Journal of Computational Science 77 (2024) 102242

Contents lists available at ScienceDirect

Journal of Computational Science

journal homepage: www.elsevier.com/locate/jocs

# Multilevel optimization for policy design with agent-based epidemic models

Jan-Hendrik Niemann [a,b,*], Samuel Uram [a], Sarah Wolf [a], Nataša Djurdjevac Conrad [b], Martin Weiser [b]

[a] Department of Mathematics and Computer Science, Freie Universität Berlin, Arnimallee 14, Berlin, 14195, Germany
[b] Zuse Institute Berlin, Takustraße 7, Berlin, 14195, Germany

ARTICLE INFO

Dataset link: https://github.com/Jan-HendrikNiemann/MLoptABM

*MSC:*
65K10
90C15
65C30
92D30
65M75

*Keywords:*
Agent-based models
Multilevel optimization
Epidemiological modeling
Gradient approximation
Optimal control

ABSTRACT

Epidemiological modeling has a long history and is often used to forecast the course of infectious diseases or pandemics. These models come in different complexities, ranging from systems of simple ordinary differential equations (ODEs) to complex agent-based models (ABMs). The former allow a fast and straightforward optimization, but are limited in accuracy, detail, and parameterization, while the latter can resolve spreading processes in detail, but are extremely expensive to optimize. Epidemiological modeling can also be used to propose and design non-pharmaceutical interventions such as lockdowns. In general, their optimal design often leads to nonlinear optimization problems. We consider policy optimization in a prototypical situation modeled as both ODE and ABM, review numerical optimization approaches, and propose a heterogeneous multilevel approach based on combining a fine-resolution ABM and a coarse ODE model. Numerical experiments, in particular with respect to convergence speed, are given for illustrative examples.

## 1. Introduction

The global COVID-19 pandemic has highlighted as never before the need for mathematical modeling to forecast infection spreading and assess consequences of various non-pharmaceutical interventions taken to counter it, such as closing of schools, social distancing rules, or travel restrictions. Such intervention policies have different direct costs as well as social or economic impacts, and differ in their effectiveness with respect to various objectives (e.g., minimizing the expected number of deaths or the number of new infections in a given time window), and in their negative side effects. While policy making must take into account many quantitatively poorly understood aspects, and therefore cannot rely solely on mathematical models of infection spreading, model-based policy optimization can be an important decision support tool.

Models used to study epidemics range from deterministic compartmental models based on ordinary differential equations (ODEs) such as variants of the well-known SIR model (considering population fractions of susceptible, infected, and recovered individuals), see, e.g., [1–4], to spatially resolved stochastic agent-based models (ABMs) [5–9], which describe the dynamics of infection as arising from a large number of discrete interactions between agents in complex interaction networks embedded in a common environment.

Each modeling approach has its own trade-offs. On the one hand, ODE models are computationally inexpensive and easy analyze. A variety of mathematical methods and tools are readily available. On the other hand, as aggregate models, ODE models provide coarse rather than detailed forecasts of infection dynamics that do not allow for uncertainty estimates of the forecasts, and policy measures are difficult to translate into changes in model parameterization. Models based on stochastic differential equations (SDEs) can serve as a means to quantify uncertainty. However, these are often still compartmental models with the same drawbacks.

In contrast, complex ABMs allow for the representation of locally adapted policies and the analysis of their consequences in terms of spatial (e.g., neighborhood, city, region), compartmental (e.g., children, students, workers, retirees), or sectoral (e.g., medical, educational, retail) dimensions, and thus provide much more detailed forecasts, including uncertainty quantification. This makes ABMs attractive for policy design, especially for locally adapted policy design in complex environments. The main drawback is that it is computationally expensive to simulate realistic ABMs. Also, the models are inherently stochastic, requiring the simulation of several to many samples to

* Corresponding author at: Department of Mathematics and Computer Science, Freie Universität Berlin, Arnimallee 14, Berlin, 14195, Germany.
*E-mail address:* niemann@zib.de (J.-H. Niemann).

https://doi.org/10.1016/j.jocs.2024.102242
Received 9 May 2023; Received in revised form 29 November 2023; Accepted 22 February 2024
Available online 23 February 2024

obtain a reliable answer. Perhaps most importantly, mathematical formulation (and hence analysis) is still the exception (for some examples see [10–15]) rather than the rule. While ABMs are widely used in various scientific disciplines such as sociology, economics, or geography, the focus is mostly on computer simulations.

Optimal control is about finding a control for a dynamic system and a given time window in such a way that an objective function is optimized, often under given constraints. Numerical methods for solving complex and nonlinear optimization problems are well established. Optimal control and stability analysis of ODE epidemic models to support policy design has a long tradition [16], but is still an active area of research, e.g., for tuberculosis [17] or HIV [18]. Most recently, optimal control has been used to study COVID-19, e.g., [19–22].

In contrast, mathematical optimization of ABMs faces significant challenges. Basic concepts necessary for optimization, such as derivatives, are not well-defined due to the discrete nature of the steps and decisions made by agents in most ABMs. A recent idea to circumvent this is internal smoothing [23]. Adjoint concepts for discontinuous option pricing [24] may also provide a direction for future research. However, most of the approaches used so far avoid the consideration of derivatives in the ABM altogether, for example, by applying genetic algorithms directly to an ABM [25,26], or by relegating the optimization problem to another model that approximates the ABM. This may be a reduced ABM, derived for example by coarsening the spatial resolution of the underlying one [27], so that again heuristic algorithms are applied for optimization. In most cases, an "equation-based" or "system-level" model is used, for which optimization methods exist. The choice of such a coarser model is based on what is already available for similar phenomena. Alternatively, a simple functional form is chosen and parameters are fitted to the ABM output (see [28,29] and references therein). The optimal controls for the reduced model are then lifted back to the ABM.

In this paper, we apply a multilevel approach that approximates the original ABM at the fine level with a ODE at the coarse level. Multilevel methods exploit the computationally cheaper optimization of coarser models to reduce the optimization effort in a more detailed "fine" model. However, the perspective taken is not that of finding optimal interventions for infectious diseases such as COVID-19, but that of numerical optimization, since finding optimal interventions is the subject of a political and societal debate that we do not intend to enter in this article. To minimize a given (prototypical and exemplary) objective, we consider derivatives and descent along them. In the (stochastic) ABM we work with the numerical approximation of derivatives by finite differences. In the multilevel approach presented here, the combination of coarse and fine levels is iterative: a coarse-level ODE model is used to generate trial steps for descent directions in the fine-level ABM. This step must then pass an acceptance test to ensure descent of the objective with high confidence. The main contributions of this paper are:

- We present gradient-based optimization approaches of complex ABMs to solve policy optimization problems. To this end, we first review numerical optimization approaches applicable to ABMs.
- Then, we present a novel hybrid multilevel optimization method that combines a fine-level ABM with a coarse-level ODE to solve the policy optimization problem more efficiently.
- We compare and illustrate the presented methods along epidemiological ABMs for the transmission dynamics of SARS-CoV-2.

The remainder is organized as follows. In Section 2, we introduce three epidemiological models and investigate their quantitative relation. Policy optimization is considered in Section 3. We briefly review state-of-the-art optimization algorithms suitable for ODEs and ABMs and discuss their computational complexity. We then introduce a heterogeneous multilevel optimization approach that combines an ABM with an ODE approximation to speed up convergence. Their relative computational efficiency is investigated in Section 4 using numerical examples of the transmission dynamics of SARS-CoV-2. Open questions and future work are discussed in Section 5.

## 2. Epidemiological models

In this section, we introduce the terminology and notation essential for the epidemiological models used in this study. Specifically, we discuss models based on ODEs, which are introduced in Section 2.1, and stochastic ABMs presented in Section 2.2. Furthermore, we analyze the relation and agreement between both modeling approaches in Section 2.3.

### 2.1. Deterministic ODE models

Mathematical models based on ODEs are commonly used to simulate epidemic spread on a macro-scale of a fully mixed population, e.g., a city or a country. These models separate the population into different compartments and count only the proportions of individuals belonging to each compartment. Compartments can, for example, be based on the infection status or demographics or both. In this paper, we consider so-called SIR models [6] and divide the population into two age groups: adults and children (compartments denoted by subscript a or c, respectively). Separation into age groups has proven beneficial due to different infection rates, see, e.g., [30] for the case of SARS-CoV-2. Individuals who have not yet been infected belong to the susceptible compartment $S_\circ$, for $\circ \in \{\mathrm{a},\mathrm{c}\}$. With a certain infection rate $r_{*\to\circ}$, for $*,\circ \in \{\mathrm{a},\mathrm{c}\}$, individuals from $S_\circ$ can get infected by an individual from $I_*$ and move to the compartment $I_\circ$. An infection is caused by interactions of the type $S_\circ + I_* \to I_\circ + I_*$, which are referred to as *second-order* interactions since two individuals are involved. Individuals in the group $I_\circ$ recover with rate $r_\circ$ and move to the compartment $R_\circ$ due to the *first-order* reaction $I_\circ \to R_\circ$. The system of ODEs describing this model is given by

$$\begin{aligned}
\dot{S}_\mathrm{a} &= -S_\mathrm{a}\,(r_{\mathrm{a}\to\mathrm{a}} I_\mathrm{a} + r_{\mathrm{c}\to\mathrm{a}} I_\mathrm{c}) \\
\dot{S}_\mathrm{c} &= -S_\mathrm{c}\,(r_{\mathrm{c}\to\mathrm{c}} I_\mathrm{c} + r_{\mathrm{a}\to\mathrm{c}} I_\mathrm{a}) \\
\dot{I}_\mathrm{a} &= S_\mathrm{a}\,(r_{\mathrm{a}\to\mathrm{a}} I_\mathrm{a} + r_{\mathrm{c}\to\mathrm{a}} I_\mathrm{c}) - r_\mathrm{a} I_\mathrm{a} \\
\dot{I}_\mathrm{c} &= S_\mathrm{c}\,(r_{\mathrm{c}\to\mathrm{c}} I_\mathrm{c} + r_{\mathrm{a}\to\mathrm{c}} I_\mathrm{a}) - r_\mathrm{c} I_\mathrm{c} \\
\dot{R}_\mathrm{a} &= r_\mathrm{a} I_\mathrm{a} \\
\dot{R}_\mathrm{c} &= r_\mathrm{c} I_\mathrm{c}
\end{aligned} \tag{1}$$

Obviously, the sum of the groups $S_\circ + I_\circ + R_\circ$ is conserved over time, such that $R_\circ$ is often omitted because of redundancy. If it is needed, it can be computed from $S_\circ$ and $I_\circ$ as $R_\circ = N_\circ - S_\circ - I_\circ$, where $N_\circ$ denotes the fraction of individuals of the respective age group. In the following, when we do not distinguish between age groups, we will refer to the fraction of susceptible $S = S_\mathrm{a} + S_\mathrm{c}$, infected $I = I_\mathrm{a} + I_\mathrm{c}$, and recovered individuals $R = R_\mathrm{a} + R_\mathrm{c}$, such that $S + I + R = 1$. A numerical solution of ODE (1) with parameter choices given in Table A.2 is shown in Fig. 1(a). For an introduction to mathematical methods for this kind of models we refer the reader to [31].

