

Optimization of Infectious Disease Prevention and Control Policies Using Artificial Life

Khalil Al Handawi and Michael Kokkolaras

Abstract—The spread of an infectious disease such as COVID-19 is governed by complex social interactions that are challenging to model. Policy makers must take measures to control the spread of infection despite the unknowns that accompany a novel epidemic. The principles of artificial life govern the intricacies of social interaction through which diseases can spread. Agent-based models can capture these complexities for a subset of the population by defining the behavior of individual agents. While they can be computationally expensive for large populations, their outcomes are stochastic. Therefore, they can be used to test disease prevention policies, that can be difficult to simulate using deterministic approaches. We developed an agent-based model that is inspired by several interactive simulations on the internet for describing the COVID-19 pandemic. We define metrics to estimate the socio-economic cost of disease prevention policies on the population. We present a policy-making tool based on blackbox optimization and evolutionary computation that provides well-rounded intervention measures in terms of socio-economic cost and disease control. Several intervention measures are suggested by the algorithms with varying degrees of disease control and socio-economic cost. Policy makers can choose an intervention measure based on their preference. This research recommends combining computational intelligence principles and the use of mathematical algorithms for identifying the critical amount of intervention necessary to control infectious diseases and formulate intervention policies that minimize socio-economic cost.

Index Terms—Artificial life, Agent-based modeling, Epidemiology, Policy-making, Optimization, Evolutionary computation, Probabilistic methods, COVID-19.

I. INTRODUCTION

THE COVID-19 pandemic affected health, daily life, and the economy worldwide. Health risks arise in the early stages of a pandemic due to lack of a vaccine or an effective treatment. As a result, non-pharmaceutical interventions (NPIs) such as social distancing, mask wearing, and closure of schools and non-essential businesses are imperative for controlling viral transmission by reducing contact rates [1]. This applies to all infectious diseases that can spread through social contact.

Khalil Al Handawi is with the Department of Mechanical Engineering, McGill University, Montreal, QC, H3A 0C3 Canada. (e-mail: khalil.alhandawi@mail.mcgill.ca).

Michael Kokkolaras is with the Department of Mechanical Engineering, McGill University, Montreal QC H3A 0C3, Canada, and also with the Research Group in Decision Analysis (GERAD), Polytechnique Montreal, Montreal QC H3T 1J4, Canada (e-mail: Michael.Kokkolaras@mcgill.ca).

This article has supplementary downloadable material available at <https://doi.org/10.1109/TETCI.2021.3107496>, provided by the authors.

(Corresponding author: Khalil Al Handawi.)

However, such measures put a strain on the economy if imposed for prolonged periods of time, causing the general well-being of the population to diminish [1]. Epidemiologists and policy makers must evaluate their policies despite the lack of data in the early stages of the pandemic. This is further exacerbated by the limited understanding of disease dynamics and the effects of various intervention measures [2].

A method to develop policies based on assessing the effect of different NPIs on the spread of the disease and their socio-economic cost is necessary [1]. However, the advent of a new disease such as COVID-19 is accompanied with a fair amount of uncertainty surrounding its transmission and impact on society. Computational intelligence principles can offer useful insights by aggregating partial and incomplete information to assess emerging behavior [3].

In this paper, we present a policy-making tool for infectious diseases based on an emerging principle of computational intelligence known as artificial life (ALife), where the behavior of the system as a whole is not explicitly programmed and only simple rules that govern the behavior of individual system entities are implemented [4], [5]. Interaction between system entities can give rise to emergent behavior at the system level. Combining ALife with other principles of computational intelligence such as probabilistic reasoning, evolutionary computation, and optimization can provide a complete framework for guiding the emergent behavior of an ALife model of infectious diseases. The optimal policies from such a framework can help inform public health policies for combating diseases such as COVID-19.

The paper is organized as follows. The next section presents relevant background on epidemiological models and pertinent ALife adaptations, known as agent-based models (ABMs) [6]. Section III describes the ALife model used for estimating the outcomes of different intervention measures. Section IV formulates the policy-making optimization problem and suggests some numerical optimization algorithms (including evolutionary algorithms). Section V shows the outcomes predicted by the ABM for different intervention measures obtained by optimization and discusses their implications with reference to possible real-life scenarios. Section VI summarizes the key takeaways from this study, provides recommendations for policy-making, and suggests some future directions.

II. THEORETICAL BACKGROUND

In epidemiology, compartmental models involve a population that is assigned to the compartments susceptible (**S**), infectious (**I**), or recovered (**R**). The compartment **R** also includes fatalities since they can no longer infect others and

is sometimes referred to as “removed”. Population members flow from one compartment to the other based on how an infectious disease spreads [7]. This flow is computed using a set of ordinary differential equations (ODEs) that assume a homogeneously mixed population [8]–[10]. The solution to these ODEs describes the size of each compartment with time and is referred to as the epidemic curve [11], [12]. Network epidemic models use a network’s structure to assign individuals’ contacts a priori. These contacts describe to whom they can transmit infection and from whom they can contract the infection [13]. Machine learning approaches can be used to track epidemic trajectories when enough observational data are available and have been used for predicting COVID-19 incidence in different countries employing deep neural networks [14].

Compartmental models make strong assumptions about random mixing of individuals; each individual has an equal chance of making contact with another individual, while network models assume a fixed set of contacts based on a priori knowledge. In both cases, an assumption regarding the behavior of the population as a whole is made to calculate the epidemic trajectory. Such assumptions do not apply during later stages of the pandemic when infection is concentrated around urban centers and local communities [15]. Data-driven approaches (such as neural networks and other machine learning models) require sufficient prior knowledge about epidemic trajectories to train the models and avoid overfitting.

An alternative approach for modeling epidemics draws inspiration from the principle of ALife, where only knowledge about the behavior of individuals is used to study emergent behavior at the population level without making any assumptions or requiring a priori knowledge about the latter.

ABMs seek to replicate ALife and study emergent behavior using soft computing and simulation techniques and are better suited for modeling local communities and urban centers since more information is available about the behavior and daily routines of individuals as opposed to the population as a whole [15]. ABMs require a number of prerequisites for modeling dynamic and stochastic phenomena such as epidemics. These prerequisites are identified in the following section based on existing literature.

A. Agent-based Models for Simulating Epidemics

ABMs are based on modeling the behavior of individual autonomous agents to make predictions of the outcomes at the population level. They involve a synthetic population with demographic characteristics (e.g., age distribution and prior medical conditions), a social contact network, and a disease model describing the transmission rate between population members [1]. We review these prerequisites for different studies in the literature and summarize them in Table I.

