365217](https://pubmed.ncbi.nlm.nih.gov/32365217/)
6. Hauser A, Counotte MJ, Margossian CC, Konstantinoudis G, Low N, Althaus CL, et al. Estimation of SARS-CoV-2 mortality during the early stages of an epidemic: A modeling study in Hubei, China, and six regions in Europe. *PLoS medicine*. 2020; 17(7):e1003189. <https://doi.org/10.1371/journal.pmed.1003189> PMID: [32722715](https://pubmed.ncbi.nlm.nih.gov/32722715/)
7. Tian H, Liu Y, Li Y, Wu CH, Chen B, Kraemer MU, et al. An investigation of transmission control measures during the first 50 days of the COVID-19 epidemic in China. *Science*. 2020; 368(6491):638–642. <https://doi.org/10.1126/science.abb6105> PMID: [32234804](https://pubmed.ncbi.nlm.nih.gov/32234804/)
8. Dehning J, Zierenberg J, Spitzner FP, Wibral M, Neto JP, Wilczek M, et al. Inferring change points in the spread of COVID-19 reveals the effectiveness of interventions. *Science*. 2020; 369(6500). <https://doi.org/10.1126/science.abb9789> PMID: [32414780](https://pubmed.ncbi.nlm.nih.gov/32414780/)
9. Li R, Pei S, Chen B, Song Y, Zhang T, Yang W, et al. Substantial undocumented infection facilitates the rapid dissemination of novel coronavirus (SARS-CoV-2). *Science*. 2020; 368(6490):489–493. <https://doi.org/10.1126/science.abb3221> PMID: [32179701](https://pubmed.ncbi.nlm.nih.gov/32179701/)
10. Kong CY, McMahon PM, Gazelle GS. Calibration of disease simulation model using an engineering approach. *Value in health*. 2009; 12(4):521–529. <https://doi.org/10.1111/j.1524-4733.2008.00484.x> PMID: [19900254](https://pubmed.ncbi.nlm.nih.gov/19900254/)
11. Martcheva M. *An introduction to mathematical epidemiology*. vol. 61. Springer; 2015.
12. Gelman A, Carlin JB, Stern HS, Dunson DB, Vehtari A, Rubin DB. *Bayesian data analysis*. CRC press; 2013.
13. Peak CM, Kahn R, Grad YH, Childs LM, Li R, Lipsitch M, et al. Individual quarantine versus active monitoring of contacts for the mitigation of COVID-19: a modelling study. *The Lancet Infectious Diseases*. 2020; 20(9):1025–1033. [https://doi.org/10.1016/S1473-3099\(20\)30361-3](https://doi.org/10.1016/S1473-3099(20)30361-3) PMID: [32445710](https://pubmed.ncbi.nlm.nih.gov/32445710/)
14. Minter A, Retkute R. Approximate Bayesian Computation for infectious disease modelling. *Epidemics*. 2019; 29:100368. <https://doi.org/10.1016/j.epidem.2019.100368>
15. Menzies NA, Soeteman DI, Pandya A, Kim JJ. Bayesian methods for calibrating health policy models: a tutorial. *Pharmacoeconomics*. 2017; 35(6):613–624. <https://doi.org/10.1007/s40273-017-0494-4> PMID: [28247184](https://pubmed.ncbi.nlm.nih.gov/28247184/)
16. Cranmer K, Brehmer J, Louppe G. The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences*. 2020; 117(48):30055–30062. <https://doi.org/10.1073/pnas.1912789117> PMID: [32471948](https://pubmed.ncbi.nlm.nih.gov/32471948/)
17. Radev ST, Mertens UK, Voss A, Ardizzone L, Köthe U. BayesFlow: Learning complex stochastic models with invertible neural networks. *IEEE Transactions on Neural Networks and Learning Systems*. 2020;. <https://doi.org/10.1109/TNNLS.2020.3042395> PMID: [33338021](https://pubmed.ncbi.nlm.nih.gov/33338021/)
18. Khailaie S, Mitra T, Bandyopadhyay A, Schips M, Mascheroni P, Vanella P, et al. Estimate of the development of the epidemic reproduction number R_t from Coronavirus SARS-CoV-2 case data and implications for political measures based on prognostics. *medRxiv*. 2020;.
19. Szegedy C, Liu W, Jia Y, Sermanet P, Reed S, Anguelov D, et al. Going deeper with convolutions. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*; 2015. p. 1–9.