### 2.2. Stochastic agent-based models

As an alternative to ODE-based epidemic models, ABMs can be used to model infection dynamics, which we will introduce in this section. In addition to infection status, individuals are often explicitly represented and characterized by, for example, their activities in specific locations, as well as other relevant characteristics such as age or sex. In simulations, infection events occur with given probabilities whenever an infected and a susceptible agent meet. Due to the probabilistic nature of infection events (and possibly the inherent randomness of the ABM itself), the numbers $S(t)$, $I(t)$, and $R(t)$ are random variables. Thus, many simulations are typically performed to compute statistical moments. We will now present our two guiding examples.

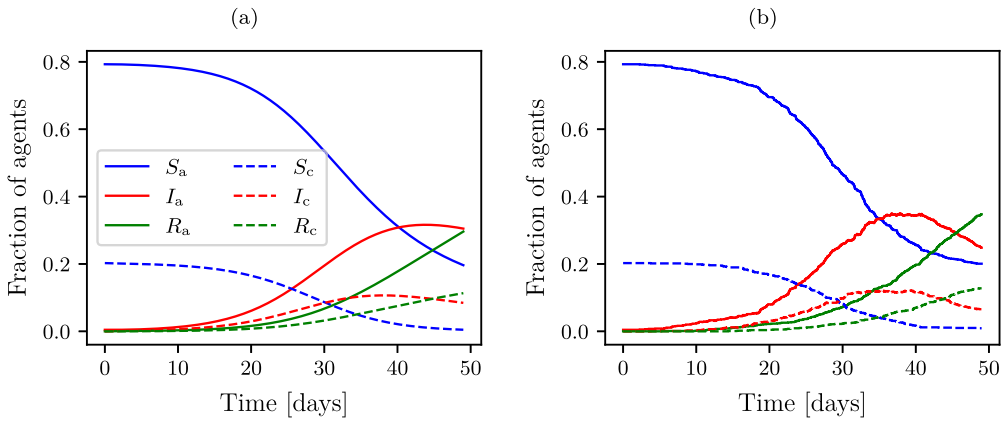


**Fig. 1.** (a) Numerical solution of the ODE (1) and (b) aggregated trajectory of a single simulation of GERDA given identical initial conditions divided into adults (solid lines) and children (dashed lines) for parameters given in Table A.2.

#### 2.2.1. Example 1 – the georeferenced demographic agent-based model

As first example, we consider the individualized **GE**o**R**eferenced **D**emographic **A**gent-based model (GERDA) for the transmission of SARS-CoV-2 and the disease dynamics of COVID-19 [32]. This model uses detailed location data and a demographically matched population of agents with realistic daily schedules to simulate contacts between people in the given locations and the resulting infection events. The disease progression of infected agents is modeled in detail with rates and times, e.g., for transition to diagnosed, hospitalized, ICU, and recovered or deceased states, depending on the agents' age group. The model has been used with data for several towns (e.g., Tepoztlán (Mexico), Zikhron Ya'akov (Israel), Gangelt (Germany)). In this work we use the data calibrated for the German municipality of Gangelt. For more details, see [32].

To simplify the disease dynamics computed by GERDA, we summarize the different subclasses of infected agents and assume that all infected agents recover. Additionally, agents are grouped by age, i.e., children (age 0 to 18 years) and adults (19+). Fig. 1(b) shows such an aggregated trajectory in terms of numbers of susceptible, infected and recovered adults and children for a single simulation of GERDA.

#### 2.2.2. Example 2 – a homogeneous ABM

In addition to GERDA, we present another but coarser ABM for modeling the spread of an infectious disease. We assume two homogeneous groups of agents, i.e., adults and children, and that any agent can be infected by any other agent at any time. Thus, we call it a *homogeneous* ABM and use the shorthand notation H/ABM to refer to it. Homogeneity allows tracking only the number of agents per compartment, making the evaluation several orders of magnitude faster than for GERDA. We represent the H/ABM as a Markov jump process that builds on the identical first- and second-order reactions as the ODE model (1). We wish to make explicit the character of an ABM and therefore write the infection rates $r_{*\to\circ}$, for $*,\circ \in \{\mathrm{a},\mathrm{c}\}$ for a population of size $N$ instead of using fractions. Additionally, since in GERDA agents can lose their immunity, we introduce a third transition $R_\circ \to S_\circ$, i.e., recovered individuals from group $R_\circ$ move to the compartment $S_\circ$ and become susceptible again. We assume that this happens independently of the age group with rate $\mu r_\circ$, for $\mu \in ]0,1[$. Note that in this work the ODE model and the H/ABM share the same parameters. The H/ABM can be simulated using Gillespie's stochastic simulation algorithm [33], which constructs exact realizations of the H/ABM in continuous time. For further details, especially with respect to convergence to ordinary or stochastic differential equations, we refer the reader to [34].

### 2.3. Model agreement

To be used in the multilevel optimization algorithm in Section 3.5, the ODE model and the ABM must describe the same phenomenon, but at different levels of detail. We will now discuss how to achieve agreement between the two models.

While many parameters in an ABM can, at least in principle, be measured or observed, e.g., frequencies of people meeting, the six ODE parameters $r_{*\to\circ}$ and $r_*$ are not directly related to measurable physical quantities. A reasonable parameterization is necessarily based on the agreement of ODE predictions with reality, where "reality" can refer to real-world data or to predictions of a different model taken as a reference (in this work GERDA). We are interested in the latter case and therefore fit the parameters $r_{*\to\circ}$ and $r_*$ in the sense of $L^2$ similarity between the ODE and GERDA results for a fixed number of agents $N_0$. A least-squares approach leads to the optimization problem

$$\min_{r_{*\to\circ}, r_*} \int_0^T \left( V_{I,\mathrm{a}}^{-1} (I_\mathrm{a}^\mathrm{ODE} - I_\mathrm{a}^\mathrm{ABM})^2 + V_{I,\mathrm{c}}^{-1} (I_\mathrm{c}^\mathrm{ODE} - I_\mathrm{c}^\mathrm{ABM})^2 \right) \mathrm{d}t, \tag{2}$$

where $I_*^\mathrm{ABM}$ is computed as a mean of a set of 1 000 independent GERDA simulations given the parameters in Table A.2. The sample variances $V_{I,*}$ of the respective infected age groups serve as weighting factors for the fit. Fig. A.8 shows the result, obtained by using a nonlinear least-squares solver (an in-built MATLAB solver) for problem (2). For comparability of GERDA and the H/ABM, we parameterize the H/ABM identically.

## 3. Policy optimization

In the following section, we will discuss the elements needed in addition to the epidemiological model to design favorable policies. Following our guiding examples, in Section 3.1 we discuss a potential design space of policies that affect the model dynamics, e.g., by changing parameters. Then, in Section 3.2, we introduce a prototypical objective to be minimized. In Sections 3.3 and 3.4 we present optimization algorithms that can be applied to the problems. Finally, in Section 3.5 we present a hybrid multilevel optimization method applied to epidemiological ABMs. We close Section 3 by a qualitative comparison with other methods.

### 3.1. Policies

It is well known that the spread of infectious diseases can be influenced not only by pharmaceutical interventions, such as medical treatment or vaccination, but also by non-pharmaceutical interventions, such as social distancing or lockdowns to reduce physical contact between individuals. For our discussion, we consider two time-dependent

controls $u_s, u_w : [0,T] \to [0,1]$ as examples representing policies for closing schools and workplaces, or homeworking for those jobs where work can be done at home, respectively. More precisely, the value of $u_s(t)$ refers to the fraction of children not going to school at time $t$, i.e., $u_s(t) = 0$ stands for the policy that no schools are closed and $u_s(t) = 1$ indicates that all schools are closed. Analogously, the control $u_w(t)$ is the fraction of workplaces closed. We denote $u = [u_s, u_w]^\top$ to refer to the combination of controls. The controls are assumed to be piece-wise constant on a time grid $0 = t_0 < \cdots < t_m = T$ to reflect the fact that policy changes take some time to be implemented and that continuously changing policies would be complicated to communicate and follow. Consequently, the admissible set of policies

$$U = \{u : [0,T] \to [0,1]^2 \mid \forall i \in \{1,\dots,m\} : u|_{]t_{i-1},t_i[} = \text{const}\},$$

of dimension $n_u = 2m$ is isomorphic to $[0,1]^{2m} \subset \mathbb{R}^{n_u}$ such that any control can be represented by $n_u$ scalar control values. Policies affect the simulated spread of infection by changing the ODE (resp. H/ABM) parameters or, in case of GERDA, by changing the agents' schedules (e.g., staying home instead of going to school). Thus, identical policies must be implemented in different ways depending on the model used. Note that the polices discussed here are for illustrative purposes only. Further policies could be taken, for example, a ban on visiting public places, wearing a mask, contact tracing and isolation, or vaccination.

**Remark 3.1.** School closure and work from home are well-established interventions, both in the original publication [32] and in practice. They have proven to be important and easily enforceable. Further and more distinct closures such as elementary schools, high schools, and universities, or businesses of everyday use, shopping malls, and leisure places are possible, but increase the parameter space of the optimization problem, resulting in higher computational costs.

#### 3.1.1. Implementing policies in the ODE

In the ODE model, infection rates are determined by the properties of the virus and the population. Additionally, we assume that they depend on the controls $u_s$ and $u_w$. For the infection rates within one age group $r_{*\to*}$ we assume quadratic dependence on the fraction of individuals since an infection event can only occur when two individuals meet, i.e.,

$$r_{a\to a} = r_{aa}(1-u_w)^2, \quad r_{c\to c} = r_{cc}(1-u_s)^2.$$

Here, the parameters $r_{aa}$ and $r_{cc}$ are independent of the policies. Quadratic dependency is widely used for general second-order interactions [35], but is obviously highly simplistic for epidemiological modeling as, for example, it neglects possible infections outside schools and workplaces. Similarly, we assume the infection rates between age groups $r_{*\to\circ}$ to depend on the controls as well, though with impact reduced by a factor of one half since many such infections will only happen within families, i.e.,

$$r_{c\to a} = r_{a\to c} = r_{ac}(1-u_w/2)(1-u_s/2).$$

The recovery rates $r_*$ are assumed to be independent of the controls.

**Remark 3.2.** Due to its strong similarity in terms of modeling, the policies for the H/ABM are implemented analogously to the ODE model. For example, using $\tilde{r}_{*\circ} := r_{*\circ}N_0$, the propensity function that characterizes the infection of an adult by an adult including the policy is given by $\tilde{r}_{a\to a} = \tilde{r}_{aa}(1-u_w)^2/N$, where $N$ is the population size used for the H/ABM. The propensity functions $\tilde{r}_{c\to c}$ and $\tilde{r}_{c\to a} = \tilde{r}_{a\to c}$ follow equivalently.

#### 3.1.2. Implementing policies in GERDA

In contrast to the ODE model, in GERDA, where contacts between agents are explicitly modeled, policies do not affect the probability of infection, but the frequency of those contacts. Temporary closure of the workplace or order to work from home is done by randomly selecting adults in a ratio equal to the control value $u_w(t)$ and changing their schedules so that they stay home. Analogously, at the rate $u_s(t)$, children stay at home instead of going to school. We assume that an adult from the same household must also be home to supervise children under the age of 13. With a high share of school closures but open workplaces, this leads to an implicit homeworking obligation for the supervisors involved. However, this effect is not reflected by the controls, since they represent only government-mandated interventions. Both policies result in a significant reduction in interactions between agents, which in turn leads to fewer infections and a slower spread of the disease.

