



Quantifying the Effects of Contact Tracing, Testing, and Containment Measures in the Presence of Infection Hotspots

25

LARS LORCH, ETH Zürich

HEINER KREMER, Max Planck Institute for Intelligent Systems

WILLIAM TROULEAU, EPFL

STRATIS TSIRTISIS, Max Planck Institute for Software Systems

ARON SZANTO, Zerobase Foundation

BERNHARD SCHÖLKOPF, Max Planck Institute for Intelligent Systems

MANUEL GOMEZ-RODRIGUEZ, Max Planck Institute for Software Systems

Multiple lines of evidence strongly suggest that infection hotspots, where a single individual infects many others, play a key role in the transmission dynamics of COVID-19. However, most of the existing epidemiological models fail to capture this aspect by neither representing the sites visited by individuals explicitly nor characterizing disease transmission as a function of individual mobility patterns. In this work, we introduce a temporal point process modeling framework that specifically represents visits to the sites where individuals get in contact and infect each other. Under our model, the number of infections caused by an infectious individual naturally emerges to be overdispersed. Using an efficient sampling algorithm, we demonstrate how to estimate the transmission rate of infectious individuals at the sites they visit and in their households using Bayesian optimization (BO) and longitudinal case data. Simulations using fine-grained and publicly available demographic data and site locations from Bern, Switzerland showcase the flexibility of our framework. To facilitate research and analyses of other cities and regions, we release an open-source implementation of our framework.

CCS Concepts: • **Computing methodologies** → **Agent/discrete models**; • **Information systems** → **Spatial-temporal systems**; Decision support systems;

Additional Key Words and Phrases: COVID-19, epidemiology, overdispersion, temporal point processes

ACM Reference format:

Lars Lorch, Heiner Kremer, William Trouleau, Stratis Tsirtsis, Aron Szanto, Bernhard Schölkopf, and Manuel Gomez-Rodriguez. 2022. Quantifying the Effects of Contact Tracing, Testing, and Containment Measures in the Presence of Infection Hotspots. *ACM Trans. Spatial Algorithms Syst.* 8, 4, Article 25 (October 2022), 28 pages.

<https://doi.org/10.1145/3530774>

This work was supported in part by SNSF under grant number 200021-182407 and the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 945719).

Authors' addresses: L. Lorch, ETH Zürich, Rämistrasse 101, Zürich 8092, Switzerland; email: lars.lorch@inf.ethz.ch; H. Kremer and B. Schölkopf, Max-Planck-Ring 4, Max Planck Institute for Intelligent Systems, Tübingen 72076, Germany; emails: {heiner.kremer, bs}@tuebingen.mpg.de; W. Trouleau, Route Cantonale, EPFL, Lausanne 1015, Switzerland; email: william.trouleau@gmail.com; S. Tsirtsis and M. G. Rodriguez, Paul-Ehrlich-Strasse 26, Max Planck Institute for Software Systems, Kaiserslautern 67663, Germany; emails: {stsirtsis, manuelgr}@mpi-sws.org; A. Szanto, Corporate Creations Network Inc. 1629 K Street, N.W. #300 Washington, DC 20006; email: aron@zerobase.io.



This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2022 Copyright held by the owner/author(s).

2374-0353/2022/10-ART25

<https://doi.org/10.1145/3530774>

1 INTRODUCTION

As countries around the world aim to counteract rising numbers of COVID-19 infections [2], overwhelmingly growing evidence suggests that few infected people in infection hotspots, or *super-spreading events* (SSEs), may be responsible for both explosive early growth of cases and sustained transmission in later stages [6, 12, 32, 40, 51, 61]. For example, in Hong Kong, the largest infection hotspots were traced back to four bars, which accounted for 32.5% of all locally acquired infections from January 23 to April 28, 2020 [6]. In South Korea, an infection hotspot linked to a church was responsible for at least 60% of all recorded cases by March 18, 2020, and over 1,000 infections were traced back to a single individual [3]. The first major outbreak in Germany occurred after an infected couple attended a carnival festivity in Heinsberg, with super spreading dynamics later verified by virus genome sequencing [74]. These lines of evidence suggest that, for COVID-19, the number of infections caused by single infectious individuals is *overdispersed*—most individuals infect a few and a few infect many, exhibiting greater variance than expected under Poisson assumptions [21, 42, 83]. Using carefully annotated tracing data, this has been identified as a root cause of SSEs [6, 32, 40, 51].

Most of the existing epidemiological models for studying containment measures, including those developed and used in the context of the COVID-19 pandemic, neither explicitly represent sites of transmission, nor do they characterize exposures as a function of individual mobility patterns. While this coarseness may be useful for fitting aggregate case trends, it makes conventional approaches unable to model the effects of granular interventions such as contact tracing or testing. Moreover, existing models either assume or result in a Poisson distribution of infections caused by an infectious individual, also called *secondary infections*, which fails to capture the high dispersion observed for COVID-19.¹ As a result, these models have been of little use for identifying conditions under which hotspots emerge [21, 40], helping design and study control measures tailored to prevent SSEs [10], and predicting where infection hotspots are most likely to occur [83].

In this work, we take a first step towards addressing the above limitations and present a data-driven framework for epidemiological modeling in the presence of overdispersed transmission dynamics and fine-grained containment measures. Our main contributions are as follows:

- (i) We introduce an event-based “check-in” mobility model that explicitly characterizes the frequency and duration of each individual’s visits to specific sites, which can be configured using a variety of publicly available data.
- (ii) We develop a novel rate of transmission at sites that quantifies the influence of these individual mobility patterns as well as environmental drivers and containment measures on the risk of exposure that each infected individual poses to others at a site. By using this transmission model and an explicit representation of the visited locations, our framework can directly characterize granular interventions that are targeted at particular sites and individuals (e.g., hygienic measures at workplaces, closures of schools, or contact tracing).
- (iii) We derive an efficient sampling algorithm for our model, which allows us to simulate the spread of COVID-19 under a variety of containment measures and counterfactual scenarios. Building on this procedure, we show how to estimate the disease transmission parameters using **Bayesian optimization (BO)** and longitudinal COVID-19 case data.

Our framework empirically scales to real-world cities and regions with hundreds of thousands of inhabitants and can be applied whenever simulated or real mobility traces as well as basic disease progression parameters (e.g., the incubation period and duration of infectiousness) are given.

¹Overdispersion has also been observed in MERS and SARS [56, 63, 66].

We showcase our approach using fine-grained demographic data and site locations from Bern, Switzerland, and other regions in Germany and Switzerland. Our results demonstrate that the number of individual disease transmissions—both overall and during a site visit—naturally emerges to be overdispersed, i.e., exhibiting higher variance than expected under the common Poisson assumption and that our model is able to robustly characterize the observed COVID-19 case trends. These findings hint at the potential of our framework as a complementary policy tool for studying the efficacy of containment measures, factors of disease transmission, and the nature of infection hotspots—hand in hand with existing societal and ethical considerations. To facilitate research and analyses in this area, we release an open-source implementation of our framework [57].

2 BACKGROUND

Temporal point processes are random processes whose realizations $\mathcal{H} = \{t_1, t_2, \dots, t_n\}$ consist of discrete events localized in time $t_i \in \mathbb{R}^+$ [5]. A temporal point process is commonly represented as a *counting process* $N(t)$, which counts the number of events that occurred before time t

$$N(t) = \sum_{t_i \in \mathcal{H}} u(t - t_i), \quad (1)$$

where $u(x)$ is the unit step function and equals 1 if $x \geq 0$ and 0 otherwise. Given a history of events $\mathcal{H}(t) = \{t_i \in \mathcal{H} \mid t_i \leq t\}$ until time t , we use a *conditional intensity function* $\lambda(t)$ to model the arrival probability of the next random event in the process. More specifically, the conditional intensity function $\lambda(t)$ models the probability of an event occurring in an arbitrarily small time window after t . We write

$$P(dN(t) = 1 \mid \mathcal{H}(t)) = \lambda(t) dt, \quad (2)$$

where the differential is defined as $dN(t) = N(t + dt) - N(t) \in \{0, 1\}$. Here, dt is an arbitrarily small time interval, and only one event can occur in $[t, t + dt)$. The intensity function $\lambda(t)$ can be interpreted as an instantaneous rate of events per unit of time, for example, $\lambda(t) = 5$ visits/week or $\lambda(t) = 1$ infection caused/hour. Note that $\lambda(t)$ may be time-varying and conditional on $\mathcal{H}(t)$.

In **stochastic differential equations (SDEs)** with jumps, the evolution of a set of state variables is characterized by the stochastic events of a set of counting processes. Jump SDEs are commonly used for modeling dynamical systems with discrete stochastic events in continuous time, such as visits to sites or infections with a disease. To illustrate, let $N(t)$ represent a counting process recording the number of emails sent to a person, and assume their inbox has a capacity of 1,000. Ignoring the deletion of emails, the change in the number of emails in the inbox $X(t)$ maybe expressed by the SDE $dX(t) = u(1,000 - X(t))dN(t)$, which increments $X(t)$ at every arrival of $N(t)$ until reaching the limit of 1,000.

3 A SPATIOTEMPORAL EPIDEMIC MODEL

In this section, our goal is to develop an agent-based, compartmental epidemiological model under which fine-grained spatiotemporal interventions can be expressed formally and the distribution of secondary infections induced by the model can exhibit overdispersion.

To this end, our framework is composed of a collection of binary state variables that determine the mobility pattern, epidemiological condition, and testing status of each single individual $i \in \mathcal{V}$. We model the state transitions using SDEs with jumps, a model class that captures (i) the stochastic nature of infection events and mobility patterns, (ii) events in continuous time, i.e., *not* in aggregate over a period, and (iii) discrete state transitions—an individual either does *or* does not get infected, visit a site, or get tested positively.

In the remainder of this section, we formally describe the dynamics of each state variable of the model. To ease the exposition, we distinguish between variables related to mobility, epidemiology, testing, and containment measures. Later, in Section 4, we then show how to generate random forward simulations of the entire model by devising an efficient sampling algorithm.

3.1 Mobility

For each individual $i \in \mathcal{V}$ and a set of sites \mathcal{S} that individuals can visit, let $P_{i,k}(t) = 1$ if the individual is at site $k \in \mathcal{S}$ at time t and $P_{i,k}(t) = 0$ otherwise. We characterize the value of the states $P_{i,k}(t)$ using the following SDE with jumps:

$$dP_{i,k}(t) = dU_{i,k}(t) - dV_{i,k}(t). \quad (3)$$

$U_{i,k}(t)$ and $V_{i,k}(t)$ are counting processes that record the events of individual i arriving at and leaving from site $k \in \mathcal{S}$, respectively. Thus, Equation (3) captures that $P_{i,k}(t)$ increments to 1 after person i arrives at site k and decrements to 0 after the person leaves. We define the dynamics of the state transitions as

$$\begin{aligned} P(dU_{i,k}(t) = 1 | \mathcal{H}(t)) &= \eta_{i,k}(t) \prod_{l \in \mathcal{S}} (1 - P_{i,l}(t)) dt, \\ P(dV_{i,k}(t) = 1 | \mathcal{H}(t)) &= v_k U_{i,k}(t) dt \end{aligned} \quad (4)$$

where $\eta_{i,k}(t)$ is the rate at which individual i visits site k and $1/v_k$ is the average duration of a visit to site k . Equation (4) makes the state variable $P_{i,k}(t) \in \{0, 1\}$ well-defined by ensuring that $P_{i,k}(t) = 1$ for only one site $k \in \mathcal{S}$ at a time.

To configure the rates $\eta_{i,k}(t)$ and average duration $1/v_k$ for every individual and site, one can resort to publicly available data. In our simulations, we configure the individual mobility statistics using the spatial distribution of real site locations, high-resolution population density data, country-specific information about the household structure, and region-specific age demographics. We also assume that the probability that an individual i visits a specific site k decreases with the distance between their household and the site, similar to the gravity model [85]. Figure 1 illustrates the sites \mathcal{S} in a mobility model of Bern, Switzerland, which will be used for the case study in Section 5.

3.2 Epidemiology

To model the health status of each individual $i \in \mathcal{V}$ while being in contact with others at sites \mathcal{S} of the mobility model, we build on recent variations of the **Susceptible-Exposed-Infected-Resistant (SEIR)** compartment models that have been introduced in the context of COVID-19 modeling [37, 54]. More specifically, we define the epidemiological condition of each individual $i \in \mathcal{V}$ using the indicator state variables $\mathbb{S}(t) = \{S_i(t), E_i(t), I_i^a(t), I_i^p(t), I_i^s(t), H_i(t), R_i(t), D_i(t)\}_{i \in \mathcal{V}}$ with each $\in \{0, 1\}$, whose meaning is specified in Table 1.

Exposure. First, we formally characterize the state transition of individual $i \in \mathcal{V}$ from being susceptible ($S_i(t)$) to being exposed ($E_i(t)$) using the following jump SDEs:

$$\begin{aligned} dS_i(t) &= -dN_i(t) \\ dE_i(t) &= dN_i(t) - dM_i(t). \end{aligned} \quad (5)$$

The counting process $N_i(t)$ models the exposure of individual $i \in \mathcal{V}$ and thus forms the core component of our epidemiological model. $M_i(t)$ models the subsequent transition to the asymptomatic or presymptomatic state.