Agent traits describe an agent’s compartment (susceptible, infectious, or removed), geographic location, demographics, and behavior during social contact [22]. Table I shows that several geographic locations can be included in a model, such as schools, workplaces, and households to provide realistic predictions for community outbreaks. Such locations have an impact on the agents’ contact rate.

Most ABMs in Table I consider population heterogeneity due to geography and demographics by assigning traits and behavior characteristics to the agents using random-number generators based on census data [22]. A fewer number of studies have considered time-dependant heterogeneity which captures seasonality effects [1], [16], [18], [21]. Several studies show that endogenous behavioral changes in the agents can occur due to fear of infection during an epidemic and can be used to explain the effectiveness of policies that spread awareness [17].

Table I distinguishes between two types of social contact: repeat and random. Repeat contact occurs frequently due to familial relations between agents, whereas random contact occurs less often and with random unrelated agents. The reviewed models consider social contact as the primary mechanism for infectious disease transmission.

The inherent randomness of traits and contact result in a range of possible outcomes as opposed to a single deterministic outcome. This is known as stochastic variability and resonates with the principle of ALife where several emergent behaviors are equally likely to be observed, i.e., ALife studies life as it could be and not necessarily its exact real-world counterpart [23], [24]. This behavior is useful for accommodating the lack of information surrounding infectious disease dynamics but makes predictions difficult to make. Additionally, input and model uncertainty can arise due to lack of knowledge surrounding the model’s parameters and assumptions, respectively.

Probabilistic reasoning principles are used to address stochastic variability of the outcome by estimating the mean and confidence interval using the average and standard deviation of several realizations [1], [11], [12], [12], [19]. This requires exhaustive sampling of the epidemiological model and can incur significant computational costs as model complexity and the number of agents grow. Input model uncertainty can be reduced by sampling the range of possible values for the model parameters which further exacerbates the computational cost for reducing uncertainty. Optimization and evolutionary computation can be used to efficiently explore the parameter space without the need for exhaustive sampling [3]. Finally, Model uncertainty can be reduced by cross-validation with observational data and other epidemiological models using the calibration techniques in Table I [23].

The computational cost of executing ABMs can be reduced using multi-threaded computation and active set modeling, where only infected agents and their contacts are actively computed [17], [18]. None of the reviewed models in Table I took advantage of graphics processing unit (GPU) parallelization which can be used to efficiently manipulate the relatively large tensors used to represent populations in ABMs [25], [26].

Pandemic ABMs can be used to explore the effect of different NPI policies on the epidemic curve [23]. Such measures include shielding at-risk people, testing and isolating infectious individuals, mask-wearing, and social distancing [1], [27]. Only one study in Table I considered the socio-economic impact of intervention measures by modeling their effect on trade flows and migration [17].

ABMs vary in terms of abstractness and their reliance on

TABLE I
SUMMARY OF SURVEY OF ABMS FOR MODELING THE SPREAD OF INFECTIOUS DISEASES

Feature	[1]	[16]	[17]	[18]	[19]	[15]	[20]	[21]	described ABM
Model prerequisites									
Landscape									
Single environment			✓	✓	✓		✓		✓
Multiple environments	✓	✓		✓		✓		✓	
Quarantine environment								✓	✓
Social contact network									
Repeat contacts	✓	✓	✓	✓		✓	✓	✓	✓
Random contacts	✓	✓	✓	✓		✓	✓	✓	✓
Population heterogeneity in behavior and traits									
Demographic	✓	✓		✓		✓	✓	✓	✓
Time dependant	✓	✓		✓				✓	✓
Geographic	✓	✓		✓		✓	✓	✓	✓
Endogenous behavioral changes			✓			✓		✓	✓
Model parameters calibration									
Obtained by minimizing error	✓								
Estimated from historical data	✓	✓				✓	✓	✓	
Computational considerations									
Active set modeling			✓	✓					
CPU parallelization	✓		✓	✓				✓	✓
GPU parallelization									✓
Applications to policy-making									
Intervention models									
Social distancing	✓							✓	✓
Reduced mobility		✓						✓	✓
School and work closures	✓					✓		✓	
Contact tracing					✓				
Testing capacity								✓	✓
Outcomes									
Cost of interventions						✓	✓		
Socio-economic impact			✓						✓
Number of fatalities	✓	✓	✓	✓	✓	✓	✓	✓	✓
Maximum number of infections	✓	✓	✓	✓	✓	✓	✓	✓	✓
Duration of epidemic		✓						✓	✓

empirical observations to tune their parameters. Detailed data-driven ABMs are used as an artificial lab for assessing the effectiveness of different intervention measures and policies, while abstract ABMs are used to test and demonstrate fundamental concepts such as adaptive behavior of agents [23].

B. Contributions

Based on the presented background, it can be argued that there is a need for a computationally-efficient framework for modeling infectious diseases and the socio-economic impact of intervention policies. Modeling the uncertainties associated with infectious diseases is the primary challenge in this research area. Empirically grounded and detailed ABMs, such as CovidSim [27], address model uncertainty by using complex social networks and assumptions based on data from observed populations; they have been used partially to guide policy making in the United Kingdom during the early stages of the COVID-19 pandemic.

This paper addresses input uncertainties and the stochastic variability of ABMs by combining multiple computational intelligence concepts pertaining to probabilistic reasoning and optimization (including evolutionary computation). This is done by adapting several deterministic optimization algorithms for the stochastic ABM by utilizing probabilistic estimates

of the emergent behavior that we wish to guide. The socio-economic impact of intervention policies and the infection trajectory are some of the emergent trends that are of interest to policy makers and the primary focus of this paper.

The level of complexity of an ABM should be motivated by the hypothesis it is being used to test. In our case, we use an abstract ABM to demonstrate the ability of stochastic optimization to minimize socio-economic cost of NPIs despite the stochastic variability of ABMs. Our choice of using an abstract ABM for solving intervention policy-making problems is based on the paradigm reported in [20], [28], [29]. The contributions of our work can be summarized as follows.

- We propose an ABM for simulating human systems and social networks, to study and guide emergent ALife behaviors related to the spread of infectious diseases and the socio-economic impact of different disease control and intervention measures.
- We provide an alternative implementation of ABMs based on particle dynamics suitable for parallelization and GPU computing.
- We formulate and solve pertinent stochastic optimization problems to support public health policy-making.
- We present comparative studies related to using other models (such as CovidSim) and algorithms (such as

genetic algorithms and deterministic direct search).