### 3.2. Objective

In this section, we define a simple example of an objective that nevertheless raises interesting optimization questions. We assume that the general aim of pandemic policy management is to reduce the number $I$ of infected people while minimizing the economic and social costs of the interventions.

The first aim is reflected in the health objective

$$c_h(u) = \int_{t=0}^{T} (I(t)/N + \exp(10(I(t) - I_{\max})/N))\,\mathrm{d}t,$$

where the number of infected agents $I(t)$ depends implicitly on the control $u$ via the ABM or ODE model, respectively. The first term represents the negative impact of infections, which grows approximately linearly with the number of infected people, assuming that the number of people requiring medical treatment is a fixed fraction of the infected. The second term represents the social impact, which increases dramatically once the capacity of the health system for proper treatment of the severely ill infected is exhausted. This is formulated in terms of the health care system's carrying capacity $I_{\max}$.

Both controls incur direct costs, interpreted as economic impact $c_w$ for homeworking and social impact $c_s$ for school closure. The economic impact becomes $+\infty$ when the homeworking rate approaches an upper bound $u_w^{\max} < 1$, reflecting the fact that not all workplaces can be closed or done from home. Thus, we define the second aim $c_w$ by

$$c_w(u) = \int_0^T -\log(u_w^{\max} - u_w(t))\,\mathrm{d}t.$$

We further assume that the social impact of school closures depends quadratically on the fraction of schools closed, i.e.,

$$c_s(u) = \int_0^T u_s(t)^2\,\mathrm{d}t.$$

Combining these three aims with weights $a_s, a_w > 0$, we define our objective

$$J(u) := c_h(u) + a_s c_s(u) + a_w c_w(u). \tag{3}$$

Note that a common factor for all three terms does not affect the minimizer, so that we can normalize the weight for $c_h$ to one.

**Remark 3.3.** The functions $c_h$, $c_w$ and $c_s$ are chosen based on qualitative arguments, such that they have a singularity at total lockdown. The constant $I_{\max}$ is selected as a threshold. Alternatively, one could model $I_{\max}$ as a state constraint. However, state constraints model abrupt changes in the objective function. Instead, our choice reflects the increasing difficulty of allocating patients to beds as bed availability diminishes. It represents the fact that in one region individuals may die due to inadequate care, even if there are unoccupied beds in another region. Although the transition is notably sharp, it does not result in an abrupt discontinuity or jump to infinity. However, any monotone

function with a small slope around 0 and a steep slope away from 0 (i.e., close to total lockdown) could be chosen. Similarly to the way $c_h$ is modeled, also $c_s$ and $c_w$ are based on qualitative reasoning. Various modeling approaches for the economic, social and health impacts have been used in the literature, e.g., [4,36–38].

If the number of infected is computed using any time stepping method for solving the SIR initial value problem (1), see, e.g., [31], the objective (3) can be evaluated directly using numerical integration. We refer to this as $J^{\mathrm{ODE}}$. For an ABM involving randomness, the number of infected is a stochastic process. Each simulation yields a different trajectory for the number of infected. In this case, we refer to (3) as $J^{\mathrm{ABM}}$, and define the actual objective as the expectation $\mathbb{E}[J^{\mathrm{ABM}}(u)]$, which can be evaluated using a simple Monte Carlo integration

$$\mathbb{E}[J^{\mathrm{ABM}}(u)] \approx \frac{1}{n}\sum_{i=1}^{n} J^{\mathrm{ABM}}(u)_i =: \mathbb{E}_n[J^{\mathrm{ABM}}(u)] \tag{4}$$

based on samples $J^{\mathrm{ABM}}(u)_i$. The accuracy of this unbiased estimate is given in terms of its standard deviation as

$$\sigma(\mathbb{E}_n[J^{\mathrm{ABM}}(u)]) \approx \sigma_n(J^{\mathrm{ABM}}(u))/\sqrt{n} \tag{5}$$

with the sample variance

$$\sigma_n^2(J^{\mathrm{ABM}}(u)) \approx \frac{1}{n-1}\sum_{i=1}^{n}(J^{\mathrm{ABM}}(u)_i - \mathbb{E}_n[J^{\mathrm{ABM}}(u)])^2.$$

The well-known slow convergence of Monte Carlo methods of order $\mathcal{O}(n^{-1/2})$ requires a large number of expensive ABM simulations for a faithful evaluation of the objective. Variance reduction methods such as control variates, quasi-Monte Carlo sequences, or hierarchical Monte Carlo methods are popular approaches to reduce the numerator in (5) and thus the sampling error [39]. Unfortunately, they are intrusive by requiring a joint probability space underlying the different random variables used, and thus a carefully designed use of random number generators for sampling. This is difficult to achieve for complex ABMs designed from scratch, and virtually impossible for existing ABMs. Thus, in general, one has to rely on the simple Monte carlo evaluation (4). We will now discuss how the two models, i.e., the ODE model and the ABM model, can be optimized, as this is key to the optimization algorithm in Section 3.5.

### 3.3. ODE optimization

For an ODE model, the policy optimization problem

$$\min_{u\in U} J^{\mathrm{ODE}}(u)$$

is a classical optimal control problem, for which several well-studied numerical solution approaches exist, see, e.g., [40,41]. We briefly outline the exact gradient descent algorithm, since subsequent algorithms build on it.

The idea is to go downhill for some step size $\alpha$ in direction of the steepest descent $s_k = -\nabla J^{\mathrm{ODE}}(u)$. Outward pointing components are projected to the set of admissible controls $U$ (here the unit cube in $\mathbb{R}^{n_u}$), i.e.,

$$P_U(x,y)_i = \begin{cases} 0, & (y_i = 0 \wedge x_i < 0) \vee (y_i = 1 \wedge x_i > 0) \\ x_i, & \text{otherwise} \end{cases}$$

for $i \in \{1,\dots,n_u\}$. The step size $\alpha$ is conceptually the largest step size that (1) satisfies a constant bound $\alpha \le \alpha_0$ and admissibility $\alpha \le \bar{\alpha}(u_k) := \max\{a \mid u_k + a s_k \in U\}$, and (2) satisfies the Armijo rule of sufficient decrease,

$$J^{\mathrm{ODE}}(u_k + \alpha s_k) \le J^{\mathrm{ODE}}(u_k) + \alpha c_1 \nabla J^{\mathrm{ODE}}(u_k) s_k. \tag{6}$$

The latter ensures that at least a certain fraction $c_1 \in ]0,1[$ of the descent promised by the first-order Taylor approximation of the objective is realized. The most efficient way to compute the gradient $\nabla J^{\mathrm{ODE}}(u)$ is to solve the adjoint equations associated with the ODE system (1) and $J^{\mathrm{ODE}}$; for details see Appendix B.

**Algorithm 1:** Basic gradient descent algorithm

```
Input u_0 ∈ U, α_0 > 0, c_1 ∈ ]0,1[
for k = 0, ..., K do
    compute s_k = P_U(−∇J^ODE(u_k), u_k)
    α = min{α_0, ᾱ(u_k)}
    while Armijo condition (6) violated do
        α = α/2
    u_{k+1} = u_k + α s_k
Output u* = u_K
```

Algorithm 1 converges, usually at a linear rate, to a stationary point, usually a local minimizer, under rather mild regularity assumptions of $J^{\mathrm{ODE}}$. Its asymptotic convergence rate deteriorates with growing condition number of the Hessian of $J^{\mathrm{ODE}}$, i.e., the ratio of largest and smallest eigenvalue. Preconditioning, i.e., replacing $\nabla J^{\mathrm{ODE}}$ by $B^{-1}\nabla J^{\mathrm{ODE}}$ with some approximation $B \approx (J^{\mathrm{ODE}})''$ of the Hessian, can improve the convergence speed significantly, but constructing effective and computationally cheap preconditioners $B^{-1}$ is not trivial. For details and more sophisticated algorithms we refer to the nonlinear optimization literature [42].

### 3.4. ABM optimization

Solving the policy optimization problem for realistic ABMs, i.e.,

$$\min_{u\in U} \mathbb{E}[J^{\mathrm{ABM}}(u)], \tag{7}$$

is much more difficult than solving it for the ODE model due to three major challenges:

(i) ABMs are much more complex, involving thousands of agents, and they often simulate the agents' activities with a higher time resolution, whereas ODE models often only represent the average population and thus omit the vast majority of details that ABMs provide. Thus, computing a single ABM trajectory is orders of magnitude more expensive.

(ii) ABMs are inherently stochastic, such that many independent simulations are required to approximate the objective with sufficient accuracy.

(iii) ABMs are inherently discontinuous due to the discrete decisions of the agents. Thus, single ABM trajectories are not differentiable with respect to the control $u$, so that efficient adjoint gradient computation, as possible for ODE models, is not directly available.

Therefore, for the time being, ABM optimization must rely on objective samples alone, inferring descent directions from noisy objective evaluations.

#### 3.4.1. Gradient approximation

For approximating directional derivatives of the objective $\mathbb{E}[J^{\mathrm{ABM}}(u)]$ in direction $v \in \mathbb{R}^{n_u}$, we use finite differences

$$\begin{aligned}\nabla\mathbb{E}[J^{\mathrm{ABM}}(u)]^\top v &\approx (2h)^{-1}\left(\mathbb{E}_n[J^{\mathrm{ABM}}(u+hv)] - \mathbb{E}_n[J^{\mathrm{ABM}}(u-hv)]\right)\\ &= \frac{1}{2hn}\sum_{i=1}^{n}\left(J^{\mathrm{ABM}}(u+hv)_i - J^{\mathrm{ABM}}(u-hv)_i\right)\\ &=: \nabla_h\mathbb{E}_n[J^{\mathrm{ABM}}(u)]^\top v\end{aligned} \tag{8}$$

for small $h > 0$. Using the unit vectors $v = e_k$ for $k = 1,\dots,n_u$, the complete gradient vector can be obtained. The approximation error is of order

$$\|\nabla\mathbb{E}[J^{\mathrm{ABM}}(u)] - \nabla_h\mathbb{E}_n[J^{\mathrm{ABM}}(u)]\| = \mathcal{O}\left(h^2 + \frac{\sigma}{h\sqrt{n}}\right),$$

which requires a careful choice of sample size $n$ of order $n = \mathcal{O}(h^{-6})$ to balance discretization and sampling error. Using

$$\nabla_{h,k} J^{\mathrm{ABM}}(u)_i := (2h)^{-1}(J^{\mathrm{ABM}}(u + he_k)_i - J^{\mathrm{ABM}}(u - he_k)_i),$$

where $e_k$ is the $k$th unit vector, in the estimator

$$(V_n)_{kl} = \frac{1}{n-1}\sum_{i=1}^{n} \left(\nabla_{h,k} J^{\mathrm{ABM}}(u)_i - \nabla_h \mathbb{E}[J^{\mathrm{ABM}}(u)]_k\right) \cdot \left(\nabla_{h,l} J^{\mathrm{ABM}}(u)_i - \nabla_h \mathbb{E}[J^{\mathrm{ABM}}(u)]_l\right)$$

for the sample covariance $V_n \in \mathbb{R}^{n_\mathrm{u} \times n_\mathrm{u}}$, the sample mean standard deviation

$$\sigma(\nabla_h J^{\mathrm{ABM}}(u)) = \sqrt{\|V_n\|/n},$$

provides an a posteriori error estimate for the gradient evaluation. Here, $\|\cdot\|$ refers to the 2-norm. If $n$ is chosen such that $2\sigma \leq \epsilon \|\nabla_h \mathbb{E}_n[J^{\mathrm{ABM}}(u)]\|$, a relative accuracy $\epsilon > 0$ of the approximate gradient can be ensured with high probability, i.e.,

$$\|\nabla \mathbb{E}[J^{\mathrm{ABM}}(u)] - \nabla_h \mathbb{E}_n[J^{\mathrm{ABM}}(u)]\| < \epsilon \|\nabla_h \mathbb{E}_n[J^{\mathrm{ABM}}(u)]\|. \tag{9}$$