To connect the epidemiological model to the mobility model, we assume that *an individual's instantaneous rate of exposure increases by a constant site-specific transmission rate β_k when in contact*

Table 1. Epidemiological State Variables $\mathbb{S}(t)$

State	Description	Infected	Contagious	Symptoms
$S_i(t)$	is susceptible	-	-	-
$E_i(t)$	is exposed	✓	-	-
$I_i^a(t)$	is asymptomatic, mild course of disease	✓	✓	-
$I_i^p(t)$	is presymptomatic, progresses to $I_i^s(t)$ later	✓	✓	-
$I_i^s(t)$	is symptomatic	✓	✓	✓
$H_i(t)$	is hospitalized	✓	✓	✓
$R_i(t)$	is resistant and recovered	-	-	-
$D_i(t)$	has died	-	-	-

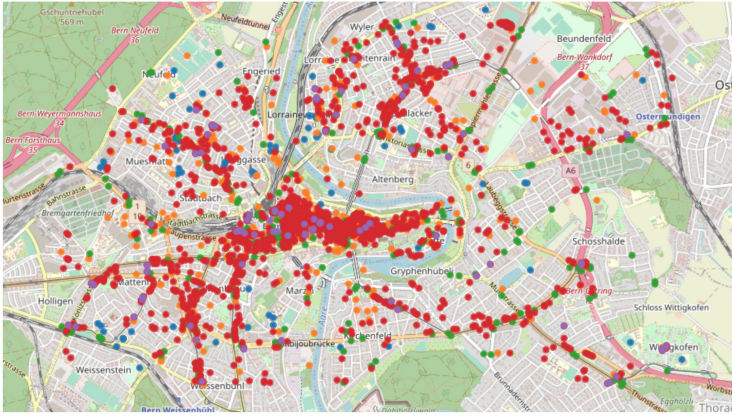


Fig. 1. Site locations by category in the mobility model of Bern, Switzerland. The colored dots depict schools and research institutes (blue), social places (orange), bus stops (green), workplaces (red), supermarkets (purple).

with another infectious individual at a site $k \in \mathcal{S}$. Consequently, the exposure rate of each individual i only depends on the individual's contacts at sites $k \in \mathcal{S}$ based on the mobility traces $P_{i,k}(t)$ —and not the contacts of others. We capture this model of exposure by the following conditional intensity function $\lambda_i(t)$ of the exposure counting process $N_i(t)$:

$$\lambda_i(t) = S_i(t) \sum_{k \in S} P_{i,k}(t) \sum_{j \in \mathcal{V} \setminus \{i\}} \int_{t-\delta}^t K_{j,k}(\tau) \gamma e^{-\gamma(t-\tau)} d\tau, \quad (6)$$

where

$$K_{j,k}(\tau) = P_{j,k}(\tau) \beta_k \left(I_j^p(\tau) + I_j^s(\tau) + \mu I_j^a(\tau) \right)$$

and $P(dN_i(t) = 1 \mid \mathcal{H}(t)) = \lambda_i(t) dt$. In the above:

- (i) $\beta_k \geq 0$ is the base transmission rate of presymptomatic and symptomatic individuals at site k . Depending on the availability of labeled and unlabeled data, one may consider sharing the same parameter for all sites $k \in \mathcal{S}$ or for sites of the same category, e.g., distinguishing only indoors and outdoors. The scale $\mu \in [0, 1]$ denotes the relative transmission rate of asymptomatic compared to (pre-)symptomatic individuals.

- (ii) $K_{j,k}(\tau) \in \{0, \mu\beta_k, \beta_k\}$ captures the effective contribution of individual j to transmission at site k and is non-zero if and only if j is both infected (symptomatic, presymptomatic, or asymptomatic) and present at site k at time τ .
- (iii) $\int_{t-\delta}^t K_{j,k}(\tau) \gamma e^{-\gamma(t-\tau)} d\tau$ models environmental transmission by accounting for the fact that the virus survives for some period of time δ on surfaces or in the air after an infected individual has left a site [71].

Thus, while traditional epidemiological models constrain the number of exposures to homogeneous Poisson distributions, our model in Equation (6) employs a stochastic and dynamically adjusting exposure rate for each individual $i \in \mathcal{V}$ based on the mobility traces $P_{i,k}(t)$, under which this constraint is lifted. Infections within households can be characterized by adding an analogous term $\lambda_{\mathcal{H}(i)}(t)$ with household transmission rate ξ to $\lambda_i(t)$, which is outlined in Appendix A.

Disease progression. After an individual $i \in \mathcal{V}$ got exposed, the subsequent state transitions are characterized by the counting processes $W_i(t)$, $Y_i(t)$, $Z_i(t)$, $R_i^a(t)$, and $R_i^s(t)$. Let $a_i \sim \text{Bern}(\alpha_a)$ indicate whether an infected individual i is asymptomatic or not. Then, in the asymptomatic case where $a_i = 1$, individual i progresses from exposed to asymptomatic (via $M_i(t)$) and ultimately to recovered (via $R_i^a(t)$)

$$dI_i^a(t) = a_i dM_i(t) - dR_i^a(t). \quad (7)$$

In the symptomatic case where $a_i = 0$, the individual progresses from exposed to presymptomatic (via $M_i(t)$), from presymptomatic to symptomatic (via $W_i(t)$) and from symptomatic to resistant (via $R_i^s(t)$). In addition, symptomatic individuals may be hospitalized or die from the disease. Let $h_i \sim \text{Bern}(\alpha_h)$ indicate whether the individual eventually requires hospitalization, and $b_i \sim \text{Bern}(\alpha_b)$ whether they eventually die from the disease. Then, a symptomatic individual may also transition from symptomatic infected to hospitalized (via $Y_i(t)$) and from symptomatic infected to dead (via $Z_i(t)$). All in all, the presymptomatic ($I_i^p(t)$), symptomatic ($I_i^s(t)$) and resistant ($R_i(t)$) state variables are characterized by

$$\begin{aligned} dI_i^p(t) &= (1 - a_i) dM_i(t) - dW_i(t) \\ dI_i^s(t) &= dW_i(t) - (1 - b_i) dR_i^s(t) - b_i dZ_i(t) \\ dR_i(t) &= a_i dR_i^a(t) + (1 - a_i) dR_i^s(t). \end{aligned} \quad (8)$$

Moreover, the hospitalized ($H_i(t)$) and dead ($D_i(t)$) states are given by

$$\begin{aligned} dH_i(t) &= h_i I_i^s(t) dY_i(t) - (1 - b_i) H_i(t) dR_i^s(t) - b_i H_i(t) dZ_i(t) \\ dD_i(t) &= b_i dZ_i(t). \end{aligned} \quad (9)$$

Since disease progression is disjoint from the mobility model, we follow the literature in modeling the above transition times using easy-to-sample log-normal distributions [52, 55]—starting at the time $E_i(t)$, $I_i^p(t)$, $I_i^a(t)$, or $I_i^s(t)$ become one, respectively, and terminating after their first event. In practice, we fix their parameters based on estimates of the mean transition durations from the clinical COVID-19 literature.

3.3 Testing

Individuals are tested according to a testing policy $\pi_{\text{test}}(t)$, e.g., testing only symptomatic or vulnerable people, at a rate $\lambda_{\text{test}}(t)$, which can be chosen to match location-specific testing statistics. The test outcomes are only known after a reporting delay Δ_{test} . Formally, the counting process $T(t)$ records the number of known test outcomes by time t . Let $T_i^+(t)$ and $T_i^-(t)$ be the number of

times an individual $i \in \mathcal{V}$ has been tested positive and negative, respectively, by time t . Then, we characterize the state variables $T_i^+(t)$ and $T_i^-(t)$ using the following SDEs:

$$\begin{aligned} dT_i^+(t) &= \left(E_i(t) + I_i^a(t) + I_i^p(t) + I_i^s(t) \right) d_i(t) dT(t + \Delta_{\text{test}}), \\ dT_i^-(t) &= \left(S_i(t) + R_i(t) \right) d_i(t) dT(t + \Delta_{\text{test}}), \end{aligned} \quad (10)$$

where $d_i(t) \in \{0, 1\} \sim \pi_{\text{test}}(t)$ indicates whether i is tested at time t according to the policy. In the above, a test result is positive if the individual is exposed ($E_i(t)$) or infected ($I_i^a(t) + I_i^p(t) + I_i^s(t)$), and negative if the individual is either susceptible ($S_i(t)$) or recovered ($R_i(t)$). This can be relaxed to account for test specificity and sensitivity.

3.4 Containment Measures

In the above context, we can formally model a variety of containment measures that not only affect the broad population \mathcal{V} but also target specific sites or individuals, possibly in a time-variant fashion. These may range from more granular (e.g., isolating individuals who have tested positive for 14 days or who had contact with a positively tested individual) to less granular (e.g., implementing a state of “lockdown” for the entire population). The effect of mobility reduction and quarantine can be characterized by reducing the rates $\eta_{i,k}(t)$ at which individuals visit sites in the mobility model. Hygienic measures (e.g., face masks) can be implemented by reducing the transmission rate β_k at specific sites (e.g., workplaces). In all cases, the measures reduce the conditional intensities $\lambda_i(t)$ of the exposure counting processes $N_i(t)$, possibly dynamically based on the values of other state variables at time t .

Moreover, if desired, we may assume that contacts between individuals at sites are registered by a peer-to-peer proximity-based tracing system, analogous to the smartphone-based Bluetooth systems that have been implemented in the context of the COVID-19 pandemic [7]. A contact between individuals i and j will be registered if (i) their visit times at a specific site $k \in \mathcal{S}$ overlap, and (ii) both opt to use the proximity-based tracing system, e.g., by means of carrying a Bluetooth device. Visit times are said to overlap when $P_{i,k}(t) = 1$ and $P_{j,k}(t) = 1$ for some site $k \in \mathcal{S}$ and time t . When an individual i is tested positive, their registered contacts may be advised to isolate or seek testing themselves as described in Section 3.3. For contact tracing, the type of intervention may depend on the risk of exposure caused by the positively tested individual, which can be estimated using our model. Appendix B provides further details.

4 MODEL SIMULATION AND ESTIMATION

4.1 Epidemiological Sampling Algorithm

Having formally defined the model dynamics in Section 3, we now introduce a procedure to sample trajectories of the individual epidemiological states $\mathbb{S}(t)$ over a time horizon $t \in [0, t_{\text{max}})$, which ultimately allows us to empirically study the spread of the disease under a variety of scenarios. The initial conditions $\mathbb{S}(0)$, a testing policy $\pi_{\text{test}}(t)$, and the mobility traces $P_{i,k}(t)$ are assumed to be fixed *a priori*—from simulations of a synthetic mobility model as in Section 3.1 or real-world data.

To sample a trajectory of the epidemiological state variables, we start by noticing that their values change at—and only change at—*events* of the counting processes that model the transitions in the model SDEs. Hence, all state variables $\mathbb{S}(t)$ are constant between two consecutive events when considering the event times of all counting processes in the model on *one* timeline. This leads us to the backbone principle for generating random realizations of the model: we initialize the state variables $\mathbb{S}(0)$, sample the next time of state transition for each $i \in \mathcal{V}$, and push these transition events onto *one* temporally-sorted priority queue Q that simultaneously tracks the next

events for *all* individuals in the model. The algorithm then repeatedly loops through: (i) popping the next event e from Q ; (ii) updating the state of individual i associated with e ; (iii) sampling the next time t of state transition e' for i ; and (iv) pushing e' to Q with priority t .

As explained in Section 3, we fix the time-to-event distributions of all processes *not* concerning exposure, i.e., excluding $\{N_i(t)\}_{i \in \mathcal{V}}$, to independent, easy-to-sample distributions as estimated by clinical COVID-19 literature. This means that once an individual is exposed, sampling the following times of state transition, e.g., to symptomatic and recovered states, is trivial. However, sampling the time of exposure of i , i.e., the first event time of $N_i(t)$, is hard because the rate $\lambda_i(t)$ dynamically interacts with all other stochastic state variables $\mathbb{S}(t)$ via the mobility model $P_{i,k}(t)$. To be able to sample from $N_i(t)$, we decompose the intensity $\lambda_i(t)$ into a sum of contributions $\lambda_{j \rightarrow i}(t)$ caused by other individuals j :

$$\lambda_i(t) = S_i(t) \sum_{j \in \mathcal{V} \setminus \{i\}} \sum_{k \in \mathcal{S}} P_{i,k}(t) \int_{t-\delta}^t K_{j,k}(\tau) \gamma e^{-\gamma(t-\tau)} d\tau =: S_i(t) \sum_{j \in \mathcal{V} \setminus \{i\}} \lambda_{j \rightarrow i}(t), \quad (11)$$

where the last summation over $j \in \mathcal{V} \setminus \{i\}$ is sparse as it only indexes over contacts of individual i after time t . Note that $\lambda_{j \rightarrow i}(t) = 0$ when i and j are not in contact directly or when j left site $k \in \mathcal{S}$ more than δ -time before i arrived.

By Equation (11), if individual i is susceptible, the counting process $N_i(t)$ can be seen as a superposition of several processes $N_{j \rightarrow i}(t)$ with intensities $\lambda_{j \rightarrow i}(t)$. This implies that the time-to-event distribution of $N_i(t)$ is equivalent to the distribution of the time to the *first* arrival of all processes $N_{j \rightarrow i}(t)$ [5, 27]. Using the temporal ordering invariant of Q , we can thus process valid exposure events on the fly. Whenever an individual j becomes infectious, i.e., $I_j^a = 1$ or $I_j^p = 1$, we sample the next exposure event that j *causes* for every individual i in contact with j in the future at rate $\lambda_{j \rightarrow i}(t)$, and push these events onto Q . Later, when an exposure event e for individual i is popped from Q in step (i), we check whether e is the *first* exposure of i by verifying $S_i(t) = 1$, and discard subsequent exposure events for i .

To sample the next event time of the subprocess $N_{j \rightarrow i}(t)$ after time t' , we use the principle of *thinning* [27]. We can generate a valid sample from $N_{j \rightarrow i}(t)$ by repeatedly adding $\tau \sim \text{Expo}(\lambda_{j \rightarrow i}^{\max})$ to t' and stopping with probability $\lambda_{j \rightarrow i}(t')/\lambda_{j \rightarrow i}^{\max}$ at a given iteration, where $\lambda_{j \rightarrow i}^{\max}$ is an upper bound on $\lambda_{j \rightarrow i}(t)$. We skip zero-intensity windows whenever reaching $\lambda_{j \rightarrow i}(t) = 0$ during thinning, which is sound by viewing $N_{j \rightarrow i}(t)$ itself as a superposition of counting processes, one for each interval of non-zero intensity, and skipping their initial zero-rate periods by the memoryless property.² If j recovers, i.e., $R_j(t) = 1$, then $\lambda_{j \rightarrow i}(t)$ in Equation (11) is dynamically set to 0 at the time of recovery because $K_{j,k}(t) = 0$. By the principle of thinning, all exposure events caused by j beyond this point, sampled back when j got infectious, are discarded on the fly, i.e., when they get popped from Q .