III. AGENT-BASED MODEL OF INFECTIOUS DISEASES AND INTERVENTION POLICIES

The proposed ABM is inspired by several models available on the internet [30]–[33]. While these models are not used to make real-life predictions on ongoing or past epidemics, they offer an interesting social contact network in which agents behave adaptively. We adopt their approach for modeling social contact since this is an important aspect that needs to be studied in more detail [23]. We were motivated by trade flow models to develop the mobility metric for estimating socio-economic cost [34]. We completely developed the other prerequisites of the ABM.

The ABM is constructed using a synthetic population of \bar{n} agents defined by a population matrix \mathbf{A}^t at time frame $t \in \{0, 1, \dots, T\}$, where T is the total number of simulation steps. \mathbf{A}^t is a $\bar{n} \times \bar{m}$ matrix composed of \bar{m} trait vectors each having \bar{n} components as follows

$$\mathbf{A}^t = [\mathbf{a}_1^t, \mathbf{a}_2^t, \dots, \mathbf{a}_{\bar{m}}^t]^T. \quad (1)$$

The components of trait vectors are given by Table II. The table also describes how they are initialized at time frame $t = 0$, i.e., provides the initial population matrix \mathbf{A}^0 . Some traits have a constant initial value, while others have their initial values sampled from the normal distribution \mathcal{N} , uniform distribution \mathcal{U} , or the binomial distribution \mathcal{B} for discrete traits. At time frame $t = 50$ the disease is seeded by choosing agent $i_0 = 1$ (patient zero) and setting their status as infected $a_{i_0,11}^{50} = I$. Since agent 1's starting position is random, the infection center is also random. The parameters n_E and v_{\max} are explained in detail in Sections III-A and III-B, respectively.

There are \bar{r} environments that describe the geographic location and behavioral pattern of each agent within them. A trait in the population matrix \mathbf{A}^t describes the membership of an agent to an environment. The population matrix \mathbf{A}^t is incrementally updated during each time frame t to simulate the behavior of the agents, the evolution of the disease's spread over time, and its impact on the population. We will describe components of our ABM starting with the social contact network. We refer to the traits describing position, velocity, and force as $z_{x,i}^t = a_{i,1}^t$, $z_{y,i}^t = a_{i,2}^t$, $v_{x,i}^t = a_{i,3}^t$, $v_{y,i}^t = a_{i,4}^t$, $F_{x,i}^t = a_{i,5}^t$, and $F_{y,i}^t = a_{i,6}^t$.

A. Social Contact Network

Social interactions in the ABM are handled by means of forces exerted on the agents that result in an alteration of their trajectories. In our model, these interaction forces arise due to the affinity of agents to repel each other. This is an example of adaptive behavior, whereby agents may attempt to distance themselves from others in a way that more closely approximates the dynamics observed in the Spanish flu. This adds another dimension of realism to the model [23].

These interactions are used to describe real-life social interactions such as social distancing. They can be adjusted according to NPIs specified by public health authorities.

We describe the repulsive force between agents using Coulomb's law for describing the interaction between electrical charges. In our model, we assume all agents possess like charges so they repel each other. We consider the force on agent i due to agent j in terms of the amplitude S_D and the distance between the agents

$$\begin{aligned} f_{x,i,j} &= \frac{S_D}{\|\mathbf{z}_i - \mathbf{z}_j\|_2^2} \cdot \frac{z_{x,i} - z_{x,j}}{\|\mathbf{z}_i - \mathbf{z}_j\|_2}, \forall j \neq i \\ f_{y,i,j} &= \frac{S_D}{\|\mathbf{z}_i - \mathbf{z}_j\|_2^2} \cdot \frac{z_{y,i} - z_{y,j}}{\|\mathbf{z}_i - \mathbf{z}_j\|_2}, \forall j \neq i, \end{aligned} \quad (2)$$

where $f_{x,i,j}$ and $f_{y,i,j}$ are the components of the force on agent i due to agent j . The vectors $\mathbf{z}_i = [z_{x,i}, z_{y,i}]^T$ and $\mathbf{z}_j = [z_{x,j}, z_{y,j}]^T$ are the coordinates of agents i and j , respectively. We obtained the net force on agent i due to $\bar{n} - 1$ agents using

$$\begin{aligned} F_{x,i} &= \sum_{j=1}^{\bar{n}} f_{x,i,j}, \quad \forall j \neq i \\ F_{y,i} &= \sum_{j=1}^{\bar{n}} f_{y,i,j}, \quad \forall j \neq i. \end{aligned} \quad (3)$$

This is illustrated in Figure 1a. Interaction forces are computed for agents that are not flagged as essential workers (who must travel frequently among the population), i.e., $a_{i,14} \neq 1$. The number of essential workers is given by the parameter n_E .

B. Dynamics

The ABM is iterated through time by advancing the time frame $t \leftarrow t + 1$. The x and y forces are used to update the x and y velocities of agents according to

$$\begin{aligned} v_{x,i}^{t+1} &= v_{x,i}^t + F_{x,i}^t \times dt \\ v_{y,i}^{t+1} &= v_{y,i}^t + F_{y,i}^t \times dt, \end{aligned} \quad (4)$$

where dt is a constant step size. A global speed limit v_{\max} is set for the agents and is chosen such that the agents do not go out of bounds with respect to their environment.