If applicable, variance reduction methods can be used to reduce the sample mean standard deviation $\sigma$ and thus the sampling error. For example, correlated sampling uses the same random number generator seeds for evaluating $J^{\mathrm{ABM}}(u - hv)_i$ and $J^{\mathrm{ABM}}(u + hv)_i$ in (8), which leads to the two values being increasingly correlated for $h \to 0$ and therefore to an expected error order of $\mathcal{O}\left(h^2 + 1/\sqrt{n}\right)$. If this increased correlation actually is realized by the ABM, the required number $n$ of samples is reduced significantly to $\mathcal{O}(h^{-4})$. Moreover, if the policy change $v$ affects only policies after some time $t_v > 0$, the trajectories corresponding to $u - hv$ and $u + hv$ coincide on $[0, t_v]$ and can be simulated just once using a checkpoint-restart ability of the ABM, if present.

#### 3.4.2. Inexact gradient descent

Unlike for the ODE optimization, the exact steepest descent direction is not available when solving the policy optimization problem (7) for the ABM objective $\mathbb{E}[J^{\mathrm{ABM}}(u)]$. Nevertheless, descent methods still converge as before if sufficient local decrease can be achieved, i.e., if the descent condition $g^\top s \leq -c_0 \|g\| \|s\|$ holds for some $c_0 > 0$, where we write $g := \nabla \mathbb{E}[J^{\mathrm{ABM}}(u)]$ for the exact gradient and $s := -\nabla_h \mathbb{E}_n[J^{\mathrm{ABM}}(u)] \approx -g$ for the search direction. We will now discuss an inexact gradient descent method to solve the policy optimization problem (7).

A sufficiently small relative error $\epsilon < 1/2$ of the computed steepest descent direction, i.e., $\|s + g\| \leq \epsilon \|s\|$, guarantees that $s$ is a descent direction, since by $\|g\| = \|g + s - s\| \geq (1 - \epsilon)\|s\|$ it holds that

$$g^\top s = g^\top(-g + s + g) \leq -\|g\|(1 - \epsilon)\|s\| + \epsilon \|g\| \|s\| \leq -(1 - 2\epsilon)\|g\| \|s\|.$$

In particular, choosing $\epsilon = 1/4$ in (9) yields $c_0 = 1/2$ and thus guarantees at least half of the local progress compared to exact gradient descent. Since $-s^\top g \leq (1 + \epsilon)\|s\|^2$ (using Cauchy–Schwarz) yields

$$-(1 + \epsilon) c_1 \alpha \|s\|^2 \leq c_1 \alpha s^\top g,$$

the Armijo condition (6) is implied by

$$\mathbb{E}[J^{\mathrm{ABM}}(u + \alpha s)] - \mathbb{E}[J^{\mathrm{ABM}}(u)] \leq -(1 + \epsilon) c_1 \alpha \|s\|^2.$$

When evaluating the objective with absolute sampling errors $e \leq 2\sigma$ as in (5), we obtain

$$\mathbb{E}[J^{\mathrm{ABM}}(u + \alpha s)] - \mathbb{E}[J^{\mathrm{ABM}}(u)] \leq \mathbb{E}_n[J^{\mathrm{ABM}}(u + \alpha s)] - \mathbb{E}_n[J^{\mathrm{ABM}}(u)] + 2e.$$

Choosing $e \leq \epsilon c_1 \alpha \|s\|^2$, the Armijo condition is satisfied with high confidence if the acceptance test

$$\mathbb{E}_n[J^{\mathrm{ABM}}(u + \alpha s)] - \mathbb{E}_n[J^{\mathrm{ABM}}(u)] \leq -(1 + 3\epsilon) c_1 \alpha \|s\|^2 \tag{10}$$

is passed for the trial step $s$.

Algorithm 2 converges similarly to the exact gradient descent to a stationary point, usually a local minimizer. However, it is computationally expensive due to the large number of samples required to compute approximate gradients (line 3) and for the line search (line 5).

**Algorithm 2:** Basic inexact gradient descent algorithm

1 **Input** $u_0 \in U$, $\alpha_0 > 0$, $c_1 \in ]0, 1[$
2 **for** $k = 0, \ldots, K$ **do**
3  compute $s_k = P_U(-\nabla_h \mathbb{E}_n[J^{\mathrm{ABM}}(u_k)], u_k)$ with accuracy $\epsilon < 1/2$
4  $\alpha = \min\{\alpha_0, \bar{\alpha}(u_k)\}$
5  **while** *acceptance test* (10) *fails* **do**
6  $\alpha = \alpha/2$
7  $u_{k+1} = u_k + \alpha s_k$
8 **Output** $u^\star = u_k$

### 3.5. Multilevel optimization

In this section, we present a hybrid multilevel optimization approach to solve the policy optimization problem (7). Nonlinear multilevel optimization approaches, such as MG/OPT [43] or recursive multilevel trust region (RMTR) [44], exploit the computationally cheap optimization of coarser models to generate trial steps that can make much more progress than possible with Taylor approximations of the fine model by respecting the underlying nonlinear structures. In the case of an ABM as fine model, a well-matched ODE model can serve as a coarse model, with the advantage of being computationally extremely cheap. Algorithm 3 implements a basic multilevel optimization algorithm. Here, it is restricted to only two levels $J^{\mathrm{ABM}}$ and $J^{\mathrm{ODE}}$.

The two main differences to the inexact gradient descent algorithm are the use of a coarse model for computing the trial step instead of taking the negative gradient direction, and by restricting the trial step to an $l^\infty$ trust region within the $n_\mathrm{u}$-dimensional box $U$ of admissible policies instead of performing a line search. For minimizing the coarse ODE model in line 6, Algorithm 1 is a viable option. The computational cost of solving an ODE optimization problem is negligible compared to evaluating ABM gradients.

In order to guarantee convergence, the step suggested by optimizing the coarse model must lead to a sufficient decrease of the fine objective as required by an Armijo type condition for some trust region radius $\rho > 0$, i.e,

$$\mathbb{E}_n[J^{\mathrm{ABM}}(u + \delta u_k)] - \mathbb{E}_n[J^{\mathrm{ABM}}(u)] \leq -(1 + 3\epsilon) c_1 \delta u_k^\top s \tag{11}$$

with $e \leq \epsilon c_1 \delta u_k^\top s$, needs to be fulfilled. This can be ensured if the coarse model is first-order consistent, i.e., its gradient coincides with the fine model gradient at the current iterate $u_k$, which is achieved by adding the linear correction $-(s_k + \nabla J^{\mathrm{ODE}}(u_k))^\top \delta u$ in line 6. Even then, the computed step does not have to lead to a decrease of the fine objective $J^{\mathrm{ABM}}$ since the coarse and fine model are different. A backtracking line search is not guaranteed to solve this issue since in highly nonlinear problems the step $\delta u_k$ does not necessarily have to be a descent direction. Instead, the minimization of the coarse model is restricted to a neighborhood of the current iterate $u_k$. Using the $l^\infty$ neighborhood $u_k + ]-\rho, \rho[^{n_\mathrm{u}}$ and intersecting it with the admissible set $U$ leads to the same type of box-constrained optimization subproblem as before. For sufficiently small $\rho > 0$, a reduction of $J^{\mathrm{ABM}}$ is ensured by first-order consistency. Note that the selection of the trust region radius $\rho$ is deliberately simple. For more efficient update strategies, see [45].

Algorithm 3 can be extended from two-levels to a true recursive multilevel scheme by solving the coarse model subproblem in line 6 not by a simple gradient method, but by minimizing some even coarser model. However, due to the extreme difference in computational effort between ABM and ODE model, the use of coarser ODE models does not promise any benefit. In contrast, the use of a hierarchy of fine ABM, coarse ABM, SDE models, and ODE models is an interesting perspective, which, however, we will not explore in this work.

**Algorithm 3:** Basic multilevel optimization

**Input** $u_0 \in U$, $\rho_0 > 0$
**for** $k = 0, \dots, K$ **do**
  compute $s_k = P_U(-\nabla_h \mathbb{E}_n[J^{\mathrm{ABM}}(u_k)], u_k)$ with accuracy $\epsilon < 1/2$
  $\rho = \rho_0$
  **repeat**
    $\delta u_k = \underset{\delta u \in (U-u_k)\cap ]-\rho,\rho[^{n_\mathrm{u}}}{\arg\min} J^{\mathrm{ODE}}(u_k+\delta u) - (s_k + \nabla J^{\mathrm{ODE}}(u_k))^\top \delta u$
    $\rho = \rho/2$
  **until** *acceptance test* (11);
  $u_{k+1} = u_k + \delta u_k$
**Output** $u^\star = u_k$

### 3.6. Advantages and limitations of multilevel optimization

Next, we discuss the advantages and limitations of multilevel optimization over other well-established gradient-free and gradient-based methods that are applicable to optimal policy design with ABMs. The main dimensions of comparison are monotonicity, gradient evaluation, and convergence speed. Both ensuring monotone decrease of the objective and making use of gradient information usually improve reliability and iteration count, but increase the computational cost.

Gradient-free optimization methods [46] like particle swarm or evolutionary algorithms are applicable but require too many ABM evaluations as they have no inherent concept of descent directions. In some approaches, approximate descent directions are obtained from averaging over populations, which incurs a large number of model evaluations. Nelder–Mead and coordinate search perform a systematic search for descent, but have a very coarse direction resolution, leading to quite inexact and suboptimal directions of the steps taken.

Gradient-based methods usually take far fewer steps, but require the evaluation of gradients, which can be costly [42]. Steepest descent takes locally optimal descent directions, but suffers from "zig-zagging" in the case of ill-conditioned objective Hessians. Newton's method converges locally very quickly, but is not applicable since second derivatives are only available as very rough estimates at exceedingly high cost. Quasi-Newton methods approximate the Hessians from implicit finite differencing of gradients, and suffer from amplification of stochastic noise.

Multilevel optimization exploits an available numerically cheap coarse model. Although every iteration of the multilevel algorithm is approximately as expensive as one of the inexact gradient descent algorithm, in the initial phase, the proposed steps can be expected to be better than steepest descent because the coarse model can capture nonlinearities of the fine model. Asymptotically, the coarse model's Hessian acts as a preconditioner for the gradient descent method on the fine model. A good agreement of both models' Hessians improves the asymptotic convergence rate on ill-conditioned problems and avoids the zig-zagging that slows down gradient methods applied to general ill-conditioned problems. The proposed steps are then closer to Newton steps and improve the asymptotic convergence rate on ill-conditioned problems.