Combining the above, we arrive at an efficient sampling procedure for the epidemiological model SDEs using a single priority queue Q , which is formally defined in Algorithms 2 and 3 of Appendix D. In this context, we note that interventions like social distancing or hygienic measures always *reduce* the rates $\lambda_{j \rightarrow i}(t)$ and can thus likewise be implemented using thinning, i.e., rejecting the affected exposure events with some probability.

Our sampling procedure is a Las Vegas algorithm, i.e., its runtime is a random variable but its output, i.e., the sampled trajectory of state variables, is always faithful. This is because the number of state transition events we sample and subsequently process depends on the number of infectious individuals, which is itself a random variable. The following proposition bounds the expected

²If $T \sim \text{Expo}(\lambda)$, then $P(T \geq t + s \mid T \geq s) = P(T \geq t)$.

runtime of our sampling procedure under some mild technical assumptions on the mobility traces $P_{i,k}(t)$. A proof is given in Appendix D.

PROPOSITION 1. *Assume that any given individual $i \in \mathcal{V}$ makes $O(t_{\max})$ visits to sites \mathcal{S} , the mobility model is sparse, i.e., every individual $i \in \mathcal{V}$ has $O(1)$ unique contact persons³, and there are no containment measures. Then, the expected runtime of our sampling procedure for generating a trajectory of the epidemiological states $\mathbb{S}(t)$ over a time horizon $[0, t_{\max})$ is given by*

$$O\left(|\mathcal{V}| \left(t_{\max} \log(t_{\max}|\mathcal{V}|) + \frac{1}{q} t_{\max} \log(t_{\max})\right)\right),$$

where $q \in (0, 1)$ is a constant known a priori that depends on the parameters of the epidemiological model.

The above result implies that, if the number of sites individuals visit increases linearly with time and the number of unique contact persons is constant, then the expected runtime of our sampling procedure is *quasilinear* in the number of individuals $|\mathcal{V}|$ and the length of the sampled trajectory t_{\max} . Moreover, it is worth pointing out that generating random rollouts of the model can be embarrassingly parallelized. Finally, in our experiments, we have empirically found that our sampling procedure scales to regions of more than one hundred thousand individuals \mathcal{V} with around one thousand sites \mathcal{S} .

4.2 Parameter Estimation

Building on the sampling algorithm, we can estimate the unspecified epidemiological parameters $\theta = \{\beta_k, \xi\}$, i.e., the transmission rate of individuals at sites and in their households, in a given epidemiological scenario. More specifically, provided a set of initial conditions $\mathbb{S}(0)$, testing policy π_{test} , a priori fixed mobility traces $P_{i,k}(t)$, and fixed parameters of the processes not concerning exposure, i.e., excluding $\{N_i(t)\}_{i \in \mathcal{V}}$, we find the parameters θ that provide the best fit to the observed COVID-19 cases in a given region. To this end, we view the model simulation as a black box and apply BO, which amounts to iteratively building a surrogate model of our objective and evaluating at promising parameter settings [16].

Following the standard BO paradigm, we interpret the *expected number of positive cases at time t* in our model as a black box function $g_t(\theta)$ where

$$g_t(\theta) := \mathbb{E}_{\mathcal{T} \sim \theta} \left[\sum_{T_i^+ \in \mathcal{T}} T_i^+(t) \right]. \quad (12)$$

The expectation in Equation (12) is defined over realizations of the testing state variables $\{T_i^+(t)\} =: \mathcal{T} \sim \theta$ of the model with exposure parameters θ . In practice, $g_t(\theta)$ is only observed via noisy evaluations at different values of θ since the expectation is approximated using a Monte Carlo estimate of J random simulations. \mathcal{T} is stochastic not only due to the counting processes, but in absence of real mobility traces also due to random seeds $\mathbb{S}(0)$ and synthetic $P_{i,k}(t)$, independently simulated for each rollout.

The objective we aim to minimize is the *mean daily squared error* of cumulative positive cases between the model predictions and the real observed COVID-19 cases of the region. This allows us to form a link between the spatiotemporal states of each individual in the model and aggregate longitudinal case data. The squared error has previously been considered in parameter estimation for black-box models [11] and in the context of COVID-19 research [23]. Let c_t^{true} be the cumulative number of real COVID-19 cases at the end of the day t as provided by the national authorities. Then,

³This implies that the individuals \mathcal{V} and sites \mathcal{S} are not considered independently. Formally stated in terms of the mobility model, we assume $\sum_{k \in \mathcal{S}} P_{i,k} = O(t_{\max})$ and $\sum_{j \in \mathcal{V}} \sum_{k \in \mathcal{S}} P_{i,k} P_{j,k} = O(1)$.

ALGORITHM 1: Parameter estimation using Bayesian optimization

Input: Black-box simulator $g_t(\theta)$, parameter domain $\text{dom}(\theta)$, time horizon t_{\max} , case data c_t^{true} , hyperparameters J, M, N

- 1: $s(\mathbf{x}) := -\sum_{t=1}^{t_{\max}} (c_t^{\text{true}} - \mathbf{x}_t)^2$
- 2: $\theta_{1:M} \leftarrow$ first M quasi-random settings
- 3: $\mathcal{D} \leftarrow \emptyset$
- 4: **for** $i \in [M]$ **do** ▷ Quasi-random initialization
 - 5: Obtain daily sim. result $g_t(\theta_i)$ from J random roll-outs
 - 6: $\mathcal{D} \leftarrow \mathcal{D} \cup \{(\theta_i, g_t(\theta_i))\}$
- 7: **while** $|\mathcal{D}| \leq N$ **do** ▷ Bayesian Optimization
 - 8: $p(g_t(\theta)) \mid \mathcal{D} \leftarrow \text{GaussianProcessPosterior}(\mathcal{D})$
 - 9: $\theta^* \leftarrow \arg \max_{\theta' \in \text{dom}(\theta)} \text{KnowledgeGradient}(p, \theta')$ [79]
 - 10: Obtain daily sim. result $g_t(\theta^*)$ from J random roll-outs
 - 11: $\mathcal{D} \leftarrow \mathcal{D} \cup \{(\theta^*, g_t(\theta^*))\}$
- 12: **return** $\arg \max_{(\theta, g_{1:t_{\max}}(\theta)) \in \mathcal{D}} s(g_{1:t_{\max}}(\theta))$

our objective f to be minimized is a composition of the squared error score and per-day black-box functions $g_t(\theta)$ averaged over a time of t_{\max} days:

$$f(\theta) = \frac{1}{T} \sum_{t=1}^{t_{\max}} \left(c_t^{\text{true}} - g_t(\theta) \right)^2. \quad (13)$$

The compositionality of f can allow for greater sample efficiency [11, 13], in particular when estimating additional parameters. However, when only estimating $\theta = \{\{\beta_k\}, \xi\}$ and β_k held constant across sites, we found it to be favorable for the BO surrogate model to directly learn $f(\theta)$, as opposed to the daily $g_t(\theta)$, as the black-box function. We use the knowledge gradient acquisition function [79] to navigate parameter proposals, which often shows favorable performance in noisy settings [13, 39]. Combining the above with the default BO procedure, our resulting parameter estimation algorithm is summarized in Algorithm 1.

5 A CASE STUDY OF BERN, SWITZERLAND

In Sections 3 and 4, we introduced a framework for epidemiological modeling and transmission parameter estimation in mobility models of any region of interest. In summary, the application of our model presupposes (i) a set of mobility traces $P_{i,k}(t)$ from synthetic mobility simulations or real-world data of the region, (ii) the time distributions for disease progression after infection, (iii) initial conditions $\mathbb{S}(0)$ or assumptions about influx of infected individuals, and (iv) a testing policy $\pi_{\text{test}}(t)$.

In the following, we showcase the flexibility of our framework in a case study of the city of Bern, Switzerland, and analyze the overdispersion of secondary infections as well as the course of the COVID-19 epidemic under various, fine-grained containment measures. We present supplementary results for additional regions of Germany and Switzerland in Appendix E. More generally, the progression we follow in Sections 5.1 and 5.2 can be viewed as step-by-step instructions to configure and calibrate the model to any desired region or disease variant.

5.1 Experimental Setup

Mobility traces. We leverage fine-grained demographic data and open-source site locations to build a mobility model for Bern, Switzerland, that contains $|\mathcal{V}| = 133,790$ individual inhabitants visiting $|\mathcal{S}| = 2,174$ real points of interest. The individuals \mathcal{V} belong to one of nine age groups according to the real demographics of the region. They are placed in households of up to five people according to their age and reported household structure in Switzerland [18]. The households

themselves are located across the spatial expansion of the city using high-resolution population density data provided by *Facebook Data for Good* [1]. To obtain relevant site locations in the regions of interest, we use geolocation data provided by *OpenStreetMap* [4]. Specifically, we retrieve the location of all sites \mathcal{S} in five site categories: *education* (schools, universities, research institutes), *social* (restaurants, cafés, bars), *transportation* (bus stops), *work* (offices, shops), and *groceries* (supermarkets, convenience stores). The sites \mathcal{S} are visualized in Figure 1.

Since real check-in traces $P_{i,k}(t)$ of the population of Bern are not publicly available, we simulate synthetic mobility traces from the model in Section 3.1 under the assumption of the gravity model [85]. In particular, we assume that each individual $i \in \mathcal{V}$ visits only a constrained set of unique sites $\mathcal{S}_i \subset \mathcal{S}$, which are selected with probability inversely proportional to the squared distance from their homes. This reflects the fact that individuals typically study or work at only one place, form habits regarding the public transportation they use, and social places or supermarkets they visit. We set the check-in rate $\lambda_{i,k}(t)$ of Section 3.1 to a constant value that depends on the individual's age group and site type; see Table 3 in the Appendix. The mean duration $1/v_k$ at sites of type education and work, social, transportation, and groceries are fixed to 120, 90, 12, and 30 minutes, respectively. We sometimes set these times to lower values than one would expect because individuals are neither exposed to all others at a site nor continuously exposed during their visit.

Disease parameters. As described in Section 3.2, we fix the parameters for disease progression after exposure based on recent estimates of the COVID-19 literature. The values we use are summarized in Table 2. We set mortality and hospitalization rates per age group using COVID-19 case data of the county-level administrative region [17] and previous studies [36].

Testing. To abstract away from testing criteria implemented in different regions, we assume that only true symptomatic individuals are registered for testing and that tests have perfect accuracy. We set the reporting delay Δ_{test} to 30 minutes, accounting for the now frequently available⁴ rapid tests [33], and assume that there is sufficient capacity to test all selected individuals. Moreover, positively tested individuals and their household members are quarantined for 14 days in isolation from each other.

5.2 Model Fit and Parameter Estimation

To estimate the transmission rate at sites β and in households ξ , we consider the time horizon from early March 2020 until May 2020 since it includes both times before and during governmental interventions in Switzerland, which occurred largely from March 16, 2020 to May 10, 2020 [28]. We use COVID-19 case data of the county-level administrative region [17] to define the objective (13) and run the procedure described in Section 4.2 for $N \geq 100$ steps with $M = 10$ initial quasi-random settings and $J = 200$ rollouts. During governmental interventions, the check-in frequencies of individuals at sites in the mobility model are reduced as estimated by Google mobility data in the region [41], and education and social sites are closed, i.e., not visited at all. In the model simulations used for estimation, each of the J realizations is randomized across realizations of the synthetic mobility traces and infection seeds. For parameter estimation, the initially exposed and infectious individuals are heuristically selected as described in Appendix C based on knowledge about the case numbers at the start date of the estimation period.

Figure 2 visualizes both the objective values $f(\theta)$ obtained at various settings for the transmission rates β and ξ as well as the model predictions for the cumulative cases during the time window of parameter estimation. The contour plot indicates that there is a single and identifiable optimal parameter regime, whose optimal values were estimated as $\beta = 0.0337$ and $\xi = 0.0038$.

⁴During parameter estimation, Δ_{test} is set to 48 hours to account for the test delay early in the pandemic.




Fig. 2. Transmission parameter estimation for the model of Bern, Switzerland. The left plot shows a contour plot of the log objective in (13) as a function of the transmission rate at sites β and in households ξ . The estimated parameters $\beta = 0.0337$ and $\xi = 0.0038$ lie in the identifiable optimal region. The right plot shows the predicted and real cumulative cases under the estimated parameters. The “lockdown” measures in Switzerland were implemented on March 16, 2020 as indicated by the dashed vertical line. The solid line indicates the mean, the shaded areas one and two standard deviations across 100 random simulations, respectively.

Furthermore, we find that the simulations using the estimated parameters are able to accurately match the observed longitudinal trend of cases during the estimation period early in the epidemic. Beyond the model of Bern, Switzerland, Appendix E provides a collection of parameter estimation results for four additional regional models of other urban and rural regions in Switzerland and Germany [19, 50, 64]. These supplementary findings confirm a similarly identifiable optimal parameter regime and demonstrate that both the epidemiological model and the transmission parameter estimation procedure are robustly applicable to other regional mobility models.

In the remainder of this section, we use the estimated transmission parameters for the model of Bern in all of our experiments. We first empirically study the degree of overdispersion in the number of secondary infections caused by infectious individuals under our mobility and fitted transmission model. We then use our framework to quantify the effects of a range of containment measures. To create a general epidemiological scenario, we assume a small but continual influx of five untraceable exogenous exposures per 100,000 inhabitants and per week and simulate the model state variables over a period of four months.