Updated velocities are adjusted according to the following condition to limit the agents' maximum speed

$$\begin{cases} v_{x,i}^{t+1} \leftarrow v_{x,i}^{t+1}, & v_i^{t+1} \leq v_{\max} \\ v_{x,i}^{t+1} \leftarrow v_{x,i}^{t+1} \times \frac{v_{\max}}{v_i^{t+1}}, & v_i^{t+1} > v_{\max} \end{cases},$$

where v^t is the speed of the agents computed from the ℓ_2 -norm of $[v_x^t, v_y^t]^T$. The same applies for the y -component of velocity $v_{y,i}^{t+1}$. The x and y coordinates of agents are updated by their velocity vectors according to

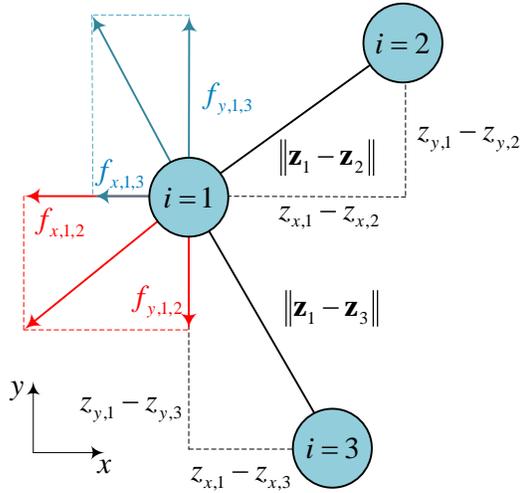
$$\begin{aligned} z_{x,i}^{t+1} &= z_{x,i}^t + v_{x,i}^{t+1} \times dt \\ z_{y,i}^{t+1} &= z_{y,i}^t + v_{y,i}^{t+1} \times dt. \end{aligned}$$

C. Environmental Interactions

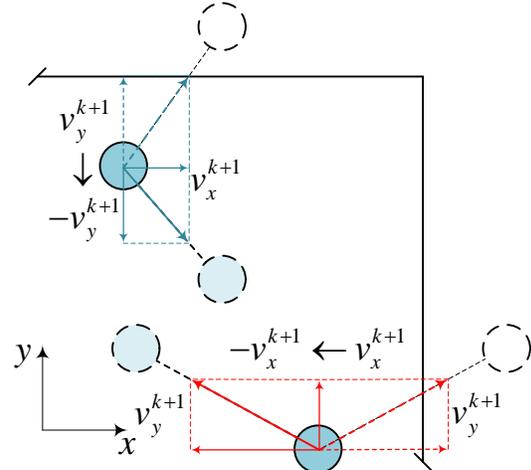
The lower and upper bounds of each environment $q \in \{1, 2, \dots, \bar{r}\}$ are given by $\mathbf{L}_q = [L_{x,q}, L_{y,q}]^T$ and $\mathbf{U}_q = [U_{x,q}, U_{y,q}]^T$, respectively. Agents are not permitted to leave the physical bounds of the current environment they are assigned to. Our model contains two environments, the main

TABLE II
DEFINITION AND INITIALIZATION OF POPULATION TRAITS

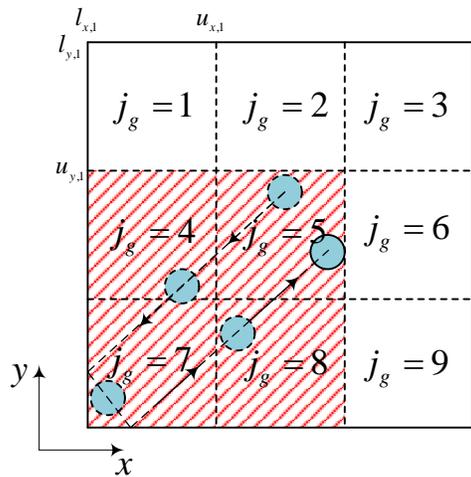
Trait	Containing set	Initial value
x -coordinate	$\{a_{i,1} \in \mathbb{R} : 0 \leq a_{i,1} \leq 1\}$	$a_{i,1}^0 \sim \mathcal{U}(a, b), a = 0, b = 1$
y -coordinate	$\{a_{i,2} \in \mathbb{R} : 0 \leq a_{i,2} \leq 1\}$	$a_{i,2}^0 \sim \mathcal{U}(a, b), a = 0, b = 1$
x -velocity	\mathbb{R}	$V_x \sim \mathcal{U}(a, b), a = -1, b = 1, a_{i,3}^0 = V_x \times v_{\max}/a_{i,7}^0$
y -velocity	\mathbb{R}	$V_y \sim \mathcal{U}(a, b), a = -1, b = 1, a_{i,4}^0 = V_y \times v_{\max}/a_{i,7}^0$
x -force	\mathbb{R}	$a_{i,5}^0 = 0$
y -force	\mathbb{R}	$a_{i,6}^0 = 0$
speed	$\{a_{i,7} \in \mathbb{R} : a_{i,7} \geq 0\}$	$a_{i,7}^0 = \sqrt{V_x^2 + V_y^2}$
infection state	$\{S, I, R, F\}$	$a_{i,8}^0 = S, a_{i,8}^{50} = I, i_0 = 1$
infected since	$\{a_{i,9} \in \mathbb{N} : 0 \leq a_{i,9} < T\}$	$a_{i,9}^0 = 0$
environment	$\{a_{i,10} \in \mathbb{N} : 1 \leq a_{i,10} \leq \bar{r}\}$	$a_{i,10}^0 = 1$
Time-independent traits		
agent index i	$\{a_{i,11} \in \mathbb{N} : 1 \leq a_{i,11} \leq \bar{n}\}$	$a_{i,11} = i$
age	$\{a_{i,12} \in \mathbb{N} : a_{i,12} \geq 0\}$	$a_{i,12} \sim \mathcal{B}(n_t, p), n_t = 81, p = 0.55$
recovery period	$\{a_{i,13} \in \mathbb{R} : 0 \leq a_{i,13} \leq 1\}$	$a_{i,13} \sim \mathcal{N}(\mu, \sigma), \mu = 0.5, \sigma = 0.5/3$
essential worker	$\{0, 1\}$	$a_{i,14} \sim \mathcal{B}(n_t, p), n_t = \bar{n}, p = n_E/\bar{n}$



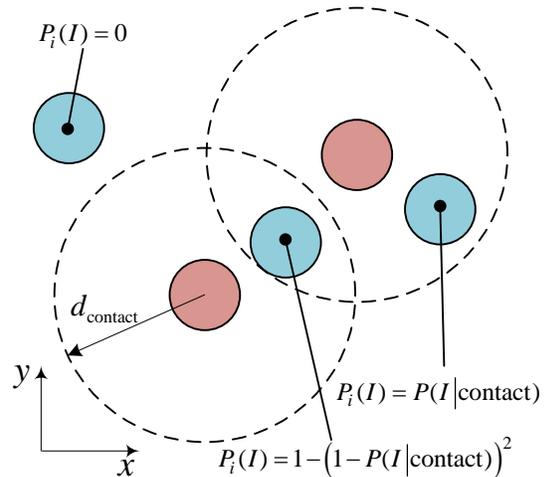
(a) Inter-agent interaction by forces



(b) Interaction of agents with host environment



(c) Mobility



(d) Infection model

Fig. 1. Illustration of a) adaptive behavior of agents using repulsive forces, b) interaction of agents with the host environment, c) mobility calculation of a single agent, and d) dynamics of disease spread from infected to susceptible agents

world and a quarantine environment referred to as “hospital”. Their physical dimensions are provided in the supplementary material (Table S.II). If the position update results in the agents being out of bounds with respect to their current environment, then the position is not updated while the velocity vector has its components reversed to simulate a “bounce”.