Non-monotone stochastic methods such as stochastic gradient descent, ADAM, AdaGrad etc. aim to provide cheap, inexact gradient approximations and omit acceptance tests [47]. While this is numerically cheap, it sacrifices monotonicity and explicit convergence control, and requires a decreasing step size, which leads to slow convergence.

In general, the stochastic nature of ABMs makes monotonicity tests and gradient evaluation by finite differencing computationally expensive, which opens up room for further research.

## 4. Numerical examples

In this Section, we will present numerical examples using the optimization techniques for ABMs given in the previous section for the example objective. We start with a brief parameterization of the models in Section 4.1. A comparison of the multilevel algorithm with the inexact gradient descent algorithm follows in Sections 4.2 and 4.3 along our guiding examples, i.e., the H/ABM and GERDA.

### 4.1. Parameterization

In order to anchor the examples in some empirical orders of magnitude, we considered figures relating to the beginning of the pandemic in Germany. For the model we choose the German municipality of Gangelt. In each experiment we consider a time frame of seven weeks, i.e., $T = 1\,176$. We set the constant $I_{\max}$ to the value 0.5%, so that it refers to the carrying capacity of the health care system in terms of total available intensive care beds. For the upper bound on the fraction of work that can be done at home, we choose an educated guess of $u_{\mathrm{w}}^{\max} = 0.81$, based on the fact that not all jobs can be moved home or closed (e.g., sales employees or health care workers). The weights given to the costs of homeworking and homeschooling versus the weight given to the number of infected are subject of a societal or political debate that we do not intend to enter here. Therefore, an example is defined rather than opting for a multicriteria analysis that would leave the choice of weights up to a potential user of the results obtained. Thus, we set $a_{\mathrm{s}} = a_{\mathrm{w}} = 1$.

Every algorithm is started from $u_0 = 0_{n_\mathrm{u}}$, where $0_{n_\mathrm{u}}$ denotes a vector of length $n_\mathrm{u}$ with zeros only. All parameters relevant to the models are summarized in Table A.2, while the parameters relevant to the optimization algorithms are summarized in Table A.3. A detailed derivation of all parameters can be found in Appendix A.

### 4.2. Optimal policies for the H/ABM

Due to the very long runtime of GERDA, a detailed numerical analysis of the algorithms would not be justifiable for this model. Therefore, in this section, we use the H/ABM introduced in Section 2.2.2 that has a strongly reduced computational cost to compare the multilevel algorithm with the inexact gradient descent algorithm. As a termination criterion, we choose either 15 iterations or when the sample size $n$ of the objective estimate (4) or gradient estimate (8) is more than one million samples, whichever comes first.

First, we present the policies computed by the multilevel and inexact gradient descent algorithm for two different cases: constant policies that do not change for the entire time horizon of seven weeks, and piece-wise constant policies that change every 168 h (i.e., weekly changing). Fig. 2 shows the computed policies as well as the expected time course of the H/ABM for a time frame of 7 weeks. The solid line represents the solution of the multilevel algorithm and the dashed line represents the solution of the inexact gradient descent algorithm. Since all solutions are qualitatively similar (i.e., $J^{\mathrm{MLO}}(u_{\mathrm{const}}) \approx 115$, $J^{\mathrm{IGD}}(u_{\mathrm{const}}) \approx 116$, $J^{\mathrm{MLO}}(u_{\mathrm{weekly}}) \approx 111$, $J^{\mathrm{IGD}}(u_{\mathrm{weekly}}) \approx 111$), we expect qualitatively similar results in the expected time courses due to conditioning, i.e., the trajectories in Fig. 2(a) for the fractions of agents are visually indistinguishable for all computed solutions. Both for constant and piece-wise constant policies, the share of homeschooling is fairly high, while the share of homeworking is quite low. For the piece-wise constant policy, the share of homeworking is close to zero after the first week, but at the expense of homeschooling. Any non-zero control causes the number of infections – and thus the trajectories in Fig. 2(a) – to be lower than in the uncontrolled case; compare with Fig. 1(b).

Next, we take a look at the path that both algorithms are taking. Fig. 3 shows the Monte Carlo evaluation of $\mathbb{E}_n[J^{\mathrm{ABM}}(u)]$ for $n = 100$ and constant policies. Note that an explicit approximation as done here is only possible for $n_\mathrm{u} = 2$. For $n_\mathrm{u} > 2$ it becomes computationally (and

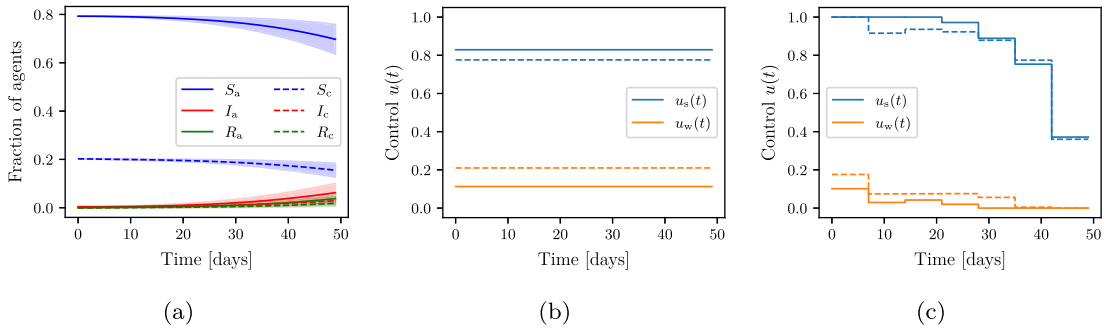


**Fig. 2.** (a) Expected time courses of the H/ABM and corresponding optimal policies computed using the multilevel algorithm (solid) compared to the policies computed using the inexact gradient descent algorithm (dashed) for (b) constant and (c) weekly changing policies. Due to conditioning, the trajectories in (a) are visually indistinguishable for the policies shown in (b) and (c). Any non-zero policy will result in a lower number of infections in (a) compared to the uncontrolled case.

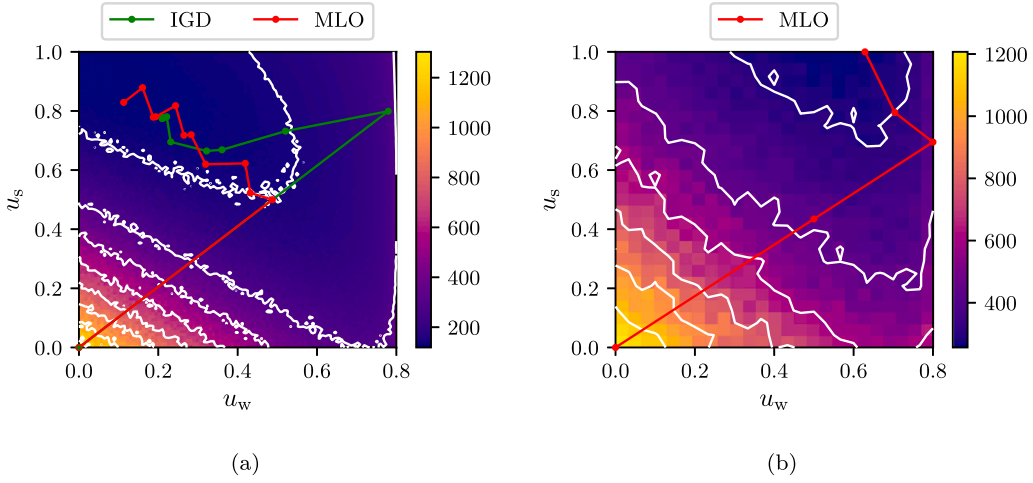


**Fig. 3.** Monte Carlo evaluations $\mathbb{E}_n[J^{\mathrm{ABM}}(u)]$ for $n = 100$ and $u = [u_s, u_w]^\top \in \mathbb{R}^2$, i.e., the policies are constant for the entire simulation time $T$, and optimization path given by iterates $u_k$ computed using the multilevel and inexact gradient descent algorithms. The objective $J^{\mathrm{ABM}}$ in Figure (a) is evaluated for the H/ABM and in Figure (b) for the GERDA model. It can be clearly seen that both objectives have different minimizers.

visually) infeasible. We see that, as expected, the multilevel algorithm is initially superior to the inexact gradient descent and takes more advantageous steps towards the minimizer $u^\star$ of $J^{\mathrm{ABM}}$ as it respects the underlying nonlinear structures of H/ABM. The first step of the inexact gradient descent algorithm is just barely still admissible. As it continues, the objectives $J^{\mathrm{ABM}}$ computed by the algorithms descend, however, at different rates. When approaching the optimum, the changes with respect to $u^\star$ get less pronounced (cf. Fig. 4). Especially the inexact gradient descent algorithm slows down significantly and barely makes progress. This can also be seen in Fig. 5, which shows the convergence. The error is defined as $\mathrm{err} := |\mathbb{E}[J^{\mathrm{ABM}}(u_k)] - \mathbb{E}[J^{\mathrm{ABM}}(u^\star)]|$, where $u^\star$ denotes the true minimizer of $J^{\mathrm{ABM}}$ and $u_k$ the current iterate. The true minimizer $u^\star$ is unknown and can only be approximated. Thus, we set $u^\star = u^{\mathrm{MLO}}_{k_{\max}}$, where $u^{\mathrm{MLO}}_{k_{\max}}$ denotes the final iterate using the multilevel algorithm since it can be shown that $J^{\mathrm{ABM}}$ is locally convex (cf. Fig. 3(a)) and that $u_{k_{\max}} \leq u^{\mathrm{IGD}}_{k_{\max}}$ holds. The figures confirm that in both cases the multilevel algorithm outperforms the inexact gradient descent algorithm in the early phase. In the later phase the increasing number of ABM simulations is driven by the choice of $e$, which controls the number $n$ required for the estimates of $\mathbb{E}[J^{\mathrm{ABM}}(u)]$ to pass the acceptance test with high confidence. Fig. 6, which shows the cumulative sum of ABM simulations, illustrates this very well. It also shows that both algorithms become expensive at about the same rate.

### 4.3. Optimal policies for GERDA

We now consider the GERDA model. To handle the massive computational effort, we use a reduced version of the municipality of Gangelt with about 1 000 agents. In this way, one agent in GERDA is roughly equivalent to 10 people in the real world (as of December 2021). This allows computing objectives and gradient estimates as in Section 3.4 for large $n$ in reasonable time on a high performance computer.