5.3 Overdispersion of Secondary Infections

As argued in Section 1, existing epidemiological models have predominantly built on homogeneous Poisson transmission dynamics that fail to capture the overdispersion of secondary infections observed for COVID-19. In addition, they do not explicitly model visits to sites where exposures occur. As a result, these models have been of little use for studying and predicting where and when infection hotspots are most likely to occur [10, 21, 40, 83].

In contrast to previous work, we find that *overdispersion of the distribution of secondary infections emerges naturally under our model*. Using the previously specified and estimated model parameters, we simulate the spread of COVID-19 under no containment measures other than the testing of symptomatic and isolation of positively tested individuals. During these simulations, we count the number of secondary infections caused by individuals that got infectious during a 7-day window after 1 month of the model simulations. Using two goodness-of-fit tests for the Poisson distribution, the Chi-squared (χ^2) and variance tests (VT) [24, 38], we are able to reject the null hypothesis that the distribution of secondary infections, both overall and when stratified per visit, follows a Poisson distribution. In particular, for both distributions of secondary infections, we obtain $p_{\chi^2} < 10^{-8}$ and $p_{VT} < 10^{-8}$. With sample variance generally significantly exceeding the

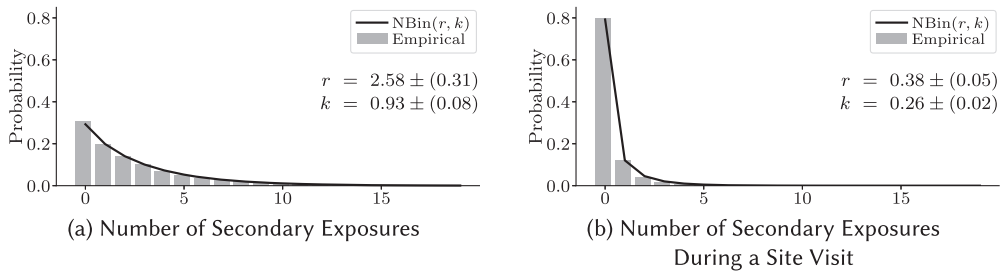


Fig. 3. Overdispersion of Secondary Infections. Panels (a) and (b) show a generalized negative binomial model fitted by maximum likelihood estimation to the secondary infections caused by an infectious individual, overall and stratified by site visits, respectively, and averaged over 100 random realizations. The secondary infections are counted for individuals that got infectious during a 7-day window after 1 month of the simulation under no containment measures other than the testing of symptomatic and isolation of positively tested individuals.

sample mean, both ways of counting the number of secondary infections naturally exhibit a higher variance than expected under the Poisson assumption and are thus *overdispersed*.

To measure the degree of overdispersion, we follow recent work in the context of COVID-19 [12, 32] and fit a generalized negative binomial distribution $\text{NBin}(r, k)$, an overdispersed generalization of the Poisson, where $r > 0$ is the mean or reproduction number, and k is the dispersion parameter. Figure 3 summarizes the results. Averaged over 100 random realizations, we find that the dispersion parameter $k < 1$ both overall and when stratified per visit ($k = 0.93 \pm (0.08)$ and $k = 0.26 \pm (0.02)$, respectively), evidence of substantial overdispersion [32]. We hypothesize that the higher overdispersion observed when aggregating per visit is a direct effect of the interaction between the stochastic check-in mobility model and our model of transmission at sites.

5.4 Efficacy of Containment Measures

Reducing contact at public sites by restricting individual mobility has been one of the most prevalent measures to counteract the spread of COVID-19 [46]. Our modeling framework allows us to faithfully study how effective various variants of this approach are at, e.g., containing the disease, reducing peak hospitalizations, or changing the effective reproduction number R_t over time. Instead of restricting the mobility of the entire population or only vulnerable groups, previous work has, for instance, proposed to divide the population into two subgroups that get isolated on alternating days [48, 58].

Figure 4 shows a comparative analysis of three of these variants: restricting the mobility of everyone, only vulnerable groups, or one of two random subgroups on alternating days. The measures are implemented as described in Sections 3.3–3.4. For each variant, we consider different levels of mobility restriction where individual check-in activity at sites in the mobility model is reduced by between 5% and 75%. In our simulations, the vulnerable groups are defined as individuals older than 60 years, who typically suffer more complications from COVID-19 [17, 64]. Our findings highlight the fact that the efficacy of each policy strongly depends on the degree to which individual movement activity is reduced. While restricting the mobility of everyone is overall clearly most effective, our findings suggest that isolating (i.e., reducing the mobility by 100%) one of two subgroups on alternating days can reduce the effective reproduction number, averaged over the phase of exponential case growth, and peak hospitalizations as much as reducing the mobility of everyone by 50%. Moreover, our results also suggest that the morally debatable strategy of quarantining only vulnerable groups does not live up to its expectation of reducing peak hospitalizations significantly.

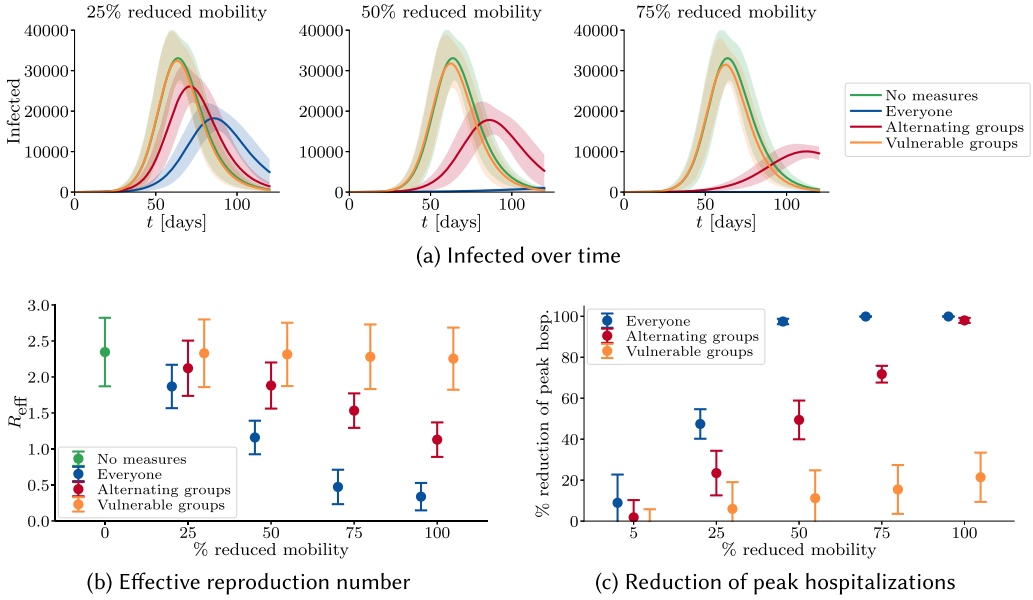


Fig. 4. Mobility restrictions. Individuals of certain groups reduce their mobility, i.e., their frequency of visits to sites by a certain proportion. Panel (a) shows the number of infected over time for different levels of mobility restrictions. Panel (b) shows the effective reproduction number during the phase of exponential growth of the number of infected. Panel (c) shows the reduction of peak hospitalizations compared to a scenario without mobility restrictions. Points and lines represent the mean over 100 rollouts of the simulation. Error bars and shaded regions correspond to plus and minus one and two standard deviations, respectively.

Orthogonal to various strategies that aim at reducing the number of contacts, the promise of digital contact tracing has been to achieve fine-grained epidemic control without severe societal or economical restrictions [37]. In this section, whenever an individual is tested positive, we use contact tracing to identify all of their contacts in the 10 days leading up to the test result (see Section 3.4). If a given contact was longer than 15 minutes—the time threshold used by the national COVID-19 tracing apps in, e.g., France, Germany, Switzerland, and the United Kingdom [26, 35, 59, 69]—the contact person is tested and isolated from everyone in the mobility model for 14 days.

We analyze the effectiveness of digital contact tracing in combination with various degrees of mobility restrictions for the entire population at different digital tracing adoption levels. The findings shown in Figure 5(a) illustrate that the adoption of the digital tracing system and the activity reduction due to social distancing have a complementary relationship in reducing the cumulative number of infections, as already argued by previous work [25]. Furthermore, the results suggest that, while contact tracing can provide a significant contribution to the mitigation of an epidemic, even at high adoption levels of 75%, it requires a combination with activity reductions of 25% and above in order to achieve epidemic control ($R_t < 1$). The effective reproduction number R_t shown in Figure 5(b) decreases over time at a constant adoption level due to the growing number of recovered individuals in the population.

6 RELATED WORK

Our work builds upon previous work on compartmental epidemiological modeling, human mobility models, and temporal point processes. Most of the classical epidemiological literature has focused on developing population models [45], unable to capture heterogeneous transmission

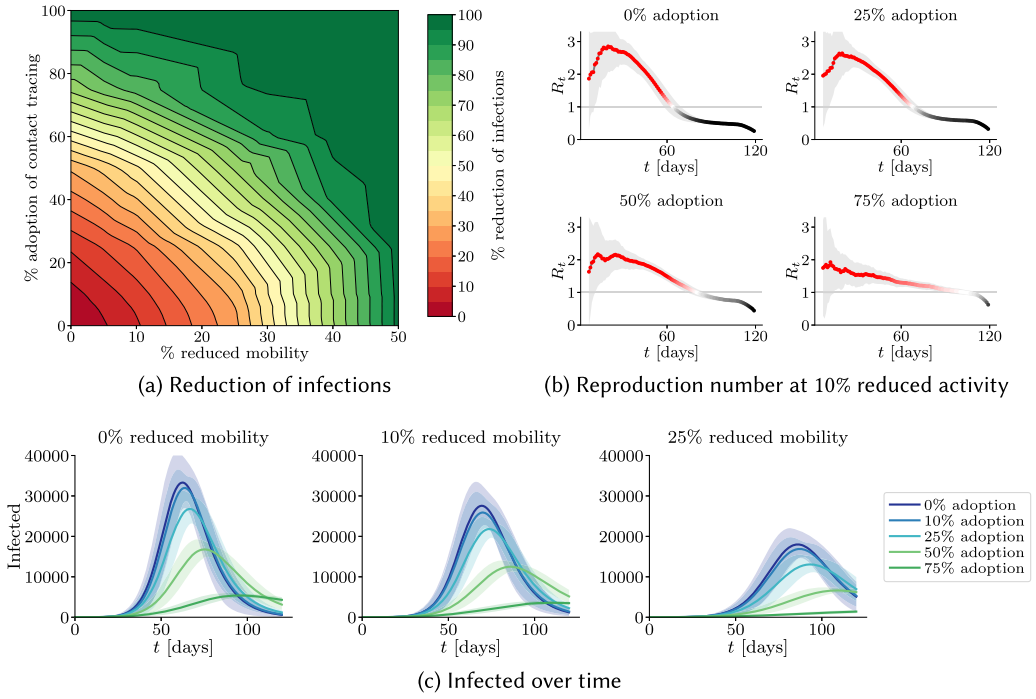


Fig. 5. Contact tracing. A certain proportion of the population has adopted a digital contact tracing system and, in addition, everyone reduces their visit frequency to sites by a certain proportion. Quarantine measures are imposed on traceable individuals who have been in contact with positively tested individuals for a minimum of 15 minutes. Panel (a) shows the reduction of the cumulative number of infections compared to a scenario without interventional measures for different adoption levels of the tracing technology and various degrees of mobility restrictions. Panel (b) shows the reproduction number over time for different adoption levels at a fixed mobility reduction of 10%. Panel (c) shows the number of infected individuals over time for different scenarios. Lines and shaded regions represent the mean and two standard deviations computed over 100 rollouts of the simulation, respectively.

dynamics at the individual level. More recently, there has been research on agent-based epidemic modeling [8, 20, 22, 72, 73], also in the context of COVID-19 [29, 36, 44, 49, 54, 60, 76]. These models predominantly use multi-layer contact networks, discrete time, metapopulation, or Poisson transmission rate assumptions to characterize individual infections, rather than the frequency and duration of each individual's visits to specific sites, as our model does. Notable exceptions are by Aleta et al. [9], who use check-in data of real sites, yet only to configure the layers of a multi-layer contact network, and Ferreti et al. [37], who employ a time-varying transmission rate, but average over individuals who infect few or many others. Chang et al. [23] consider specific points of interest in US cities but only model transmission dynamics of metapopulations of up to 3,000 people rather than among single individuals. Ultimately, none of the above models, including these three exceptions, can be faithfully used to characterize the dispersion of the number of individual infections during a visit, or to straightforwardly study the course of a disease under fine-grained intervention policies such as, e.g., contact tracing or testing. As a result, these models have not been useful for studying conditions under which hotspots emerge [21, 40], analyzing measures to prevent SSEs [10], or predicting where infection hotspots are most likely to occur [83].

The literature on human mobility models has a rich history, which has been extensively reviewed by Barcosa et al. [14]. In our experiments, the spatial distribution of visits of our *event-based* “check-in” mobility model follows the gravity model [85]. Analogous to previous COVID-19 research, these visits are synthetically generated in each simulation [15, 31, 48]. However, our formulation is not restricted to this specific choice and one could think of designing event-based mobility models with a spatial distribution of visits following, e.g., the radiation model [65] or population-weighted opportunities model [81]. That said, the configuration of visit types, frequencies, and durations are specific to event-based models like ours, where the existing body of work on mobility models provides only very limited guidance [47, 82].

Finally, there has been a flurry of work on temporal point process modeling in the machine learning literature in recent years [70, 80, 84, 86]. They have been particularly successful in predicting information propagation in social networks and the web, where they have achieved state-of-the-art performances [30, 34]. However, the development of compartmental epidemiological models based on temporal point processes has been lacking.

7 DISCUSSION

Motivated by multiple lines of evidence that strongly suggest for infection hotspots to play a key role in the transmission dynamics of COVID-19, we have introduced a spatiotemporal epidemic model that explicitly represents sites where infections occur and hotspots may emerge. Through a case study that used fine-grained demographic data, site locations, mobility data as well as COVID-19 case data from Bern, Switzerland, we have demonstrated that our model can allow individuals and policy-makers to make more effective decisions concerning the implementation of containment measures, contact tracing, and testing—at the individual level and in the presence of overdispersion. To facilitate this, we have released an easy-to-use implementation of the entire framework necessary to perform experiments for any desired region [57].