$$\left\{ \begin{array}{l} v_{x,i}^{t+1} \leftarrow -v_{x,i}^{t+1}, \quad \text{if } \begin{array}{l} z_{x,i}^{t+1} - L_{x,q} < 0 \\ -z_{x,i}^{t+1} + U_{x,q} < 0 \end{array} \\ v_{y,i}^{t+1} \leftarrow -v_{y,i}^{t+1}, \quad \text{if } \begin{array}{l} z_{y,i}^{t+1} - L_{y,q} < 0 \\ -z_{y,i}^{t+1} + U_{y,q} < 0 \end{array} \end{array} \right. , \quad (5)$$

This is illustrated in Figure 1b.

D. Mobility

We introduce a new epidemiological metric for estimating the socio-economic cost of intervention measures on society in the context of ABMs. This metric is given by a matrix that tracks the mobility of agents within the “world” environment. Mobility is given by the number of unique locations visited by a particular agent. Impeding travel prohibits migration from occurring thereby, restricting trade flow [17], [34]. This means that agents’ mobility is a reflection of their ability to contribute to the economy and societal relations and is inversely related to the socio-economic cost. This aspect has been over-looked in previous studies since the majority focus on modeling disease dynamics rather than socio-economic cost [1], [15], [20].

We partition the world environment into $n_{\text{grids}} = n_g \times n_g$ grids. We use the grids to compute mobility as a metric that is impacted by the disease and intervention measures. An $\bar{n} \times j_g$ matrix tracks the number of discrete grids visited within the main environment by each agent up to time frame t

$$\mathbf{G}^t = [\mathbf{g}_1^t, \mathbf{g}_2^t, \dots, \mathbf{g}_n^t]^T. \quad (6)$$

The components of \mathbf{g}_i^t are g_{i,j_g}^t , where $j_g \in \{1, 2, \dots, n_{\text{grids}}\}$. These components indicate whether agent i has visited grid j_g . The matrix \mathbf{G}^t is initialized to $\mathbf{0}$ at $t = 0$.

Grids are represented using the matrix $\mathbf{B} = [\mathbf{l}_x, \mathbf{l}_y, \mathbf{u}_x, \mathbf{u}_y]^T$, where \mathbf{l}_x , \mathbf{l}_y and \mathbf{u}_x , \mathbf{u}_y denote the lower and upper bounds of each grid j_g , respectively. The components of \mathbf{l}_x , \mathbf{l}_y , \mathbf{u}_x , and \mathbf{u}_y are calculated as

$$l_{x,j_g} = \frac{j_g \bmod n_g}{n_g}, \quad l_{y,j_g} = \frac{j_g/n_g}{n_g}, \\ u_{x,j_g} = \frac{j_g \bmod n_g + 1}{n_g}, \quad \text{and} \quad u_{y,j_g} = \frac{j_g/n_g + 1}{n_g},$$

respectively. The modulo operator $a \bmod b$ returns the remainder of the quotient a/b .

Element g_{i,j_g}^t of \mathbf{G}^t is calculated by comparing agent i ’s position with l_{x,j_g} , l_{y,j_g} , u_{x,j_g} , and u_{y,j_g}

$$\left\{ \begin{array}{l} g_{i,j_g}^{t+1} = 1, \quad \text{if } \begin{array}{l} l_{x,j_g} < z_{x,i}^{t+1} \leq u_{x,j_g} \\ l_{y,j_g} < z_{y,i}^{t+1} \leq u_{y,j_g} \end{array} \\ g_{i,j_g}^{t+1} = g_{i,j_g}^t, \quad \text{otherwise} \end{array} \right. ,$$

where $z_{x,i}^{t+1}$ and $z_{y,i}^{t+1}$ are the current x and y coordinates of agent i . We compute the mobility of agent i as the ratio of the number of visited grids to the total number of grids

$$\mu_i^t = \frac{1}{n_{\text{grids}}} \sum_{j_g=1}^{n_{\text{grids}}} g_{i,j_g}^t.$$

The average mobility of all agents is calculated as

$$M^t = \frac{1}{\bar{n}} \sum_{i=1}^{\bar{n}} \mu_i^t. \quad (7)$$

The negative of mobility $-M^t$ is used as a metric for socio-economic cost in this paper. We graphically demonstrate the computation of such a metric for a single agent in Figure 1c. The example shows that the agent has a mobility of $4/9 = 0.44$ up to the current time frame. Although the agent visited grid $j_g = 5$ twice throughout the course of its trajectory, it only counts once towards its mobility.

E. Disease Model

We model an infectious disease by a random event $P(I|\text{contact}) = p_{\text{infection}}$ that has a constant probability. A draw from the random event is made for every time frame t a susceptible agent i makes contact with an infectious agent j . Contact occurs when the distance between the susceptible and infectious agents is less than a certain threshold distance d_{contact} . This condition is given by

$$\|\mathbf{z}_i^t - \mathbf{z}_j^t\|_2 \leq d_{\text{contact}}, \forall j \neq i.$$

If agent i makes contact with several infectious agents, then the probability of agent i getting infected at time frame t is given by the probability of at least one draw of the random event $P(I|\text{contact})$ succeeding.

$$P(I) = 1 - (1 - P(I|\text{contact}))^{j_i^t},$$

where j_i^t is the number of infectious contacts that agent i has made at time frame t . This scaling of the probability of infection is similar to the scaling used by Hoertel et al. [1] for the duration of contact.

Once infected, mortality and recovery of agent i are also modeled by a random event $P(F|I) = p_{\text{mortality},i}$ and is dependent on agent i ’s age $a_{i,12}$ and host environment $a_{i,10}^t$. Age-dependant mortality is given by the piecewise function

$$\left\{ \begin{array}{l} p_{\text{mortality},i} = p_{\text{critical}} \left(1 - \frac{C - a_{i,12}}{C - R} \right), \quad R \leq a_{i,12} < C \\ p_{\text{mortality},i} = p_{\text{critical}}, \quad a_{i,12} \geq C \\ p_{\text{mortality},i} = 0, \quad a_{i,12} < R \end{array} \right. , \quad (8)$$

where p_{critical} is the critical mortality rate, C is the critical age, and R is the risk age. Agents whose age lies between the risk and critical ages are at risk of dying from the disease.