20. Gers FA, Schmidhuber J, Cummins F. Learning to forget: Continual prediction with LSTM. *Neural computation*. 2000; 12(10):2451–2471. <https://doi.org/10.1162/089976600300015015> PMID: [11032042](https://pubmed.ncbi.nlm.nih.gov/11032042/)
21. Van Erven T, Harremoës P. Rényi divergence and Kullback-Leibler divergence. *IEEE Transactions on Information Theory*. 2014; 60(7):3797–3820. <https://doi.org/10.1109/TIT.2014.2320500>
22. Tang B, Wang X, Li Q, Bragazzi NL, Tang S, Xiao Y, et al. Estimation of the transmission risk of the 2019-nCoV and its implication for public health interventions. *Journal of Clinical Medicine*. 2020; 9(2):462. <https://doi.org/10.3390/jcm9020462> PMID: [32046137](https://pubmed.ncbi.nlm.nih.gov/32046137/)

23. Papamakarios G, Sterratt D, Murray I. Sequential neural likelihood: Fast likelihood-free inference with autoregressive flows. In: The 22nd International Conference on Artificial Intelligence and Statistics. PMLR; 2019. p. 837–848.
24. Greenberg D, Nonnenmacher M, Macke J. Automatic posterior transformation for likelihood-free inference. In: International Conference on Machine Learning. PMLR; 2019. p. 2404–2414.
25. Csilléry K, Blum MG, Gaggiotti OE, François O. Approximate Bayesian computation (ABC) in practice. *Trends in Ecology & Evolution*. 2010; 25(7):410–418. <https://doi.org/10.1016/j.tree.2010.04.001> PMID: 20488578
26. Sunnåker M, Busetto AG, Numminen E, Corander J, Foll M, Dessimoz C. Approximate bayesian computation. *PLoS computational biology*. 2013; 9(1):e1002803. <https://doi.org/10.1371/journal.pcbi.1002803> PMID: 23341757
27. Raynal L, Marin JM, Pudlo P, Ribatet M, Robert CP, Estoup A. ABC random forests for Bayesian parameter inference. *Bioinformatics*. 2018; 35(10):1720–1728. <https://doi.org/10.1093/bioinformatics/bty867>
28. Radev ST, Mertens UK, Voss A, Köthe U. Towards end-to-end likelihood-free inference with convolutional neural networks. *British Journal of Mathematical and Statistical Psychology*. 2020; 73(1):23–43. <https://doi.org/10.1111/bmsp.12159> PMID: 30793299
29. Talts S, Betancourt M, Simpson D, Vehtari A, Gelman A. Validating Bayesian inference algorithms with simulation-based calibration. *arXiv:180406788*. 2018;.
30. Blumberg S, Funk S, Pulliam JR. Detecting differential transmissibilities that affect the size of self-limited outbreaks. *PLoS Pathog*. 2014; 10(10):e1004452. <https://doi.org/10.1371/journal.ppat.1004452> PMID: 25356657
31. Klinger E, Rickert D, Hasenauer J. pyABC: distributed, likelihood-free inference. *Bioinformatics*. 2018; 34(20):3591–3593. <https://doi.org/10.1093/bioinformatics/bty361> PMID: 29762723
32. Harries M, Gornyk D, Glöckner S, Strengert M, Kerrinnes T, Bojara G, et al. SARS-CoV-2 seroprevalence in Germany—a population based sequential study in five regions. *medRxiv*. 2021;.
33. Pritsch M, Radon K, Bakuli A, Le Gleut R, Olbrich L, Noller G, et al. Prevalence and risk factors of infection in the representative COVID-19 cohort Munich. *International Journal of Environmental Research and Public Health*. 2021; 18(7):3572. <https://doi.org/10.3390/ijerph18073572> PMID: 33808249
34. Streeck H, Schulte B, Kümmerer BM, Richter E, Höller T, Fuhrmann C, et al. Infection fatality rate of SARS-CoV2 in a super-spreading event in Germany. *Nature communications*. 2020; 11(1):1–12. <https://doi.org/10.1038/s41467-020-19509-y>
35. Santos-Hövenner C, Neuhauser HK, Rosario AS, Busch M, Schlaud M, Hoffmann R, et al. Serology-and PCR-based cumulative incidence of SARS-CoV-2 infection in adults in a successfully contained early hotspot (CoMoLo study), Germany, May to June 2020. *Eurosurveillance*. 2020; 25(47):2001752.