While the purpose of this work does not lie in providing mechanistic forecasts, we have shown that an identifiable pair of only two fitted parameters, the transmission rates at sites and households, provides reasonable predictiveness over our estimation window. Importantly, we find that our epidemiological model empirically exhibits overdispersion in the number of secondary infections, which suggests that our formulation characterizes the transmission dynamics at infection hotspots—an epidemiological driver that effective containment measures would demand preventing [10, 21]. In this context, we do not intend to argue that our approach allows for a more accurate fit to aggregate case data than existing meta-population or network-based compartmental models. Instead, our results in Section 5 demonstrate that we are able to formally model fine-grained interventions and perform analyses that would not be possible within the mathematical formulation used by existing meta-population models.

In this work, we have used fine-grained demographic data and site locations to configure our mobility model. However, if contact tracing data become accessible to researchers, we believe that the variance of our predictions could be lowered and that it would be possible to use our framework to identify areas with higher risk of infection in real time. Beyond legal compliance and gaining societal acceptance, the use of epidemic models with high spatiotemporal resolution such as ours should respect each individual’s privacy. It is hence important to highlight that, both during parameter estimation and contact tracing, we only need to compute the contact *duration* of individuals with an infected person—the identity of the infected person is not required. As a result, there are reasons to believe that such computations can be made in a decentralized and privacy-preserving manner [68]. Ultimately, although our model has greater resolution than many of those in use today, its predictions can only be faithfully considered when being aware of the high variance observed across random realizations.

APPENDICES

A HOUSEHOLD EXPOSURES

If information about households $\mathcal{H}(i)$ that each individual $i \in \mathcal{V}$ belongs to is available, one can account for exposures within households analogously to exposures at sites \mathcal{S} by adding an additional rate $\lambda_{\mathcal{H}(i)}(t)$ to the conditional intensity function $\lambda_i(t)$ of the exposure counting process $N_i(t)$:

$$\lambda_{\mathcal{H}(i)}(t) = S_i(t) \xi \sum_{j \in \mathcal{H}(i) \setminus i} \int_{t-\delta}^t K_{i,j}^{\mathcal{H}}(\tau) \gamma e^{-\gamma(t-\tau)} d\tau, \quad (14)$$

where

$$K_{i,j}^{\mathcal{H}}(\tau) = \left(I_j^s(\tau) + I_j^p(\tau) + \mu I^a(\tau) \right) \prod_{k \in \mathcal{S}} (1 - P_{i,k}(\tau))(1 - P_{j,k}(\tau)), \quad (15)$$

where $\xi \geq 0$ is the base transmission rate within households. This intensity function models our assumption that individuals within a household are in contact as long as they are not visiting any site.

Exposure events caused by $\lambda_{\mathcal{H}(i)}(t)$ can be sampled analogously to the principles for sampling exposure times introduced in Section 4.1. Their superposition with exposures at sites is handled by the priority queue.

B EMPIRICAL PROBABILITY OF EXPOSURE

The exposure risk of others caused by an infectious individual can be computed under our model and empirically approximated using location or contact data, e.g., from (manual) contact tracing. Specifically, the probability of exposure $\hat{p}_{i \leftarrow j}(t_0, t_f)$ during a time window $[t_0, t_f]$ associated with j in the process $N_i(t)$ is given by

$$\hat{p}_{j \leftarrow i}(t_0, t_f) = 1 - \exp \left(-K_{j,i}^{\text{risk}}(t_0, t_f) \right), \quad (16)$$

with

$$K_{j,i}^{\text{risk}}(t_0, t_f) = \sum_{k \in \mathcal{S}} \beta_k \int_{t_0}^{t_f} P_{i,k}(t') \int_{t'-\delta}^{t'} P_{j,k}(\tau) \gamma e^{-\gamma(t'-\tau)} d\tau dt', \quad (17)$$

and follows from the survival probability in a temporal point process [27]. The estimated probability of exposure is conservatively high by assuming that all contacts are (pre-)symptomatic and not considering a possible scaling of μ for asymptomatic individuals. See Section 3.2.

C STATE VARIABLE INITIALIZATION

During the parameter estimation period, it is necessary to specify initial epidemiological conditions $\mathbb{S}(0)$ that are consistent with the COVID-19 case data used in the objective. To this end, we set the number of initially symptomatic individuals $I_{\text{init}}^s = \sum_{i \in \mathcal{V}} I_i^s(0)$ equal to the real observed COVID-19 cases in a region at the start date, or scaled proportionally to the population size in an administrative region, and set all to be positively tested. Based on the above, we seed $I_{\text{init}}^a = \alpha_a / (1 - \alpha_a) I_{\text{init}}^s$ individuals to be initially asymptomatic to obtain a proportion of recently estimated $\alpha_a = 0.4$ asymptomatic seeds [37, 53, 62]. Assuming that infectious individuals have exposed R_0 others on average, we seed $E_{\text{init}} = R_0 (I_{\text{init}}^a + I_{\text{init}}^s)$ initially exposed individuals, using recent estimates of the basic reproduction number of approximately $R_0 \approx 2.0$ [37, 67, 78]. In any simulation done for parameter estimation, $E_i(0)$, $I_i^a(0)$, $I_i^s(0)$ are seeded uniformly at random following the above heuristic counts. Neither asymptomatic nor symptomatic seeds cause further exposures, and for simplicity, no other states are initially seeded.

D SAMPLING PROCEDURE

D.1 Algorithms

ALGORITHM 2: Sampling algorithm for generating a trajectory of the epidemiological states $\mathbb{S}(t)$ ⁵

Input: Initial conditions $\mathbb{S}(0)$, mobility traces $P_{i,k}(t)$, parameters $\gamma, \delta, \alpha_a, \alpha_b, \alpha_h, \mu$ and β_k , rates $\lambda_{(\cdot)}(t)$ of the counting processes

- 1: $t_{\text{now}} \leftarrow 0, S_i \leftarrow 1, Q \leftarrow$ priority queue processing in temporal order of events
- 2: **for** all $i \in \mathcal{V}$ s.t. $S_i = 0$ **do** ▷ Initial conditions
- 3: Push initial state transition $(0, \cdot, i, \emptyset, \emptyset)$ to Q
- 4: **while** Q not empty **do** ▷ Priority queue Q contains (time, transition, i , infector, site) events
- 5: $(t_{\text{now}}, e, i, j, k) \leftarrow$ pop earliest from Q
- 6: **if** e is dE **and** $R_j(t_{\text{now}}) = 0$ **and** $D_j(t_{\text{now}}) = 0$ **and** $S_i = 1$ **then**
- 7: **if** $\text{INTERVENTIONS}(i, j, k, t_{\text{now}})$ **then** ▷ Reject event and re-sample arrival time of event?
- 8: Call Algorithm 3 with $(P, j, i, t_{\text{now}}, r = 1 - (1 - \mu)I_j^a(t_{\text{now}}))$
- 9: **else** ▷ Person i exposed by infector j
- 10: $E_i \leftarrow 1, S_i \leftarrow 0, \Delta_M \sim \text{Expo}(\lambda_M(t_{\text{now}})), u \sim \text{Unif}(0, 1)$
- 11: **if** $u \leq \alpha_a$ **then**
- 12: Push $(t_{\text{now}} + \Delta_M, dI^a, i, \emptyset, \emptyset)$ event to Q
- 13: **else**
- 14: Push $(t_{\text{now}} + \Delta_M, dI^p, i, \emptyset, \emptyset)$ event to Q
- 15: **else if** e is dI^p **then** ▷ Person i presymptomatic
- 16: $I_i^p \leftarrow 1, E_i \leftarrow 0, \Delta_Z \sim \text{Expo}(\lambda_W(t_{\text{now}}))$
- 17: Push $(t_{\text{now}} + \Delta_Z, dI^s, i, \emptyset, \emptyset)$ event to Q
- 18: **for** u such that $S_u = 1$ **do**
- 19: Call Algorithm 3 with arguments $(P, i, u, t_{\text{now}}, r = 1)$
- 20: **else if** e is dI^s **then** ▷ Person i symptomatic
- 21: $I_i^s \leftarrow 1, I_i^p \leftarrow 0, u, v \sim \text{Unif}(0, 1)$
- 22: **if** $u \leq \alpha_h$ **then**
- 23: $\Delta_Y \sim \text{Expo}(\lambda_Y(t_{\text{now}})), \text{Push } (t_{\text{now}} + \Delta_Y, dH, i, \emptyset, \emptyset)$ to Q
- 24: **if** $v \leq \alpha_b$ **then**
- 25: $\Delta_Z \sim \text{Expo}(\lambda_Z(t_{\text{now}})), \text{Push } (t_{\text{now}} + \Delta_Z, dD, i, \emptyset, \emptyset)$ to Q
- 26: **else**
- 27: $\Delta_R \sim \text{Expo}(\lambda_{R^s}(t_{\text{now}})), \text{Push } (t_{\text{now}} + \Delta_R, dR, i, \emptyset, \emptyset)$ to Q
- 28: **else if** e is dI^a **then** ▷ Person i asymptomatic
- 29: $I_i^a \leftarrow 1, E_i \leftarrow 0, \Delta_R \sim \text{Expo}(\lambda_{R^a}(t_{\text{now}}))$
- 30: Push $(t_{\text{now}} + \Delta_R, dR, i, \emptyset, \emptyset)$ event to Q
- 31: **for** u such that $S_u = 1$ **do**
- 32: Call Algorithm 3 with arguments $(P, i, u, t_{\text{now}}, r = \mu)$
- 33: **else if** e is dH **then** ▷ Person i hospitalized
- 34: $H_i \leftarrow 1$
- 35: **else if** e is dR **then** ▷ Person i resistant
- 36: $R_i \leftarrow 1, I_i^a \leftarrow 0, I_i^s \leftarrow 0, H_i \leftarrow 0$
- 37: **else if** e is dD **then** ▷ Person i died
- 38: $D_i \leftarrow 1, I_i^s \leftarrow 0, H_i \leftarrow 0$

⁵For simplicity, we omit details about the procedure $\text{INTERVENTIONS}(i, j, k, t)$, which applies thinning as explained in Section 4.1 for possible interventional measures. Details can be found in our publicly available implementation [57].

ALGORITHM 3: Push the next event of individual i exposing individual j in time window $[t, t_{\max})$ by considering only the contribution $\lambda_{i \rightarrow j}(t)$ in (11) as described in Section 4.1.

Input: P, i, j, t, r

```

1: procedure INCONTACT( $u, v, \tau$ )
2:   return True if  $\exists k \in \mathcal{S}$  s.t.  $P_{u,k}(\tau) = 1$  and  $\exists \tau' \in [\tau - \delta, \tau]$  s.t.  $P_{v,k}(\tau') = 1$ 
   else return False
3: procedure NEXTCONTACT( $u, v, \tau$ )
4:   return  $\min_{\tau' > \tau} \tau'$  s.t. INCONTACT( $u, v, \tau'$ )
5: procedure WILLBEINCONTACT( $u, v, \tau$ )
6:   return True if  $\exists \tau' \in [\tau, t_{\max}]$  s.t. INCONTACT( $u, v, \tau'$ )
   else return False
7:  $\tau \leftarrow t$ 
8: while WILLBEINCONTACT( $j, i, \tau$ ) do                                      $\triangleright$  Thinning loop
9:    $b \leftarrow$  INCONTACT( $j, i, \tau$ )
10:  if not  $b$  then
11:     $\tau \leftarrow$  NEXTCONTACT( $j, i, \tau$ )
12:     $\Delta E_j \sim \text{Expo}(\max_k \{\beta_k\} r \int_{\tau-\delta}^{\tau} \gamma e^{-\gamma(\tau-v)} dv)$ 
13:     $\tau \leftarrow \tau + \Delta E_j$ 
14:    if INCONTACT( $j, i, \tau$ ) then
15:       $k \leftarrow$  site of contact
16:       $p \leftarrow (\beta_k r \int_{\tau-\delta}^{\tau} \gamma e^{-\gamma(\tau-v)} P_{i,k}(v) dv) / (\max_k \{\beta_k\} r \int_{\tau-\delta}^{\tau} \gamma e^{-\gamma(\tau-v)} dv)$ 
17:       $u \sim \text{Unif}(0, 1)$ 
18:      if  $u \leq p$  then                                                  $\triangleright$  Accept/reject sampled event time
19:        Push  $(\tau, dE, j, i, k)$  event to  $Q$ 
20:      break

```

D.2 Proof of Proposition 1

The sampling algorithm and its subroutine are formally defined in Algorithms 2 and 3. In the following, we assume that: (i) a given individual $i \in \mathcal{V}$ makes $O(t_{\max})$ visits to sites \mathcal{S} over the horizon $[0, t_{\max})$; (ii) the mobility model is *sparse*, i.e., every individual $i \in \mathcal{V}$ has $O(1)$ unique contact persons; and, (iii) there are no containment measures. This implies that there are a total number of $O(t_{\max})$ contact windows of i with all other individuals $j \in \mathcal{V}$. Following our implementation [57], we assume that the mobility traces $P_{i,k}(t)$ are given as an unsorted list of time intervals $[t_0, t_1]$, where each time interval indicates a visit of an individual $i \in \mathcal{V}$ to a site $k \in \mathcal{S}$ during the simulated time period.

Event queue. In any possible trajectory of the epidemiological state variables, there is a constant number of events *not* concerning exposure that can be pushed to the event queue Q per individual. This is because every individual transitions through at most a finite set of states. In addition, since by assumption (ii) the mobility model is sparse, there is a constant number of exposure events caused by and thus pushed per individual. Thus, the overall number of events pushed to the event queue throughout the simulation is $O(|\mathcal{V}|)$. This is an upper bound on the size of the queue at any point in the simulation. Using the standard heap implementation of a priority queue, pushing to and popping from the temporally-sorted event queue Q hence incur cost $O(\log(|\mathcal{V}|))$ in the worst case.