If the host environment is the “hospital”, i.e., $a_{i,10}^t = 2$, then $p_{\text{mortality},i} \leftarrow p_{\text{mortality},i} \times f_{\text{treatment}}$, where $f_{\text{treatment}} \in (0, 1)$, i.e., the mortality rate is reduced when agent i is hospitalized.

A draw from $P(F|I)$ is made if sufficient time has passed since the infection of agent i . This is done by comparing the time at which agent i has been infected $a_{i,9}$ with the current

time frame to give $\delta^t = t - a_{i,9}$. The maximum and minimum incubation periods are given by δ_{\max} and δ_{\min} , respectively. A draw from the mortality event is made if

$$\frac{\delta^t - \delta_{\min}}{\delta_{\max} - \delta_{\min}} \geq a_{i,13}. \quad (9)$$

The normal distribution of $a_{i,13}$ accounts for the statistical variance of reported incubation periods of diseases.

A subset of agents with n_T members is randomly selected for testing at every time frame t . If found infectious, and the number of hospitalizations is less than the healthcare capacity H_{\max} , then the agent is hospitalized $a_{i,10}^t \leftarrow 2$. This implies that only detected infectious individuals are isolated from the rest of the population and receive treatment. The random event where agent i gets tested is denoted by $P(T) = n_T/\bar{n}$.

At the end of time frame t , the number of infections, fatalities, and recoveries is tallied. These outcomes are denoted as n_I^t , n_F^t , and n_R^t , respectively.

F. Synthesis of the Agent-based Model

At time frame t , social interaction forces are applied on the agents, environmental interactions are used to confine agents to their host environment, and mobility, the spread of infection, and mortality are computed. The matrices \mathbf{A}^t and \mathbf{G}^t fully describe the aforementioned characteristics at time frame t .

Figure 2 provides an overview of the process followed throughout a single time step t of the agent-based simulation. The inner-most loops in this process are parallelized using a GPU program. A pseudo-algorithm is provided in the supplementary material (Section S.I). The simulation terminates when the frame counter reaches the specified number of frames T . The reviewed models in this paper terminate when zero infections remain or herd immunity is achieved. We terminate the model at a fixed time that is sufficient for allowing the epidemic to run its course. The reason for this is that the socio-economic cost continues to accumulate even when no infections are present. This makes it easy to compare the cost of different intervention measures. In a real-world setting, the true socio-economic cost can manifest years after a pandemic is over. The nominal value of all the ABM parameters are provided in the supplementary material (Table S.I).

Several emergent outcomes are of interest for policy-making (shaded in green in Figure 2). These outcomes are the maximum number of infections $n_{I,\max}$, given by

$$n_{I,\max} = \max\{n_I^t : 0 \leq t \leq T\},$$

and the average mobility M^t at frame $t = T$ denoted as M^T . Note that mobility is the opposite of cost given by $-M^T$. It can be seen that the algorithm in Figure 2 results in a stochastic process due to the random variables in Table II used for initializing the model and the random events $P(I|\text{contact})$, $P(F|I)$, and $P(T)$ (shaded in red in Figure 2).

IV. NUMERICAL OPTIMIZATION FOR POLICY-MAKING SUPPORT

Some researchers constructed empirically-grounded ABMs and used their models to predict the trajectory of an epidemic

for a few possible combinations of intervention policies [1], [15]–[17], [19], [27]. This approach may overlook some of the more effective policies for controlling the pandemic.

Numerical optimization can be used to explore the policy design space iteratively for NPIs that minimize a cost function based on the emergent outcomes of the epidemic model. Such outcomes include but are not limited to the number of averted infections, health-adjusted life years (HALYs) gained, decreased incidence rate, intensive care unit (ICU) bed occupancy, mortality rate, lockdown time, or any other aggregate measure of public health [1], [14], [20], [28], [35]. The cost of NPIs or budget constraints can be applied as well.