Fig. 7 shows the optimal policies for both scenarios, i.e., constant and weekly changing policies. As for the H/ABM, the difference in the effects of constant and piece-wise constant policies on the trajectories are rather negligible for GERDA. However, despite appropriately adjusted parameterization of the two models, we notice that the minimizers $u^\star_{\mathrm{H/ABM}}$ and $u^\star_{\mathrm{GERDA}}$ are different; compare the solutions shown in Figs. 3 or 7. More precisely, in both scenarios the policies $u = [u_s, u_w]^\top$ are higher, especially the share of homeworking $u_w$. Remarkably, for the weekly changing scenario, there is an extreme jump in both policies in the second week of approximately 80 percentage points. The "delay" of one week and the subsequent jump are due to the incubation period modeled in GERDA. In the model this results in low values for the policies; in the real world, with infection rates in the low double digits, a hard lockdown in the first seven days as shown in Fig. 2 for the H/ABM could lead to resentment or even noncompliance from the public. In the further course, both policies are steadily scaled back. The policies calculated for GERDA are more in line with what was implemented by many governments in early 2020. This shows the need to optimize detailed ABMs directly, since coarser models – even well-fitted ABMs – seem to exhibit a crucial loss of information.

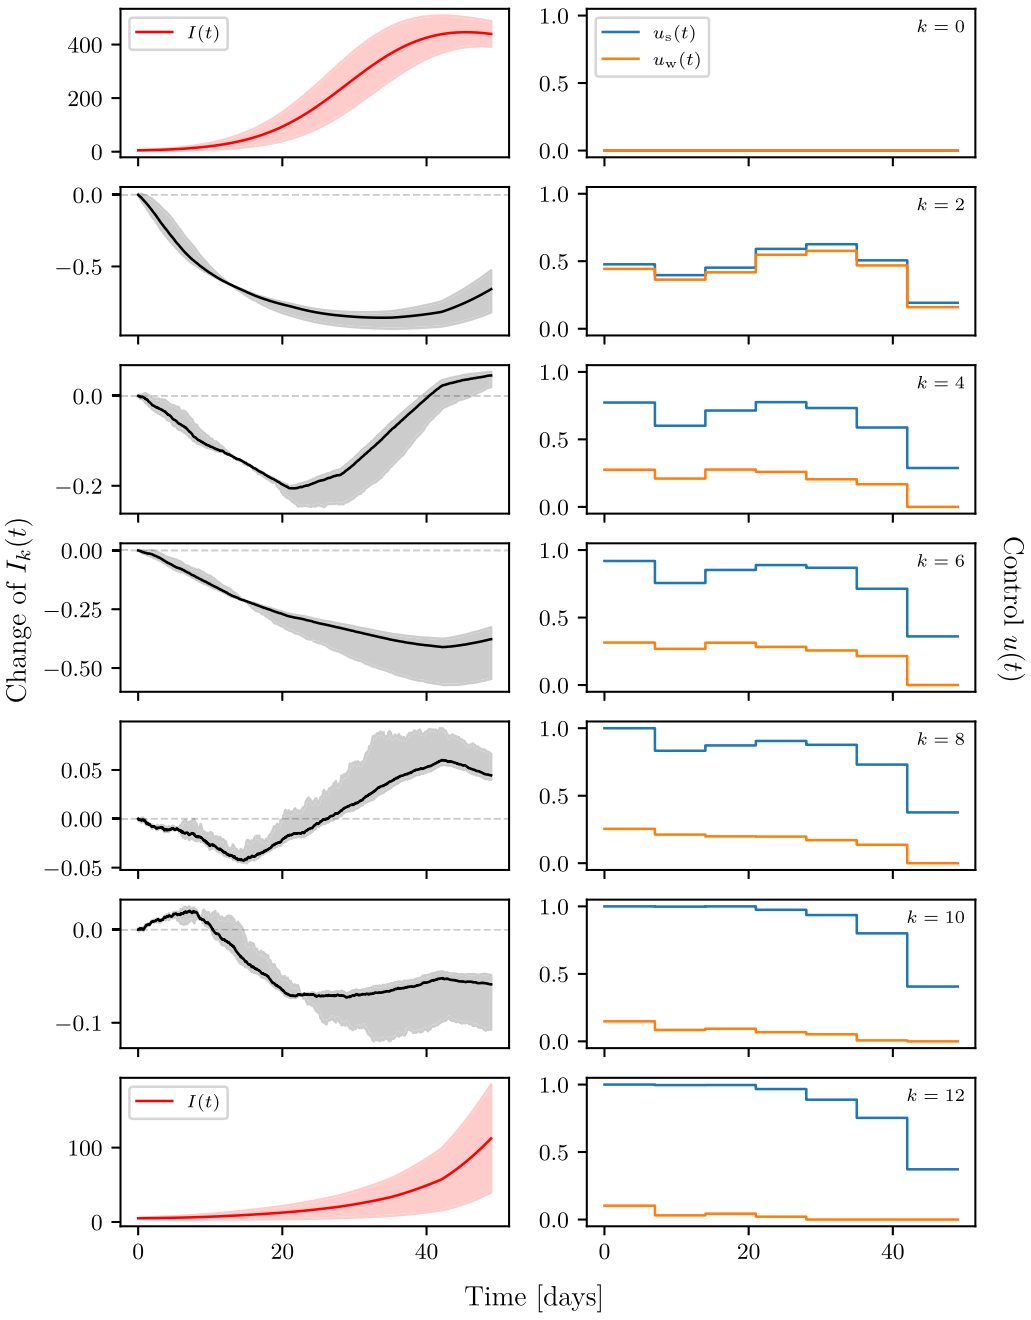


**Fig. 4.** Expected number (red/solid) and standard deviation (red/shaded area) of infected agents $I(t)$ for the corresponding policies at iteration $k$. Expected relative change (black/solid) and standard deviation (gray/shaded area) to the iteration $k-2$ (i.e., the previous figure). Data estimated from 10 000 H/ABM simulations.

The GERDA model computations emphasize the need for an efficient algorithm that quickly gets close to a solution: While for the H/ABM the inexact gradient descent algorithm is a viable option, it is not for GERDA. Here, the multilevel optimization approach is more efficient, although it also requires approximately 1.2 million GERDA simulations after the third iteration step. However, it comes with a significant reduction from $J^{\mathrm{ABM}}(u_0) \approx 1\,189$ to $J^{\mathrm{ABM}}(u_3) \approx 157$ after three iterations.

In the later phase, the agreement of the Hessian matrices of the fine and coarse models dominates the asymptotic convergence. Table 1 summarizes theoretical and experimental convergence rates for different coarse models acting as preconditioner. For the case $n_{\mathrm{u}} = 2$, the theoretical convergence rates were obtained by computing the condition numbers $\kappa$ of the Hessian matrices $H$ of a two-dimensional polynomial fit of degree five of the objective $J^{\mathrm{ABM}}$ over the entire feasible optimization domain $U$ for the respective models. For the case $n_{\mathrm{u}} = 4$, we used a local four-dimensional quadratic polynomial fit of $J^{\mathrm{ABM}}$ in an $\varepsilon$-environment around $u^\star$. In all cases the preconditioned problem leads to a faster convergence. The results are confirmed by the convergence plots (cf. also Figs. 5 for the H/ABM).

## 5. Conclusion

We presented a heterogeneous multilevel optimization approach combining a fine-level ABM with a coarse-level ODE to find exemplary non-pharmaceutical interventions in epidemic policy design. We compared this method with state-of-the-art algorithms applicable to our optimization problem. The multilevel algorithm is expected to be faster with respect to the number of iterations as it (i) captures the nonlinear structures better than a first- order second-order Taylor

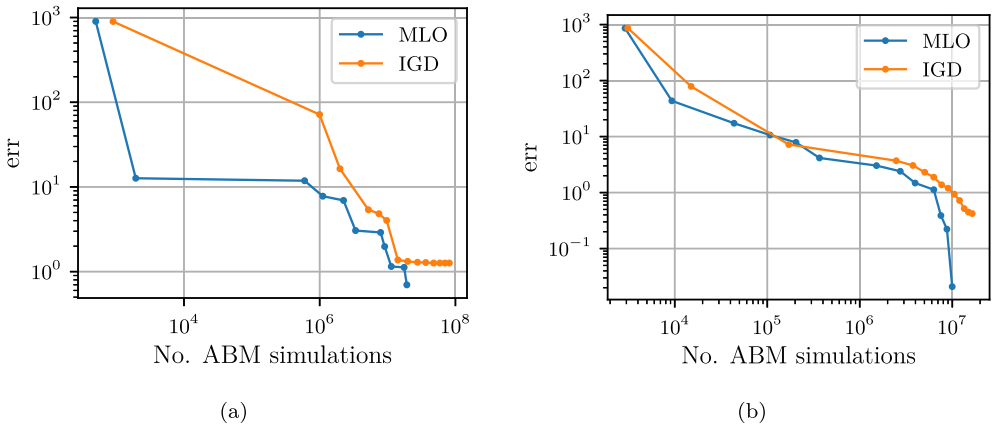


**Fig. 5.** Convergence of the multilevel and inexact gradient descent algorithms with respect to the number of ABM simulations for (a) constant and (b) piece-wise constant policies, i.e., weekly changing.

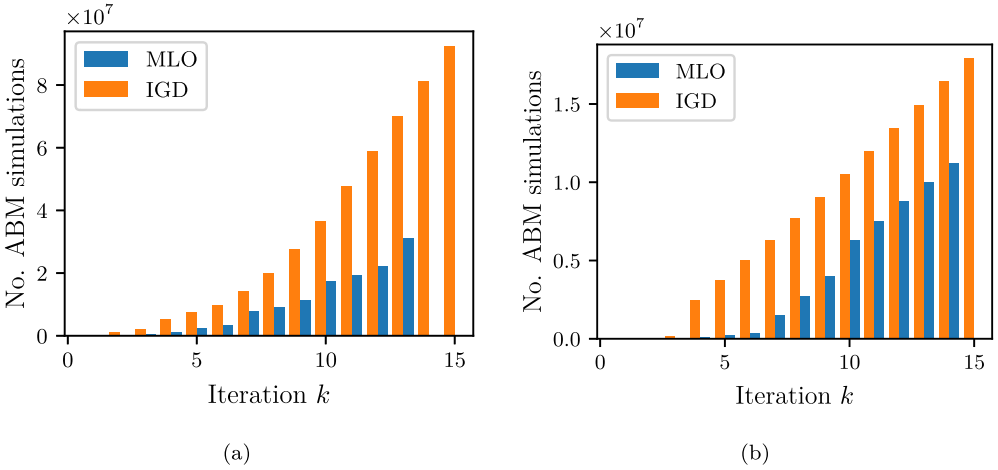


**Fig. 6.** Number of ABM evaluations for the H/ABM with (a) constant controls and (b) piece-wise constant controls, precisely weekly changing controls. The multilevel algorithm performs better than the inexact gradient descent algorithm at the beginning in terms of the number of iterations. In the further course, both algorithms become expensive at about the same rate.