Preprocessing of contacts. Sampling exposures caused by an infectious individual i relies on querying the contacts with other individuals j by checking their overlapping visits to sites \mathcal{S} . To do this efficiently, the mobility traces are preprocessed into efficient interval data structures called

interval trees.⁶ For this, we initialize two dictionaries that store two kinds of interval trees, visits by individuals and visits to sites, respectively. Both dictionaries are populated by iterating once over all $O(t_{\max}|\mathcal{V}|)$ site visits in the simulated period. For each visit, its time interval is inserted both into the tree of visits by the corresponding individual $i \in \mathcal{V}$ as well as the tree associated to the site $k \in \mathcal{S}$. Then, by assumption (i), the interval trees stratified by *individual* have size $O(t_{\max})$ and intervals do not overlap by construction. Moreover, the interval trees stratified by *site* contain $O(t_{\max}|\mathcal{V}|)$ visits and, by assumption (ii), any interval overlaps with $O(1)$ others. Thus, the total time incurred for constructing all visit interval trees is $O(t_{\max}|\mathcal{V}| \log(t_{\max}|\mathcal{V}|))$.

Using these two sets of *visit* interval trees, we build a collection of $O(1)$ *contact* interval trees for each individual $i \in \mathcal{V}$. These contain the contact windows from i to each of its unique contact persons j .⁷ To generate the contact trees for i , we iterate over all $O(t_{\max})$ visits of i . For each visit, we query the interval tree of the visited site in time $O(\log(t_{\max}|\mathcal{V}|))$ to retrieve the $O(1)$ contact persons j during that visit. Given this, we update the individual contact interval tree from i to j in time $O(\log(t_{\max}))$. Like individual visit traces, the contact intervals do not overlap by construction. The overall preprocessing cost remains $O(t_{\max}|\mathcal{V}| \log(t_{\max}|\mathcal{V}|))$.

Handling events. The backbone of the sampling procedure in Algorithm 2 consists of processing state transition events in the temporal order. All generic state transitions in the model, i.e., those *not* transitioning to an infectious state, consist of updating the correct indicator variables of the corresponding individual i or discarding events that became invalid due to thinning in constant time. In addition, we push the next state transition of i to Q , which takes time $O(\log(|\mathcal{V}|))$. Since there are $O(|\mathcal{V}|)$ generic events in the worst case, handling all of them takes an overall time of $O(|\mathcal{V}| \log(|\mathcal{V}|))$.

When an individual i first transitions to an *infectious* state, i.e., the presymptomatic $I_i^p = 1$ or asymptomatic $I_i^a = 1$ state, an additional time cost is incurred because we sample the times of the exposure events *caused* by i to all of its unique contact persons j in the future. This corresponds to calls of Algorithm 3, where we continually iterate over all $O(t_{\max})$ contact windows i has with j after some time t *until the first valid exposure event is sampled*. Specifically, we sample a next time t as $t \leftarrow t + \tau$ with $\tau \sim \text{Expo}(\lambda_{\max})$. If i is still in contact with j at time t , and if the event is not rejected using thinning due to a lower site-specific exposure rate β_k or asymptomatic infectiousness γ , the exposure time is valid and we push the event to Q in time $O(\log(|\mathcal{V}|))$. Otherwise, we repeat. Since there are at most $O(t_{\max})$ contact windows of i with j , each query to INCONTACT as formalized in Algorithm 3 incurs time $O(\log(t_{\max}))$ using the interval tree.

Let U be the random variable representing the runtime of Algorithm 3 incurred by *one* contact window from i to j . In addition, let $q \in (0, 1)$ be the probability that a given thinning sample gets accepted. By the memoryless property of thinning, the expected value of U is given by

$$\begin{aligned} \mathbb{E}[U] &= O(\log(t_{\max})) + q O(\log(|\mathcal{V}|)) + (1 - q) \mathbb{E}[U] \\ &= \sum_{n=0}^{\infty} (1 - q)^n \left(O(\log(t_{\max})) + q O(\log(|\mathcal{V}|)) \right) = O(\log(|\mathcal{V}|)) + \frac{1}{q} O(\log(t_{\max})). \end{aligned} \quad (18)$$

In the worst case, thinning is done for all $O(t_{\max})$ contact windows from i to j until an exposure event time gets accepted. Overall, Algorithm 3 is called $O(1)$ times per individual. Thus, the processing of state transitions to infectious states of all $O(|\mathcal{V}|)$ infectious individuals incurs an additional

⁶An interval tree containing n intervals allows for $O(\log n)$ insertion time. Using binary search, retrieving the subset of stored intervals that intersect with a query interval $[t_0, t_1]$ takes time $O(m + \log n)$, where m is the number of intersecting intervals.

⁷We denote a contact as being *from* i *to* j to be precise about non-contemporaneous infections (cf. Equation (6)). There is an exposure-relevant contact from i to j if i left a site less than δ -time before j arrived.

overall cost of $O(|\mathcal{V}|t_{\max}(\log(|\mathcal{V}|) + \frac{1}{q}t_{\max}\log(t_{\max})))$. This also accounts for the cost of sampling household exposures, which can be viewed as visits to an additional site with an additional set of $O(1)$ household contacts. We note that q is a constant that depends only on the exposure rate of the epidemiological model, and any lower bound thereof across sites and individuals suffices for Equation (18).

Expected runtime. Combining the preprocessing cost, the handling of all generic state transitions, and the handling of transitions to infectious states, Algorithm 2 has a total expected runtime of

$$O\left(|\mathcal{V}|\left(t_{\max}\log(t_{\max}|\mathcal{V}|) + \frac{1}{q}t_{\max}\log(t_{\max})\right)\right). \quad (19)$$

□

Table 2. Epidemiological Model Parameters in Units of Days

Counting process	Starts when	$\log \mathcal{N}$ parameters	Source
$M_i(t)$	$dE_i(t) = 1$	$(0.9470, 0.6669)^\dagger$	[52]
$R_i^s(t), R_i^a(t)$	$dI_i^s(t) = 1$	$(2.6365, 0.0713)^\ddagger$	[43, 77, 78]
$W_i(t)$	$dI_i^p(t) = 1$	$(0.7463, 0.4161)^\ddagger$	[43]
$Y_i(t)$	$dI_i^s(t) = 1$	$(1.9358, 0.1421)^\ddagger$	[75]
$Z_i(t)$	$dI_i^s(t) = 1$	$(2.5620, 0.0768)^\ddagger$	[55]

	Value	Description	Source
μ	0.55	relative asymptomatic transmission rate	[54]
γ	$6.3013 h^{-1}$	decay of infectiousness at sites [§]	[71]
δ	$0.3654 h$	non-contact contamination window [¶]	[71]
α_a	0.4	proportion of asymptom. individuals	[37, 53, 62]

Hospitalization and Fatality Rates α_h and α_b Mentioned in the Main Text are Estimated from COVID-19 Case Data in the Region and is Age-dependent. Log-normal Parameters Denote the Underlying Normal Mean and Standard Deviation.

[†]Incubation period from [52], corrected by presymptomatic infectiousness [43].

[‡]Approximate log-normal parameters constructed because COVID-19 literature results only reported using mean or median time estimates.

[§]Assumes that transmission decays with a half-life 10 times shorter than estimated for aerosols under laboratory conditions.

[¶]For computational purposes, set from γ by the time when rate of transmission drops below 10% after leaving a site.

Table 3. Assumed Mean Number of Visits Per Week Per Site
Type by Individuals of Different Age Groups for Our
Event-based Gravity Mobility Model [85]

Age group	Education	Social	Transport.	Work	Groceries
0–4	5	1	-	-	-
5–14	5	2	3	-	-
15–34	2	2	3	3	1
35–59	-	2	1	5	1
60–79	-	3	2	-	1
80+	-	2	1	-	1

See Section 5.1.

Table 4. Summary and Estimated Parameters for Towns and Regions
Studied in Germany and Switzerland

Region	Country	$ \mathcal{V} $	$ \mathcal{S} $	Estimation period [†]	Lockdown	β	ξ
Bern	CH	133,790	2,174	03/06–05/10	03/16	0.0337	0.0038
Tübingen	GER	90,539	1,446	03/12–05/03	03/23	0.0402	0.0664
Canton Jura	CH	73,416	729	03/09–05/10	03/16	0.0131	0.0080
Rheingau-Taunus-Kreis	GER	187,163	2,352	03/10–05/03	03/23	0.0010	0.0500
Kaiserslautern	GER	104,044	1,525	03/15–05/03	03/23	0.0279	0.0061

Recall that β Denotes the Individual Transmission Rate at Public Sites and ξ in Households Estimated as Described in Sections 4.2 and 5.1.

[†] Chosen such that a given region had approximately five to ten confirmed COVID-19 cases, allowing for non-degenerate and comparable initial conditions. Dates are in 2020.

E ESTIMATION RESULTS FOR ADDITIONAL REGIONAL MODELS

Figure 6 summarizes the parameter estimation results for four additional regions in Germany and Switzerland: the cities of Tübingen and Kaiserslautern as well as the Canton of Jura and the district Rheingau-Taunus. As for the model of Bern, the estimation procedure was executed as described in Section 5.1. Table 2 lists the estimated optimal parameters as well as additional details about each regional model.

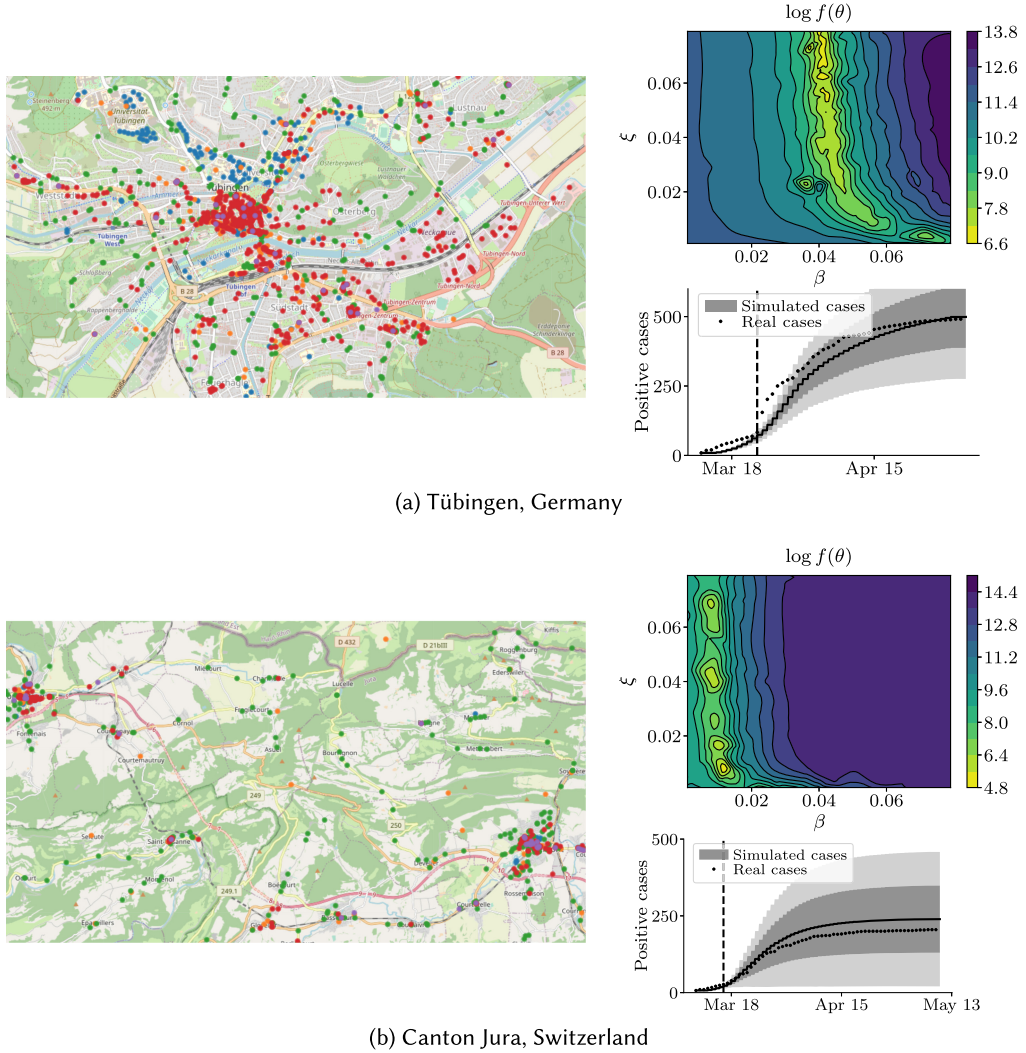
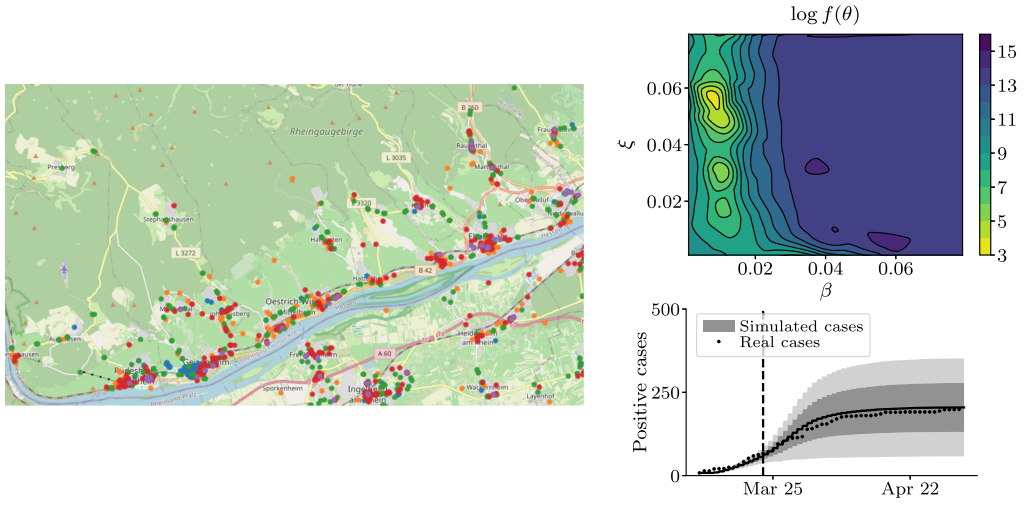
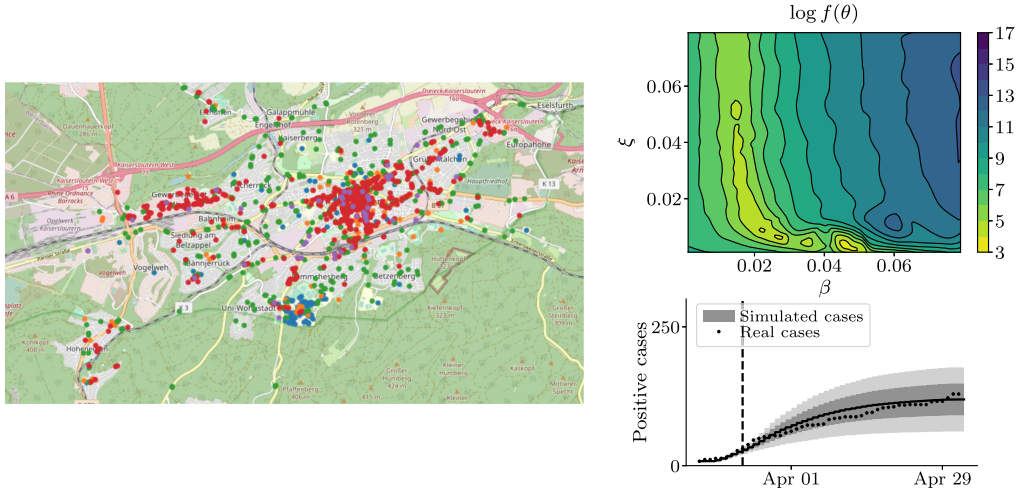