36. Wagner R, Peterhoff D, Beileke S, Guenther F, Berr M, Einhauser S, et al. Estimates and determinants of SARS-CoV-2 seroprevalence and infection fatality ratio using latent class analysis: the population-based Tirschenreuth study in the hardest-hit German county in spring 2020. *Viruses*. 2021; 13(6):1118. <https://doi.org/10.3390/v13061118> PMID: 34200766
37. Guan Wj, Ni Zy, Hu Y, Liang Wh, Ou Cq, He Jx, et al. Clinical characteristics of coronavirus disease 2019 in China. *New England journal of medicine*. 2020; 382(18):1708–1720. <https://doi.org/10.1056/NEJMoa2002032>
38. World Health Organization. Coronavirus disease 2019 (COVID-19): situation report, 73. 2020;.
39. Gao Z, Xu Y, Sun C, Wang X, Guo Y, Qiu S, et al. A systematic review of asymptomatic infections with COVID-19. *Journal of Microbiology, Immunology and Infection*. 2021; 54(1):12–16. <https://doi.org/10.1016/j.jmii.2020.05.001> PMID: 32425996
40. Xu X, Sun J, Nie S, Li H, Kong Y, Liang M, et al. Seroprevalence of immunoglobulin M and G antibodies against SARS-CoV-2 in China. *Nature Medicine*. 2020; 26(8):1193–1195. <https://doi.org/10.1038/s41591-020-0949-6> PMID: 32504052
41. Pollán M, Pérez-Gómez B, Pastor-Barriuso R, Oteo J, Hernán MA, Pérez-Olmeda M, et al. Prevalence of SARS-CoV-2 in Spain (ENE-COVID): a nationwide, population-based seroepidemiological study. *The Lancet*. 2020; 396(10250):535–544. [https://doi.org/10.1016/S0140-6736\(20\)31483-5](https://doi.org/10.1016/S0140-6736(20)31483-5) PMID: 32645347
42. Stringhini S, Wisniak A, Piumatti G, Azman AS, Lauer SA, Baysson H, et al. Seroprevalence of anti-SARS-CoV-2 IgG antibodies in Geneva, Switzerland (SEROCoV-POP): a population-based study. *The Lancet*. 2020; 396(10247):313–319. [https://doi.org/10.1016/S0140-6736\(20\)31304-0](https://doi.org/10.1016/S0140-6736(20)31304-0) PMID: 32534626
43. Knabl L, Mitra T, Kimpel J, Roessler A, Volland A, Walsler A, et al. High SARS-CoV-2 seroprevalence in children and adults in the Austrian ski resort of Ischgl. *Communications Medicine*. 2021; 1(1):1–9. <https://doi.org/10.1038/s43856-021-00007-1>

44. Pullano G, Di Domenico L, Sabbatini CE, Valdano E, Turbelin C, Debin M, et al. Underdetection of cases of COVID-19 in France threatens epidemic control. *Nature*. 2021; 590(7844):134–139. <https://doi.org/10.1038/s41586-020-03095-6> PMID: [33348340](https://pubmed.ncbi.nlm.nih.gov/33348340/)
45. Neher RA, Dyrda R, Druelle V, Hodcroft EB, Albert J. Potential impact of seasonal forcing on a SARS-CoV-2 pandemic. *Swiss medical weekly*. 2020; 150 (1112).
46. Baker RE, Yang W, Vecchi GA, Metcalf CJE, Grenfell BT. Susceptible supply limits the role of climate in the early SARS-CoV-2 pandemic. *Science*. 2020; 369(6501):315–319. <https://doi.org/10.1126/science.abc2535>
47. Ferguson NM, Cummings DA, Fraser C, Cajka JC, Cooley PC, Burke DS. Strategies for mitigating an influenza pandemic. *Nature*. 2006; 442(7101):448–452. <https://doi.org/10.1038/nature04795> PMID: [16642006](https://pubmed.ncbi.nlm.nih.gov/16642006/)
48. Davies NG, Kucharski AJ, Eggo RM, Gimma A, Edmunds WJ, Jombart T, et al. Effects of non-pharmaceutical interventions on COVID-19 cases, deaths, and demand for hospital services in the UK: a modelling study. *The Lancet Public Health*. 2020; 5(7):e375–e385. [https://doi.org/10.1016/S2468-2667\(20\)30133-X](https://doi.org/10.1016/S2468-2667(20)30133-X) PMID: [32502389](https://pubmed.ncbi.nlm.nih.gov/32502389/)
49. Flaxman S, Mishra S, Gandy A, Unwin HJT, Mellan TA, Coupland H, et al. Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe. *Nature*. 2020; p. 1–8. PMID: [32512579](https://pubmed.ncbi.nlm.nih.gov/32512579/)
50. Radev ST, D'Alessandro M, Bürkner PC, Mertens UK, Voss A, Köthe U. Amortized Bayesian model comparison with evidential deep learning. *arXiv:2004.10629*. 2020;.