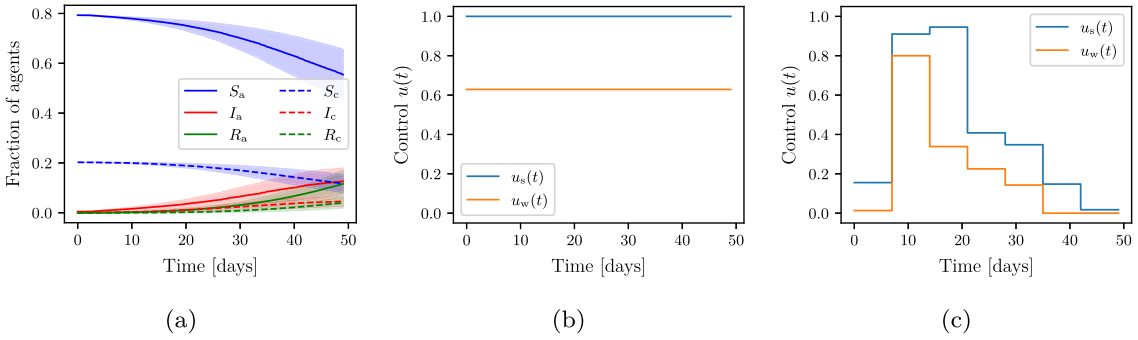


**Fig. 7.** (a) Expected time courses of GERDA and corresponding optimal policies computed using the Multilevel Optimization Algorithm 3 for (b) constant and (c) weekly changing policies.

model in the early iteration, and (ii) a well-matched coarse model serves as preconditioner and thus reduces "zig-zagging" slowing down gradient methods.

However, the theoretical speedup of the multilevel algorithm compared to the inexact gradient descent algorithm is mostly compensated by the exploding number of samples needed for a high confidence in the descent direction the closer the current iterate gets to the true minimizer. This is also the main bottleneck for the inexact gradient descent algorithm. In the initial phase, though, the superiority of the multilevel algorithm is apparent, as it leads to solutions that may already be sufficient in a real-world scenario, where the mathematical optimum need not or cannot be reached.

Additionally, we showed that the optimal policies obtained using directly the ABM differ drastically from one using the ODE model or H/ABM, which suggests that using detailed ABMs directly for the design of optimal policies is beneficial for good results. The above optimization framework can be applied not only to epidemiological models, but to any ABM where appropriate and meaningful controls can be found and applied, and suitable reduced models are available. Future research will

**Table 1**
Theoretical and experimental convergence rates for the multilevel algorithm for constant policies, i.e., $n_u = 2$ and $T = 1\,176$, and 25 days constant policies, i.e., $n_u = 4$ and $T = 1\,200$. The empirical convergence rate $\rho_{\text{ODE, fine}}$ is obtained using the ODE model as coarse-level and the H/ABM respectively GERDA as fine-level models.

| | | Fine model | | |
|---|---|---|---|---|
| | | H/ABM | GERDA | H/ABM |
| Dimension | $n_u$ | 2 | 2 | 4 |
| Condition | $\kappa(H_{\text{ODE}}^{-1} H_{\text{fine}})$ | 3.0300 | 116.8909 | 21.9801 |
| | $\kappa(H_{\text{H/ABM}}^{-1} H_{\text{fine}})$ | – | 7.0511 | – |
| | $\kappa(H_{\text{fine}})$ | 35.1841 | 235.8803 | 131.8860 |
| Convergence rate | $\rho_{\text{ODE}}$ | 0.5037 | 0.9830 | 0.9130 |
| | $\rho_{\text{H/ABM}}$ | – | 0.7516 | – |
| | $\rho_{\text{fine}}$ | 0.9447 | 0.9916 | 0.9850 |
| Experimental convergence rate | $\rho_{\text{ODE, fine}}$ | 0.7433 | 0.1934 | 0.6854 |

**Table A.2**
Parameters used for the model simulations. Here, the infection rates $r_{*\circ}$ denote the control independent parameters $r_{*\circ}$ for the ODE and $\tilde{r}_{*\circ}$ for the H/ABM, respectively. Note that the number of agents for the GERDA model is subject to random changes with each newly generated world.

| Parameter | Model | | |
|---|---|---|---|
| | GERDA | ODE | H/ABM |
| Time $T$ [hours] | 1 176 | 1 176 | 1 176 |
| Number of agents $N$ | 1 091 | – | 1 091 |
| Initially infected $I_a(0)$ | 5 | 5/1 091 | 5 |
| Initially infected $I_c(0)$ | 0 | 0 | 0 |
| World | Gangelt (reduced) | – | – |
| General infectivity | 0.175 | – | – |
| Gen. interaction frequency | 1 | – | – |
| Infection rate $r_{aa}$ | – | $1.0252 \times 10^{-12}$ | $1.1185 \times 10^{-9}$ |
| Infection rate $r_{ac}$ | – | $4.8804 \times 10^{-4}$ | $5.3246 \times 10^{-1}$ |
| Infection rate $r_{cc}$ | – | $6.1482 \times 10^{-13}$ | $6.7077 \times 10^{-10}$ |
| Recovery rate $r_a$ | – | $4.2148 \times 10^{-2}$ | $4.2148 \times 10^{-2}$ |
| Recovery rate $r_c$ | – | $4.3427 \times 10^{-2}$ | $4.3427 \times 10^{-2}$ |
| Immunity loss $\mu$ | – | – | 0.2 |

**Table A.3**
Parameters used for the optimization algorithms.

| Parameter | |
|---|---|
| Fraction of descent $c_1$ | 0.1 |
| Maximum trust region radius $\rho_0$ | 0.5 |
| Maximum finite differencing step size $h_{\max}$ | 0.1 |
| Accuracy $\epsilon$ | 0.25 |
| Health care system's capacity threshold $I_{\max}$ | $0.005N$ |
| Threshold economic impact $u_w^{\max}$ | 0.81 |
| Initial $u_0$ | $0_{n_u}$ |
| Weights $a_s = a_w$ | 1 |

address the efficient gradient approximation of stochastic dynamical systems such as ABMs.

## Funding

This research has been funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy MATH+: The Berlin Mathematics Research Center (EXC-2046/1, project ID: 390685689).

## CRediT authorship contribution statement

**Jan-Hendrik Niemann:** Investigation, Methodology, Software, Visualization, Writing – original draft, Writing – review & editing. **Samuel Uram:** Investigation, Methodology, Software, Visualization, Writing – original draft. **Sarah Wolf:** Conceptualization, Investigation, Methodology, Writing – original draft, Writing – review & editing. **Nataša Djurdjevac Conrad:** Conceptualization, Investigation, Methodology, Writing – original draft, Writing – review & editing. **Martin Weiser:** Conceptualization, Investigation, Methodology, Writing – original draft, Writing – review & editing.

## Declaration of competing interest

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

## Data availability

The data that support the findings of this study are generated based on the models given in Section 2 and do not use any external data. The algorithms are implemented in MATLAB and Python and can be found at https://github.com/Jan-HendrikNiemann/MLoptABM.

## Appendix A. Detailed parameterization

According to [48], who analyze intensive care unit loads for Berlin, Madrid and Lombardy, the fraction of infected in Berlin that needed intensive care in the first wave of the pandemic was 6 %. We use this number as an approximation for Germany and hence Gangelt. From mid April 2020, intensive care beds had to be registered with a central agency. In the second half of April, the numbers of beds registered ranged around 32 000 in Germany, see the daily reports for this time span under [49]. With this value of intensive care beds in Germany, if 6 % of the infected should not exceed this number, the maximum of infected that could be allowed would be around 530 000 persons, which corresponds to about 0.6% of the population of about 83 million. As not all intensive care beds can be allocated to COVID-19 patients since there are other reasons for needing such a bed, we reduce this number to 0.5%. Thus, we set the health care system's carrying capacity to $I_{\max} = 0.005$ N. For the upper bound on the fraction of work that can be done at home, we choose an educated guess of $u_w^{\max} = 0.81$, based on the following considerations. According to [50], the share of office jobs is below 36.7%, but a lockdown can also affect non-office jobs, e.g., sales or industrial employees, or employees in the service sector. On the other hand, there are jobs that even in a lockdown cannot be closed down or done from home. We consider the examples of workers in the health care system (5.7 million employed), food sales (552 200 employed), as well as police and fire departments (1.7 million, all numbers from [51]). These sum to about 8.5 million, representing about 17.5% of the approximately 45 million people employed in Germany. To account for further cases not included here, we increase this number to 19 %, meaning that the economic costs explode when approaching 81 % of work not done at the workplaces.

## Appendix B. Adjoint gradient computation

The most efficient way to compute the gradient $\nabla J^{\text{ODE}}(u)$ in line 3 in Algorithm 1 is to solve the adjoint equations associated with the ODE system (1) and $J^{\text{ODE}}$. Writing (1) compactly as

$$\dot{y} = f(y, u), \quad y = [S_a, S_c, I_a, I_c]^\top,$$

where we may safely neglect the recovered compartments, the adjoint equation is the terminal value problem

$$-\dot{\lambda} = f_y(y, u)^\top \lambda + J_y^{\text{ODE}}(y)^\top, \quad \lambda(T) = 0,$$

where the simulated trajectory $y$, i.e., the numerically approximated ODE solution, enters as data. Then, the gradient is given by

$$\nabla J^{\text{ODE}}(u) = \int_0^T f_u(y, u)^\top \lambda \, \mathrm{d}t.$$

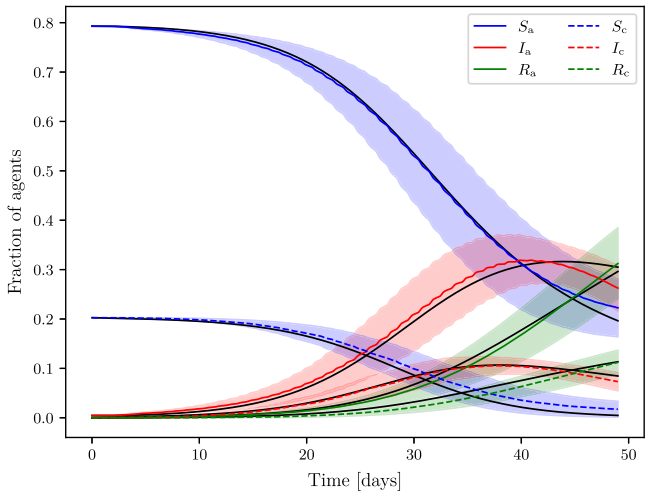


**Fig. A.8.** Solution of the fitted ODE system (1) (black, solid lines) compared to expected aggregated trajectory of GERDA (mean over 1 000 independent simulations, sample standard deviation indicated as shaded areas) for parameters given in Table A.2.

The adjoint equation can be derived from the chain rule using a trivial but clever rearrangement of terms. For more details we refer to [52].

## References


[1] L. Dell'Anna, Solvable delay model for epidemic spreading: the case of Covid-19 in Italy, Sci. Rep. 10 (2020).

[2] S. Moore, E.M. Hill, L. Dyson, M.J. Tildesley, M.J. Keeling, Modelling optimal vaccination strategy for SARS-CoV-2 in the UK, PLoS Comput. Biol. 17 (5) (2021) 1–20.

[3] M.J. Kühn, D. Abele, T. Mitra, W. Koslow, M. Abedi, K. Rack, M. Siggel, S. Khailaie, M. Klitz, S. Binder, L. Spataro, J. Gilg, J. Kleinert, M. Häberle, L. Plötzke, C.D. Spinner, M. Stecher, X.X. Zhu, A. Basermann, M. Meyer-Hermann, Assessment of effective mitigation and prediction of the spread of SARS-CoV-2 in Germany using demographic information and spatial resolution, Math. Biosci. 339 (2021) 108648.