Fig. 6. Transmission parameter estimation for additional regional models in Germany and Switzerland. In each panel, the left plot shows the sites of the mobility model with colors used as in Figure 1. The right plots show the contour plot of the log objective in Equation (13) as a function of the transmission rates β and ξ as well as the predicted and real cumulative cases under the estimated parameters, analogous to Figure 2.



(c) Rheingau-Taunus-Kreis, Germany



(d) Kaiserslautern, Germany

Fig. 6. Continued

ACKNOWLEDGMENTS

We thank the Robert-Koch-Institute, OpenStreetMaps, Google, and Facebook for providing data to make this work possible. We thank Brian Karrer from Facebook for his insightful comments and suggestions regarding Bayesian optimization, Kevin Murphy, Yusef Shafi and others from Google for helpful discussions, and Yannik Schaelte for useful comments on a preliminary version of this work. We thank Cansu Culha and the Stanford Future Bay Initiative as well as Pavol Harar from the University of Vienna for working with us to improve our publicly available implementation.

REFERENCES

- [1] 2020. Facebook Data for Good. Retrieved 30 April 2021 from <https://dataforgood.fb.com>.
- [2] 2020. Johns Hopkins University & Medicine: coronavirus resource center. Retrieved 30 April 2021 from <https://coronavirus.jhu.edu/>.

- [3] 2020. The Korean clusters: How coronavirus cases exploded in South Korean churches and hospitals. Retrieved from <https://graphics.reuters.com/CHINA-HEALTH-SOUTHKOREA-CLUSTERS/0100B5G33SB/index.html>.
- [4] 2020. OpenStreetMap. Retrieved from <https://www.openstreetmap.org/>.
- [5] Odd O. Aalen, Ørnulf Borgan, and Håkon K. Gjessing. 2008. *Survival and Event History Analysis*. Springer New York.
- [6] Dillon C. Adam, Peng Wu, Jessica Y. Wong, Eric H. Y. Lau, Tim K. Tsang, Simon Cauchemez, Gabriel M. Leung, and Benjamin J. Cowling. 2020. Clustering and superspreading potential of SARS-CoV-2 infections in Hong Kong. *Nature Medicine* 26, 11 (2020), 1714–1719.
- [7] N. Ahmed, R. A. Michelin, W. Xue, S. Ruj, R. Malaney, S. S. Kanhere, A. Seneviratne, W. Hu, H. Janicke, and S. K. Jha. 2020. A survey of COVID-19 contact tracing apps. *IEEE Access* 8 (2020), 134577–134601. DOI : <https://doi.org/10.1109/ACCESS.2020.3010226>
- [8] Hyoungh Jun Ahn and Babak Hassibi. 2013. Global dynamics of epidemic spread over complex networks. In *Proceedings of the 52nd IEEE Conference on Decision and Control*. IEEE, 4579–4585.
- [9] Alberto Aleta, David Martin-Corral, Ana Pastore y Piontti, Marco Ajelli, Maria Litvinova, Matteo Chinazzi, Natalie E. Dean, M. Elizabeth Halloran, Ira M. Longini Jr, Stefano Merler, et al. 2020. Modelling the impact of testing, contact tracing and household quarantine on second waves of COVID-19. *Nature Human Behaviour* 4, 9 (2020), 964–971.
- [10] Benjamin M. Althouse, Edward A. Wenger, Joel C. Miller, Samuel V. Scarpino, Antoine Allard, Laurent Hébert-Dufresne, and Hao Hu. 2020. Stochasticity and heterogeneity in the transmission dynamics of SARS-CoV-2. arXiv:2005.13689. Retrieved from <https://arxiv.org/abs/2005.13689>.
- [11] Raul Astudillo and Peter I. Frazier. 2019. Bayesian optimization of composite functions. In *International Conference on Machine Learning*. PMLR, 354–363.
- [12] Siva Athreya, Nitya Gadhiwala, and Abhiti Mishra. 2021. Effective reproduction number and dispersion under contact tracing and lockdown on COVID-19 in Karnataka. *Journal of the Indian Society for Probability and Statistics* 22, 2 (2021), 319–342.
- [13] Maximilian Balandat, Brian Karrer, Daniel Jiang, Samuel Daulton, Ben Letham, Andrew G. Wilson, and Eytan Bakshy. 2020. BoTorch: a framework for efficient Monte-Carlo Bayesian optimization. *Advances in Neural Information Processing Systems* 33 (2020), 21524–21538.
- [14] Hugo Barbosa, Marc Barthelemy, Gourab Ghoshal, Charlotte R. James, Maxime Lenormand, Thomas Louail, Ronaldo Menezes, José J. Ramasco, Filippo Simini, and Marcello Tomasini. 2018. Human mobility: Models and applications. *Physics Reports* 734 (2018), 1–74.
- [15] Per Block, Marion Hoffman, Isabel J. Raabe, Jennifer Beam Dowd, Charles Rahal, Ridhi Kashyap, and Melinda C. Mills. 2020. Social network-based distancing strategies to flatten the COVID-19 curve in a post-lockdown world. *Nature Human Behaviour* 4, 6 (2020), 588–596.
- [16] Eric Brochu, Vlad M. Cora, and Nando de Freitas. 2010. A Tutorial on Bayesian Optimization of Expensive Cost Functions, with Application to Active User Modeling and Hierarchical Reinforcement Learning. arXiv:1012.2599. Retrieved from <https://arxiv.org/abs/1012.2599>.
- [17] Bundesamt für Gesundheit, Schweiz. 2020. BAG COVID19 situation summary. Retrieved 30 April 2021 from <https://www.bag.admin.ch/bag/en/home/krankheiten/ausbrueche-epidemien-pandemien/aktuelle-ausbrueche-epidemien/novel-cov/situation-schweiz-und-international.html>.
- [18] Bundesamt für Statistik, Schweiz. 2020. Households. Retrieved 30 April 2021 from <https://www.bfs.admin.ch/bfs/de/home/statistiken/bevoelkerung/stand-entwicklung/haushalte.html>.
- [19] Bundesländer der Bundesrepublik Deutschland. 2020. Corona-Regelungen in den Bundesländern. Retrieved 30 April 2021 from <https://www.bundesregierung.de/breg-de/themen/coronavirus/corona-bundeslaender-1745198>.
- [20] Eric Cator and Piet Van Mieghem. 2012. Second-order mean-field susceptible-infected-susceptible epidemic threshold. *Physical Review E* 85, 5 (2012) 056111.
- [21] Muge Cevik et al. 2020. SARS-CoV-2 transmission dynamics should inform policy. *Clinical Infectious Diseases* 73, Supplement_2 (2021), S170–S176.
- [22] Deepayan Chakrabarti, Yang Wang, Chenxi Wang, Jurij Leskovec, and Christos Faloutsos. 2008. Epidemic thresholds in real networks. *ACM Transactions on Information and System Security (TISSEC)* 10, 4 (2008), 1–26.
- [23] Serina Chang, Emma Pierson, Pang Wei Koh, Jaline Gerardin, Beth Redbird, David Grusky, and Jure Leskovec. 2021. Mobility network models of COVID-19 explain inequities and inform reopening. *Nature* 589, 7840 (2021), 82–87.
- [24] William G. Cochran. 1954. Some methods for strengthening the common χ^2 tests. *Biometrics* 10, 4 (1954), 417–451.
- [25] Sebastian Contreras, Jonas Dehning, Matthias Loidolt, Johannes Zierenberg, F. Paul Spitzner, Jorge H. Urrea-Quintero, Sebastian B. Mohr, Michael Wilczek, Michael Wibral, and Viola Priesemann. 2021. The challenges of containing SARS-CoV-2 via test-trace-and-isolate. *Nature Communications* 12, 1 (2021), 1–13.
- [26] Corona-Warn-App Team. 2020. Corona-Warn-App: Documentation. Retrieved 30 April 2021 from <https://github.com/corona-warn-app/cwa-documentation/>.

- [27] Abir De, Utkarsh Upadhyay, and Manuel Gomez-Rodriguez. 2019. *Temporal Point Processes*. Technical Report. Saarland University.
- [28] Der Schweizerische Bundesrat. 2020. Verordnung 2 über Massnahmen zur Bekämpfung des Coronavirus (COVID-19) (COVID-19-Verordnung 2). Retrieved 30 April 2021 from <https://www.admin.ch/opc/de/classified-compilation/20200744/index.html>.
- [29] Laura Di Domenico, Giulia Pullano, Chiara E. Sabbatini, Pierre-Yves Boëlle, and Vittoria Colizza. 2020. Impact of lockdown on COVID-19 epidemic in Île-de-France and possible exit strategies. *BMC Medicine* 18, 1 (2020), 1–13. Expected impact of school closure and telework to mitigate COVID-19 epidemic in France.
- [30] Nan Du, Hanjun Dai, Rakshit Trivedi, Utkarsh Upadhyay, Manuel Gomez-Rodriguez, and Le Song. 2016. Recurrent marked temporal point processes: Embedding event history to vector. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 1555–1564.
- [31] Daniel Duque, David P. Morton, Bismark Singh, Zhanwei Du, Remy Pasco, and Lauren Ancel Meyers. 2020. Timing social distancing to avert unmanageable COVID-19 hospital surges. *Proceedings of the National Academy of Sciences* 117, 33 (2020), 19873–19878.
- [32] Akira Endo et al. 2020. Estimating the overdispersion in COVID-19 transmission using outbreak sizes outside China. *Wellcome Open Research* 5 (2020).
- [33] European Center for Disease Prevention and Control. 2020. An overview of the rapid test situation for COVID-19 diagnosis in the EU/EEA. Retrieved 30 April 2021 from <https://www.ecdc.europa.eu/sites/default/files/documents/Overview-rapid-test-situation-for-COVID-19-diagnosis-EU-EEA.pdf>.
- [34] Mehrdad Farajtabar, Yichen Wang, Manuel Gomez-Rodriguez, Shuang Li, Hongyuan Zha, and Le Song. 2017. Coevolve: A joint point process model for information diffusion and network evolution. *The Journal of Machine Learning Research* 18, 1 (2017), 1305–1353.
- [35] Federal Office of Public Health, Switzerland. 2020. How the SwissCovid app issues warnings. Retrieved 30 April 2021 from <https://www.bag.admin.ch/bag/en/home/krankheiten/ausbrueche-epidemien-pandemien/aktuelle-ausbrueche-epidemien/novel-cov/swisscovid-app-und-contact-tracing.html#-1674020142>.
- [36] Neil M. Ferguson, Daniel Laydon, Gemma Nedjati-Gilani, Natsuko Imai, Kylie Ainslie, Marc Baguelin, Sangeeta Bhatia, Adhiratha Boonyasiri, Zulma Cucunubá, Gina Cuomo-Dannenburg, et al. 2020. Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand. (2020).
- [37] Luca Ferretti, Chris Wymant, Michelle Kendall, Lele Zhao, Anel Nurtay, Lucie Abeler-Dörner, Michael Parker, David Bonsall, and Christophe Fraser. 2020. Quantifying SARS-CoV-2 transmission suggests epidemic control with digital contact tracing. *Science* 368, 6491 (2020), eabb6936.
- [38] Ronald Aylmer Fisher. 1992. Statistical methods for research workers. In *Proceedings of the Breakthroughs in statistics*. Springer, 66–70.
- [39] Peter I. Frazier. 2018. A Tutorial on Bayesian Optimization. arXiv:1807.02811. Retrieved from <https://arxiv.org/abs/1807.02811>.
- [40] Medics Sans Frontiers. 2020. Too little, Too late: The unacceptable neglect of the elderly in care homes during the COVID-19 epidemic in Spain. Retrieved 30 April 2021 from <https://www.msf.org/sites/default/files/2020-08/Executive%20summary%20of%20MSF%27s%20report.pdf>.
- [41] Google. 2020. COVID-19 Community Mobility Reports. Retrieved 30 April 2021 from <https://www.google.com/covid19/mobility/>.
- [42] Agus Hasan, Hadi Susanto, Muhammad Firmansyah Kasim, Nuning Nuraini, Bony Lestari, Dessy Triany, and Widyastuti Widyastuti. 2020. Superspreading in early transmissions of COVID-19 in Indonesia. *Scientific Reports* 10, 1 (2020), 1–4.
- [43] Xi He, Eric HY Lau, Peng Wu, Xilong Deng, Jian Wang, Xinxin Hao, Yiu Chung Lau, Jessica Y. Wong, Yujuan Guan, Xinghua Tan, et al. 2020. Temporal dynamics in viral shedding and transmissibility of COVID-19. *Nature Medicine* 26, 5 (2020), 672–675.
- [44] Ralf Herbrich, Rajeev Rastogi, and Roland Vollgraf. 2020. CRISP: A probabilistic model for individual-level COVID-19 infection risk estimation based on contact data. arXiv:2006.04942. Retrieved from <https://arxiv.org/abs/2006.04942>.
- [45] Herbert W. Hethcote. 2000. The mathematics of infectious diseases. *SIAM Review* 42, 4 (2000), 599–653.
- [46] Nazrul Islam, Stephen J. Sharp, Gerardo Chowell, Sharmin Shabnam, Ichiro Kawachi, Ben Lacey, Joseph M. Massaro, Ralph B. D’Agostino, and Martin White. 2020. Physical distancing interventions and incidence of coronavirus disease 2019: Natural experiment in 149 countries. *BMJ* 370 (2020).
- [47] Martin Jankowiak and Manuel Gomez-Rodriguez. 2017. Uncovering the spatiotemporal patterns of collective social activity. In *Proceedings of the 2017 SIAM International Conference on Data Mining*. SIAM, 822–830.
- [48] Omer Karin, Yinon M. Bar-On, Tomer Milo, Itay Katzir, Avi Mayo, Yael Korem, Boaz Dudovich, Eran Yashiv, Amos J. Zehavi, Nadav Davidovitch, et al. 2020. Adaptive cyclic exit strategies from lockdown to suppress COVID-19 and allow economic activity. *medRxiv* (2020).