A. Policy-making Problem Formulation

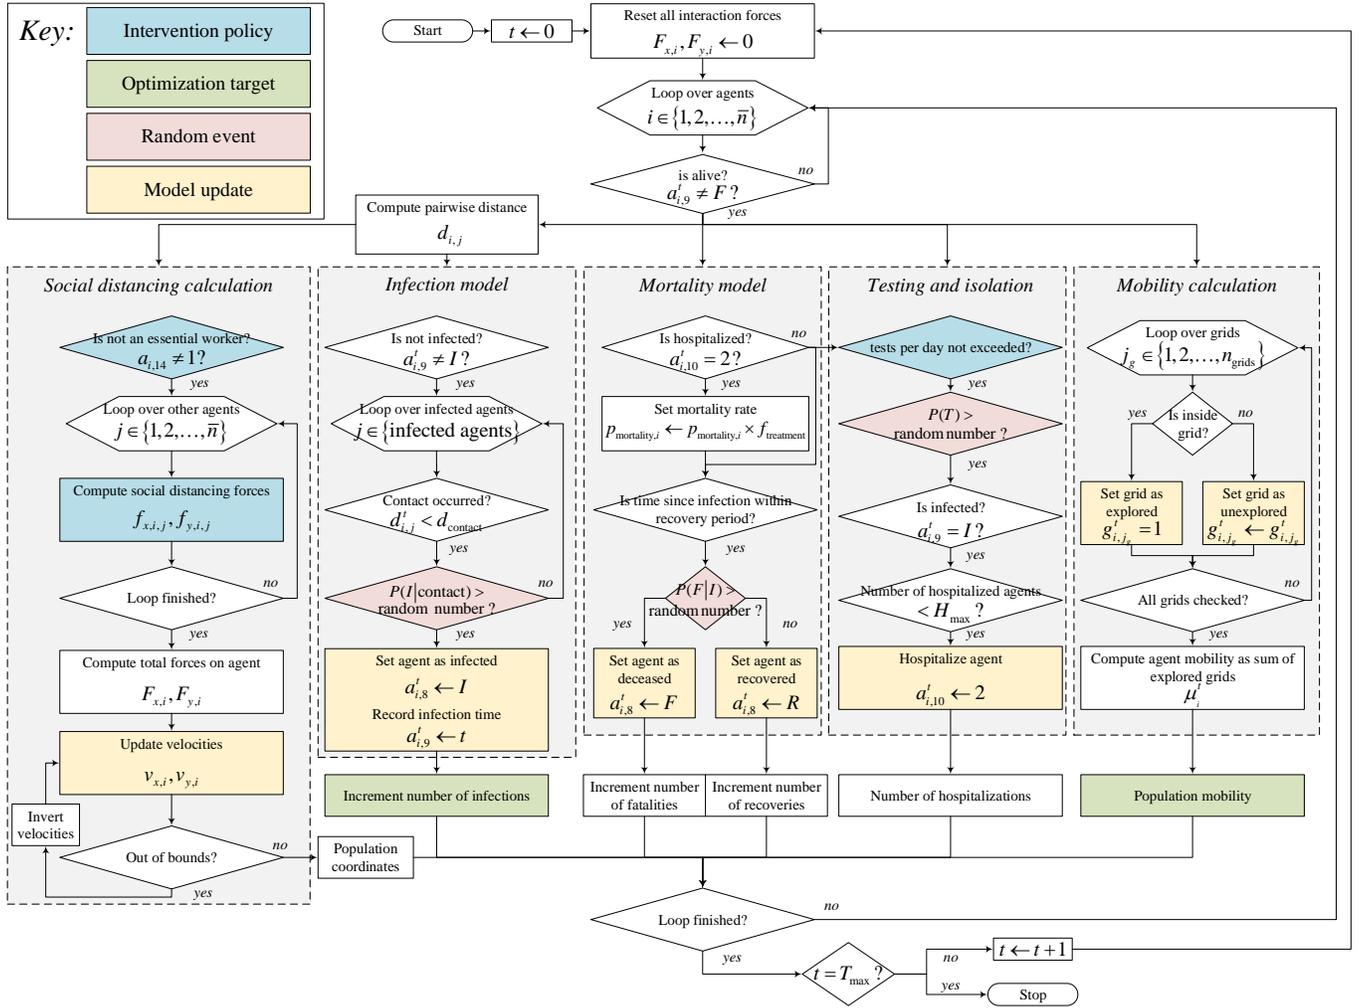
We identify some of the input parameters of the ABM in Algorithm 1 as potential optimization variables for a policy-making problem. These parameters are the number of essential workers n_E , the amplitude of the repulsive social interaction force S_D (referred to as the social distancing factor), and the number of random tests during each time frame n_T (shaded in blue in Figure 2). Increasing the parameter S_D has been cited as an effective intervention measure since it spaces agents apart reducing contact rates [1]. Increasing n_E results in more agents travelling freely among the population potentially spreading the infection. These agents are referred to as essential workers in this study. The reasoning behind this is inspired by the outbreaks in long-term care (LTC) facilities around the world due to asymptomatic infected staff working in multiple facilities [36]. Finally, n_T represents the number of daily tests and is used to test the efficacy of testing campaigns for quelling the spread of infection in its early stages. The parameter n_T is different from contact tracing since the former is a proactive measure while the latter is a reactive measure.

The stochastic optimization problem for policy-making is formulated as

$$\begin{aligned} \min_{\mathbf{x}} \quad & f(\mathbf{x}) = \mathbb{E}_{\Theta_0} [f_{\Theta_0}(\mathbf{x}) = -M^T] \\ \text{subject to} \quad & c(\mathbf{x}) = \mathbb{E}_{\Theta_1} [c_{\Theta_1}(\mathbf{x}) \equiv n_{I,\max} - H_{\max}] \leq 0, \\ \text{where,} \quad & 16 \leq n_E \leq 101, \quad 0 \leq S_D \leq 0.15, \quad 10 \leq n_T \leq 51, \\ & H_{\max} = 90, \end{aligned} \quad (10)$$

and $\mathbf{x} = [n_E, S_D, n_T]^T$ are the decision variables and represent an intervention measure. Θ_0 is a realization of the random variables and events described in Figure 2 and \mathbb{E}_{Θ_m} denotes the expectation with respect to the random variable Θ_j for all $j \in \{0, 1, \dots, m\}$, where $m = 1$ is the number of constraints. The same random variables govern the objective and constraint functions, i.e., $\Theta_0 = \Theta_1 = \Theta$. The solution to this problem results in a measure that maximizes mobility while ensuring that the health system is not overwhelmed with cases. These two emergent outcomes are of particular interest to policy makers. The maximum number of infections was set as a constraint to avoid healthcare worker infections that could lead to a vicious cycle of reduced healthcare capacity [37].

The stochastic optimization problem in Equation (10) is also solved using the CovidSim model [21], [38]. Although CovidSim does not provide a measure for socio-economic

Fig. 2. Flow chart of simulation process for a single time step t of the ABM.

impact, it does provide detailed information about infections, fatalities, and critical number of cases requiring ICU care. The objective function is provided by our ABM while the constraint is computed using CovidSim. To make this distinction, the subscript of the random variables Θ governing the objective and constraint functions is not dropped since each model is governed by different random variables and events, i.e., $\Theta_0 \neq \Theta_1$ when CovidSim is used to compute the number of infections. This problem is denoted as the unified optimization problem.

Furthermore, the variables $\mathbf{x} = [n_E, S_D, n_T]^T$ that belong to our ABM cannot be directly applied to CovidSim. We provide a mapping from the variable space of our ABM to that of CovidSim in Table III. The compliance rate, spatial contact rate, and the percentage of infectious cases isolated, denoted r_{comp} , r_{contact} , and c_I , respectively are mapped linearly onto the ABM variables n_E , S_D , and n_T , respectively, such that their lower and upper bounds coincide. The variable n_E has an upper bound equal to 10% of the ABM population and therefore corresponds to a compliance rate of 90%. A similar intuition was applied for mapping the other two variables. The other parameters of CovidSim are based on the parameters and

interventions applied in Report 9 [27].