[4] H. Wulkow, T.O.F. Conrad, N. Djurdjevac Conrad, S.A. Müller, K. Nagel, C. Schütte, Prediction of Covid-19 spreading and optimal coordination of counter-measures: From microscopic to macroscopic models to Pareto fronts, PLoS One 16 (4) (2021) 1–29.

[5] M. Marathe, N. Ramakrishnan, Recent advances in computational epidemiology, IEEE Intell. Syst. 28 (4) (2013) 96–101.

[6] F. Brauer, Mathematical epidemiology: Past, present, and future, Infect. Dis. Model. 2 (2017) 113–127.

[7] M. Tracy, M. Cerdá, K.M. Keyes, Agent-based modeling in public health: Current applications and future directions, Annu. Rev. Public Health 39 (2018) 77–94.

[8] E. Cuevas, An agent-based model to evaluate the COVID-19 transmission risks in facilities, Comput. Biol. Med. 121 (2020) 103827.

[9] A. Rodríguez, E. Cuevas, D. Zaldivar, B. Morales-Castañeda, R. Sarkar, E.H. Houssein, An agent-based transmission model of COVID-19 for re-opening policy design, Comput. Biol. Med. 148 (2022) 105847.

[10] L.R. Izquierdo, S.S. Izquierdo, J.M. Galan, J.I. Santos, Techniques to understand computer simulations: Markov chain analysis, J. Artif. Soc. Soc. Simul. 12 (1) (2009).

[11] F. Hinkelmann, D. Murrugarra, A.S. Jarrah, R. Laubenbacher, A mathematical framework for agent based models of complex biological networks, Bull. Math. Biol. (2011).

[12] S. Banisch, R. Lima, T. Araújo, Agent based models and opinion dynamics as Markov chains, Social Networks 34 (2012) 549–561.

[13] N. Djurdjevac Conrad, L. Helfmann, J. Zonker, S. Winkelmann, C. Schütte, Human mobility and innovation spreading in ancient times: A stochastic agent-based simulation approach, EPJ Data Sci. 7 (1) (2018) 24.

[14] L. Helfmann, N. Djurdjevac Conrad, A. Djurdjevac, S. Winkelmann, C. Schütte, From interacting agents to density-based modeling with stochastic PDEs, Commun. Appl. Math. Comput. Sci. 16 (1) (2021) 1–32.

[15] J.-H. Niemann, S. Winkelmann, S. Wolf, C. Schütte, Agent-based modeling: Population limits and large timescales, Chaos 31 (3) (2021) 033140.

[16] K. Wickwire, Mathematical models for the control of pests and infectious diseases: A survey, Theor. Popul. Biol. 11 (2) (1977) 182–238.

[17] D. Moualeu, M. Weiser, R. Ehrig, P. Deuflhard, Optimal control for a tuberculosis model with undetected cases in Cameroon, Commun. Nonlinear Sci. Numer. Simul. 20 (3) (2015) 986–1003.

[18] S. Duwal, S. Winkelmann, C. Schütte, M. von Kleist, Optimal treatment strategies in the context of 'treatment for prevention' against HIV-1 in resource-poor settings, PLoS Comput. Biol. 11 (4) (2015) 1–30.

[19] A. Olivares, E. Staffetti, Robust optimal control of compartmental models in epidemiology: Application to the COVID-19 pandemic, Commun. Nonlinear Sci. Numer. Simul. 111 (2022) 106509.

[20] S. Ullah, M.A. Khan, Modeling the impact of non-pharmaceutical interventions on the dynamics of novel coronavirus with optimal control analysis with a case study, Chaos Solitons Fractals 139 (2020) 110075.

[21] A. Charpentier, R. Elie, M. Laurière, V.-C. Tran, COVID-19 pandemic control: balancing detection policy and lockdown intervention under ICU sustainability, Math. Model. Nat. Phenom. 15 (2020) 57.

[22] T. Lazebnik, S. Bunimovich-Mendrazitsky, The signature features of COVID-19 pandemic in a hybrid mathematical model – implications for optimal work-school lockdown policy, Adv. Theory Simul. 4 (5) (2021) 2000298.

[23] P. Andelfinger, Differentiable agent-based simulation for gradient-guided simulation-based optimization, in: Proceedings of the 2021 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation, SIGSIM-PADS '21, Association for Computing Machinery, New York, NY, USA, 2021, pp. 27–38.

[24] C.P. Fries, Stochastic algorithmic differentiation of (expectations of) discontinuous functions (indicator functions), Int. J. Comput. Math. (2021) 1–23.

[25] S. Schutte, Optimization and Falsification in Empirical Agent-Based Models, J. Artif. Soc. Soc. Simul. 13 (1) (2010) 2.

[26] R. Miikkulainen, O. Francon, E. Meyerson, X. Qiu, D. Sargent, E. Canzani, B. Hodjat, From prediction to prescription: Evolutionary optimization of nonpharmaceutical interventions in the COVID-19 pandemic, IEEE Trans. Evol. Comput. 25 (2) (2021) 386–401.

[27] M. Oremland, R. Laubenbacher, Optimization of agent-based models: Scaling methods and heuristic algorithms, J. Artif. Soc. Soc. Simul. 17 (2) (2014) 6.

[28] G. An, B.G. Fitzpatrick, S. Christley, P. Federico, A. Kanarek, R.M. Neilan, M. Oremland, R. Salinas, R. Laubenbacher, S. Lenhart, Optimization and control of agent-based models in biology: A perspective, Bull. Math. Biol. 79 (2017) 63–87.

[29] S. Koshy-Chenthittayil, P. Mendes, R. Laubenbacher, Optimization of Agent-Based Models Through Coarse-Graining, Lett. Biomath. 8 (1) (2021) 167–178.

[30] J. She, L. Liu, W. Liu, COVID-19 epidemic: Disease characteristics in children, J. Med. Virol. 92 (7) (2020) 747–754.

[31] P. Deuflhard, F.A. Bornemann, Scientific Computing with Ordinary Differential Equations, second ed., in: Texts in Applied Mathematics, vol. 42, Springer, New York, 2002.

[32] B. Goldenbogen, S.O. Adler, O. Bodeit, J.A.H. Wodke, X. Escalera-Fanjul, A. Korman, M. Krantz, L. Bonn, R. Morán-Torres, J.E.L. Haffner, M. Karnetzki, I. Maintz, L. Mallis, H. Prawitz, P.S. Segelitz, M. Seeger, R. Linding, E. Klipp, Control of COVID-19 outbreaks under stochastic community dynamics, bimodality, or limited vaccination, Adv. Sci. 9 (23) (2022) 2200088.

[33] D.T. Gillespie, A general method for numerically simulating the stochastic time evolution of coupled chemical reactions, J. Comput. Phys. 22 (4) (1976) 403–434.

[34] T.G. Kurtz, Strong approximation theorems for density dependent Markov chains, Stochastic Process. Appl. 6 (3) (1978) 223–240.

[35] P. Atkins, J. de Paula, Atkins' Physical Chemistry, Oxford University Press, 2008.

[36] A. Shirin, Y.T. Lin, F. Sorrentino, Data-driven optimized control of the COVID-19 epidemics, Sci. Rep. 11 (1) (2021) 6525.

[37] M. Navascués, C. Budroni, Y. Guryanova, Disease control as an optimization problem, PLoS One 16 (9) (2021) 1–32.

[38] E. Hojjat, Anti-coronavirus optimization algorithm, Soft Comput. 26 (11) (2022) 4991–5023.

[39] Z. Botev, A. Ridder, Variance Reduction, American Cancer Society, 2017, pp. 1–6.

[40] A.V. Rao, A survey of numerical methods for optimal control, Adv. Astronaut. Sci. 135 (1) (2009) 497–528.

[41] J.T. Betts, Practical Methods for Optimal Control and Estimation using Nonlinear Programming, SIAM, 2010.

[42] J. Nocedal, S. Wright, Numerical Optimization, Springer, 2006.

[43] S.G. Nash, A multigrid approach to discretized optimization problems, Optim. Methods Softw. 14 (1–2) (2000) 99–116.

[44] S. Gratton, A. Sartenaer, P.L. Toint, Recursive trust-region methods for multiscale nonlinear optimization, SIAM J. Optim. 19 (1) (2008) 414–444.

[45] A.R. Conn, N.I.M. Gould, P.L. Toint, Trust-Region Methods, SIAM, 2000.

[46] J. Larson, M. Menickelly, S. Wild, Derivative-free optimization methods, Acta Numer. (2019) 287–404.

[47] M. Chau, M.C. Fu, An overview of stochastic approximation, in: M.C. Fu (Ed.), Handbook of Simulation Optimization, Springer, 2015, pp. 149–178.

[48] M. Ritter, D.V.M. Ott, F. Paul, J.-D. Haynes, K. Ritter, COVID-19: a simple statistical model for predicting intensive care unit load in exponential phases of the disease, Sci. Rep. 11 (518) (2021).
[49] Robert Koch Institut, Epidemiologischer Steckbrief zu SARS-CoV-2 und COVID-19, 2020, https://www.rki.de/DE/Content/InfAZ/N/Neuartiges_Coronavirus/Situationsberichte/Archiv_2020_tab.html.
[50] A. Hammermann, M. Voigtländer, Bürobeschäftigte in deutschland, in: Vierteljahresschrift Zur Empirischen Wirtschaftsforschung, Jg. 47 Institut der deutschen Wirtschaft IW-Trends 3/2020 (2020.
[51] Statistisches Bundesamt, Employees by industry, 2021, https://www.destatis.de/DE/Themen/Arbeit/Arbeitsmarkt/Erwerbstaetigkeit/Tabellen/arbeitnehmer-wirtschaftsbereiche.html.
[52] A. Griewank, A. Walther, Evaluating Derivatives: Principles and Techniques of Algorithmic Differentiation, SIAM, Philadelphia, 2008.

**Dr. Jan-Hendrik Niemann** studied Technomathematics at the Technische Universität Hamburg and the Universität Hamburg. He received his Ph.D. in Mathematics from the Freie Universität Berlin in 2022. His work focuses on data-driven model reduction, dynamical systems, and optimization.

**Mr. Samuel Uram** studied Mathematical Modeling of Complex Systems at the Universität Koblenz–Landau. He pursued different roles as a researcher with a focus on epidemic modeling and optimization.

**Dr. Sarah Wolf** studied Mathematics at Humboldt Universität zu Berlin and obtained her Ph.D. from Freie Universität in 2010. Since her Ph.D. studies at the Potsdam Institute for Climate Impact Research, she has always been working in interdisciplinary and sustainability-related contexts. Her research interests are the mathematical foundations of agent-based models and the use of agent-based models in discussions with stakeholders on topics like green growth, a sustainable mobility transition, or pandemics.

**Dr. Nataša Djurdjevac Conrad** is a group leader at Zuse Institute Berlin. She obtained her Ph.D. in Mathematics at the Freie Universität Berlin as a scholar of Berlin Mathematical School. Her main research interests are modeling, simulation and computational analysis of complex social systems, with a particular focus on mathematical agent-based modeling, dynamics on/of networks, model reduction approaches.

**Dr. Martin Weiser** studied mathematics and computer science at the Freie Universität Berlin, where he obtained his Ph.D. in 2001. He is head of the Modeling and Simulation of Complex Processes department at the Zuse Institute Berlin, focusing on finite element methods, inverse problems, and optimization.