- [49] Adam J. Kucharski et al. 2020. Effectiveness of isolation, testing, contact tracing, and physical distancing on reducing transmission of SARS-CoV-2 in different settings: A mathematical modelling study. *The Lancet Infectious Diseases* 20, 10 (2020), 1151–1160.
- [50] Land Baden-Württemberg. 2020. Verordnung der Landesregierung über infektionsschützende Massnahmen gegen die Ausbreitung des Virus SARS-Cov-2 (Corona-Verordnung). Retrieved 30 April 2021 from <https://www.baden-wuerttemberg.de/de/service/aktuelle-infos-zu-corona/aktuelle-corona-verordnung-des-landes-baden-wuerttemberg/>.
- [51] Max S. Y. Lau, Bryan Grenfell, Michael Thomas, Michael Bryan, Kristin Nelson, and Ben Lopman. 2020. Characterizing superspreading events and age-specific infectiousness of SARS-CoV-2 transmission in Georgia, USA. *Proceedings of the National Academy of Sciences* 117, 36 (2020), 22430–22435.
- [52] Stephen A. Lauer, Kyra H. Grantz, Qifang Bi, Forrest K. Jones, Qulu Zheng, Hannah R. Meredith, Andrew S. Azman, Nicholas G. Reich, and Justin Lessler. 2020. The Incubation Period of Coronavirus Disease 2019 (COVID-19) From Publicly Reported Confirmed Cases: Estimation and Application. *Annals of Internal Medicine* 172, 9 (2020), 577–582.
- [53] Enrico Lavezzo, Elisa Franchin, Constanze Ciavarella, Gina Cuomo-Dannenburg, Luisa Barzon, Claudia Del Vecchio, Lucia Rossi, Riccardo Manganelli, Arianna Loregian, Nicolo Navarin, et al. 2020. Suppression of COVID-19 outbreak in the municipality of Vo'. *Nature* 584, 7821 (2020), 425–429.
- [54] Ruiyun Li et al. 2020. Substantial undocumented infection facilitates the rapid dissemination of novel Coronavirus (SARS-CoV2). *Science* (2020).
- [55] Natalie M. Linton, Tetsuro Kobayashi, Yichi Yang, Katsuma Hayashi, Andrei R. Akhmetzhanov, Sung-mok Jung, Baoyin Yuan, Ryo Kinoshita, and Hiroshi Nishiura. 2020. Incubation period and other epidemiological characteristics of 2019 novel Coronavirus infections with right truncation. *Journal of Clinical Medicine* 9, 2 (2020), 538.
- [56] James O. Lloyd-Smith, Sebastian J. Schreiber, P. Ekkehard Kopp, and Wayne M. Getz. 2005. Superspreading and the effect of individual variation on disease emergence. *Nature* 438, 7066 (2005), 355–359.
- [57] Lars Lorch, Heiner Kremer, William Trouleau, Stratis Tsirtsis, Aron Szanto, Bernhard Schölkopf, and Manuel Gomez-Rodriguez. 2020. Data and code from “Quantifying the Effects of Contact Tracing, Testing, and Containment Measures in the Presence of Infection Hotspots”. Retrieved from <https://github.com/covid19-model/simulator>.
- [58] Dror Meidan, Nava Schulmann, Reuven Cohen, Simcha Haber, Eyal Yaniv, Ronit Sarid, and Baruch Barzel. 2020. An alternating lock-down strategy for sustainable mitigation of COVID-19. *Nature Communications* 12, 1 (2021), 1–12.
- [59] Ministre des Solidarités et de la Santé. 2020. TousAntiCovid: réponses à vos questions. Retrieved 30 April 2021 from <https://solidarites-sante.gouv.fr/soins-et-maladies/maladies-infectieuses/coronavirus/tousanticovid>.
- [60] Seyed M. Moghadas, Affan Shoukat, Meagan C. Fitzpatrick, Chad RWells, Pratha Sah, Abhishek Pandey, Jeffrey D. Sachs, Zheng Wang, Lauren A. Meyers, Burton H. Singer, et al. 2020. Projecting hospital utilization during the COVID-19 outbreaks in the United States. *Proceedings of the National Academy of Sciences* 117, 16 (2020), 9122–9126.
- [61] Mirco Nanni, Gennady Andrienko, Albert-László Barabási, Chiara Boldrini, Francesco Bonchi, Ciro Cattuto, Francesca Chiaromonte, Giovanni Comandé, Marco Conti, Mark Coté, et al. 2021. Give more data, awareness and control to individual citizens, and they will help COVID-19 containment. *Ethics and Information Technology* 23, 1 (2021), 1–6.
- [62] Hiroshi Nishiura, Tetsuro Kobayashi, Takeshi Miyama, Ayako Suzuki, Sung-mok Jung, Katsuma Hayashi, Ryo Kinoshita, Yichi Yang, Baoyin Yuan, Andrei R Akhmetzhanov, et al. 2020. Estimation of the asymptomatic ratio of novel coronavirus infections (COVID-19). *International Journal of Infectious Diseases* 94 (2020), 154–155.
- [63] Myoung-don Oh, Pyoeng Gyun Choe, Hong Sang Oh, Wan Beom Park, Sang-Min Lee, Jinkyong Park, Sang Kook Lee, Jeong-Sup Song, and Nam Joong Kim. 2015. Middle East respiratory syndrome coronavirus superspreading event involving 81 persons, Korea 2015. *Journal of Korean Medical Science* 30, 11 (2015), 1701–1705.
- [64] Robert Koch-Institut (RKI). 2020. COVID19 data set. Retrieved from <https://npgco-corona-npgeo-de.hub.arcgis.com/datasets/>.
- [65] Filippo Simini, Marta C. González, Amos Maritan, and Albert-László Barabási. 2012. A universal model for mobility and migration patterns. *Nature* 484, 7392 (2012), 96–100.
- [66] Richard A. Stein. 2011. Super-spreaders in infectious diseases. *International Journal of Infectious Diseases* 15, 8 (2011), e510–e513.
- [67] Lauren C. Tindale, Jessica E. Stockdale, Michelle Coombe, Emma S. Garlock, Wing Yin Venus Lau, Manu Saraswat, Louxin Zhang, Dongxuan Chen, Jacco Wallinga, and Caroline Colijn. 2020. Transmission interval estimates suggest pre-symptomatic spread of COVID-19. *Elife* 9 (2020), e57149.
- [68] Carmela Troncoso, Mathias Payer, Jean-Pierre Hubaux, Marcel Salathé, James Larus, Edouard Bugnion, Wouter Lueks, Theresa Stadler, Apostolos Pyrgelis, Daniele Antonioni, et al. 2020. Decentralized privacy-preserving proximity tracing. arXiv preprint arXiv:2005.12273 (2020).
- [69] United Kingdom National Health Service. 2020. NHS COVID-19 app risk-scoring algorithm. Retrieved 30 April 2021 from <https://covid19.nhs.uk/risk-scoring-algorithm.html>.

- [70] Utkarsh Upadhyay, Abir De, and Manuel Gomez-Rodriguez. 2018. Deep reinforcement learning of marked temporal point processes. *Advances in Neural Information Processing Systems* 31 (2018).
- [71] Neeltje van Doremalen, Trenton Bushmaker, Dylan H. Morris, Myndi G. Holbrook, Amandine Gamble, Brandi N. Williamson, Azaibi Tamin, Jennifer L. Harcourt, Natalie J. Thornburg, Susan I. Gerber, James O. Lloyd-Smith, Emmie de Wit, and Vincent J. Munster. 2020. Aerosol and Surface Stability of SARS-CoV-2 as Compared with SARS-CoV-1. *New England Journal of Medicine* 382, 16 (2020), 1564–1567. <https://doi.org/10.1056/NEJMc2004973> arXiv: <https://doi.org/10.1056/NEJMc2004973>
- [72] Piet Van Mieghem. 2011. The N-intertwined SIS epidemic network model. *Computing* 93, 2 (2011), 147–169.
- [73] Piet Van Mieghem, Jasmina Omic, and Robert Kooij. 2008. Virus spread in networks. *IEEE/ACM TON* 17, 1 (2008), 1–14.
- [74] Andreas Walker, Torsten Houwaart, Tobias Wienemann, Malte Kohns Vasconcelos, Daniel Strelow, Tina Senff, Lisanna Hülse, Ortwin Adams, Marcel Andree, Sandra Hauka, et al. 2020. Genetic structure of SARS-CoV-2 reflects clonal superspreading and multiple independent introduction events, North-Rhine Westphalia, Germany, February and March 2020. *Eurosurveillance* 25, 22 (2020), 2000746.
- [75] Dawei Wang, Bo Hu, Chang Hu, Fangfang Zhu, Xing Liu, Jing Zhang, Binbin Wang, Hui Xiang, Zhenshun Cheng, Yong Xiong, et al. 2020. Clinical Characteristics of 138 Hospitalized Patients With 2019 Novel Coronavirus-Infected Pneumonia in Wuhan, China. *JAMA* 323, 11 (2020), 1061–1069.
- [76] Chad R. Wells, Pratha Sah, Seyed M. Moghadas, Abhishek Pandey, Affan Shoukat, Yaning Wang, Zheng Wang, Lauren A. Meyers, Burton H. Singer, and Alison P. Galvani. 2020. Impact of international travel and border control measures on the global spread of the novel 2019 coronavirus outbreak. *PNAS Proceedings of the National Academy of Sciences* 117, 13 (2020), 7504–7509.
- [77] Roman Woelfel, Victor Max Corman, Wolfgang Guggemos, Michael Seilmaier, Sabine Zange, Marcel A. Mueller, Daniela Niemeyer, Patrick Vollmar, Camilla Rothe, Michael Hoelscher, et al. 2020. Clinical presentation and virological assessment of hospitalized cases of Coronavirus disease 2019 in a travel-associated transmission cluster. *medRxiv* (2020).
- [78] World Health Organization. 2020. Report of the WHO-China Joint Mission on Coronavirus Disease 2019 (COVID-19). Retrieved 30 April 2021 from <https://www.who.int/docs/default-source/coronaviruse/who-china-joint-mission-on-covid-19-final-report.pdf>.
- [79] Jian Wu and Peter Frazier. 2016. The parallel knowledge gradient method for batch bayesian optimization. In *Proceedings of the Advances in Neural Information Processing Systems* 29. Curran Associates, Inc., 3126–3134.
- [80] Hongteng Xu, Mehrdad Farajtabar, and Hongyuan Zha. 2016. Learning granger causality for hawkes processes. In *Proceedings of the International Conference on Machine Learning*. PMLR, 1717–1726.
- [81] Xiao-Yong Yan, Chen Zhao, Ying Fan, Zengru Di, and Wen-Xu Wang. 2014. Universal predictability of mobility patterns in cities. *Journal of The Royal Society Interface* 11, 100 (2014), 20140834.
- [82] Ali Zarezade, Sina Jafarzadeh, and Hamid R. Rabiee. 2018. Recurrent spatio-temporal modeling of check-ins in location-based social networks. *Plos One* 13, 5 (2018), e0197683.
- [83] Yunjun Zhang et al. 2020. Evaluating transmission heterogeneity and super-spreading event of COVID-19 in a metropolis of China. *International Journal of Environmental Research and Public Health* 17, 10 (2020), 3705.
- [84] Ke Zhou, Hongyuan Zha, and Le Song. 2013. Learning triggering kernels for multi-dimensional hawkes processes. In *Proceedings of the International Conference on Machine Learning*. PMLR, 1301–1309.
- [85] George Kingsley Zipf. 1946. The P 1 P 2/D hypothesis: On the intercity movement of persons. *American Sociological Review* 11, 6 (1946), 677–686.
- [86] Simiao Zuo, Haoming Jiang, Zichong Li, Tuo Zhao, and Hongyuan Zha. 2020. Transformer hawkes process. In *Proceedings of the International Conference on Machine Learning*. PMLR, 11692–11702.

Received 29 April 2021; revised 28 March 2022; accepted 4 April 2022