TABLE III
OPTIMIZATION PROBLEM DEFINITIONS

Parameter	Notation	Value
Optimization problem 1		
Objective	$f_{\Theta}(n_E, S_D, n_T)$	$-M^T$
Constraint	$c_{\Theta}(n_E, S_D, n_T)$	$n_{I,\text{max}} - H_{\text{max}}$
Essential workers	n_E	$16 \leq n_E \leq 101$
Social distancing factor	S_D	$0 \leq S_D \leq 0.15$
Tests/time step	n_T	$10 \leq n_T \leq 51$
Healthcare capacity	H_{max}	90
Optimization problem 2 (unified optimization problem)		
Objective	$f_{\Theta_0}(n_E, S_D, n_T)$	$-M^T$
Constraint	$c_{\Theta_1}(r_{\text{comp}}, r_{\text{contact}}, c_I)$	$n_{I,\text{max}} - H_{\text{max}}$
Compliance rate	r_{comp}	$1 \geq r_{\text{comp}} \geq 0.9$
Contact rate	r_{contact}	$7 \geq r_{\text{contact}} \geq 1$
% Cases isolated	c_I	$0 \leq c_I \leq 1$
Healthcare capacity	H_{max}	9% of population size

B. Optimization of Stochastic Simulation-based Problems

Linear and integer programming has been used to optimize NPIs by approximating the outcomes of ODE compartmental

models using a Taylor series expansion [39]. This assumes independence of interventions with respect to the outcome that is being optimized which seldom holds for epidemics. Other studies used network-based epidemic models [29] and ABMs [20], [28] to allocate interventions for an epidemic. Several realizations of the stochastic models are used to approximate the objective and constraint functions. The previous studies have relied on deterministic gradient-based methods and/or approximations to their solutions to optimize NPIs. These approximations and optimization algorithms use deterministic algorithmic objects to ensure improvements. Such algorithms do not present convergence theory when stochastic objective and constraint functions are used resulting in suboptimal solutions.

We use the stochastic version of a direct search algorithm that is well-suited for blackbox optimization problems involving expensive simulation-based models [40], [41]. This method has the advantage that no gradient information is needed to find descent directions since gradients cannot be estimated accurately with reasonable computational effort for stochastic epidemiological models [20], [28], [29], [42]. The algorithm is grounded in martingale theory and has rigorous convergence properties [43].

The algorithm, denoted StoMADS-PB uses a progressive barrier approach for handling constraints and can be used for solving constrained stochastic optimization problems of the form

$$\begin{aligned} & \underset{x \in \mathcal{X}}{\text{minimize}} && f(x) = \mathbb{E}_{\Theta_0} [f_{\Theta_0}(x)], \\ & \text{subject to} && c_1(x) = \mathbb{E}_{\Theta_1} [c_{\Theta_1}(x)] \leq 0 \\ & && c_2(x) = \mathbb{E}_{\Theta_2} [c_{\Theta_2}(x)] \leq 0 \\ & && \vdots \\ & && c_m(x) = \mathbb{E}_{\Theta_m} [c_{\Theta_m}(x)] \leq 0, \end{aligned} \quad (11)$$

The random variables Θ_j for all $j \in \{0, 1, \dots, m\}$ are supposed to be independent with unknown possibly different distributions. The numerically unavailable objective $f(\cdot)$ and constraints $c_j(\cdot)$, are given by their noisy computable versions $f_{\Theta}(\cdot)$ and $c_{\Theta_j}(\cdot)$, respectively [41].

Feasibility of a candidate solution in the space of possible solutions $\mathcal{X} \subset \mathbb{R}^n$ is measured by the constraint violation function given as

$$h(x) := \sum_{j=1}^m \max \{c_j(x), 0\}, \quad (12)$$

where $x \in \mathcal{X}$ is feasible if and only if $h(x) = 0$.

StoMADS-PB, involves search and poll steps. The search step allows the user to engage any technique, algorithm, or heuristic (including doing nothing at all) to find promising candidates for the next iterate, while the poll step performs a localized search around the incumbent solution x_0^k defined at iteration k of the algorithm. During both steps, a finite number of trial points are generated on a discretization of the space of variables called the *mesh* defined for iteration k of the algorithm.

We focus on the implementation of the poll step in this paper since the convergence analysis of the algorithm relies on it.

The poll step generates candidate solutions $x_s^k \in \mathcal{P}^k \subset \mathcal{X}$, where \mathcal{P}^k is denoted as a frame with directions that form a positive spanning set [44].

The StoMADS-PB algorithm seeks to find a trial point whose objective function value is less than that at the current incumbent, i.e., $f(x_s^k) < f(x_0^k)$, and is feasible $h(x_s^k) = 0$. However, the true value of the objective and constraint violation functions is unknown and only realizations of their values are available to the algorithm. StoMADS-PB uses estimates of the objective and constraint violation functions at the current incumbent solution and trial points to determine whether a trial point can result in an improvement in the true objective function, constraint violation function (if incumbent is infeasible), or both. The estimates of the objective and constraint violation functions are obtained by taking the mean of n^k realizations of the noisy computable versions and are denoted by f^k and h^k , respectively. The estimate for the objective, constraint, and constraint violation functions at a point $x^k \in \mathcal{X}$ is given by

$$\begin{aligned} f^k &= \frac{1}{n^k} \sum_{i=1}^{n^k} f_{\Theta_{0,i}}(x^k), \\ c_j^k &= \frac{1}{n^k} \sum_{i=1}^{n^k} c_{\Theta_{j,i}}(x^k), \text{ and} \\ h_j^k &= \sum_{j=1}^m \max \{c_j^k, 0\}, \end{aligned} \quad (13)$$

respectively. These estimates are based on the principle of probabilistic reasoning, where the mean of several observations is used to approximate the expected outcome [3].

An improvement towards optimality is achieved if

$$f_s^k - f_0^k \leq -\gamma \varepsilon (\delta_p^k)^2 \quad (14)$$

for the objective function, while an improvement in feasibility is defined as

$$h_s^k - h_0^k \leq -\gamma m \varepsilon (\delta_p^k)^2, \quad (15)$$

where γ , and ε are constants. Feasibility is defined as

$$u_s^k := \sum_{j=1}^m \max \{c_j^k + \varepsilon (\delta_p^k)^2, 0\} = 0. \quad (16)$$

The term $-\gamma \varepsilon (\delta_p^k)^2$ (denoted as the accuracy threshold) in Equations (14), (15), and (16) ensures that any improvement in feasibility or optimality of a trial point is due to an improvement in the true unknown values of the objective and constraint violation functions and not a result of the uncertainty in the estimates f^k and c_j^k . The algorithm dynamically reduces the mesh and poll size δ_p^k by proxy if no trial point satisfying these conditions is found resulting in a smaller frame \mathcal{P}^k centered around the incumbent solution. This causes the algorithm to sample more points within the vicinity of the incumbent solution, reducing the uncertainty in its estimate.

These conditions set the StoMADS-PB algorithm apart from other deterministic optimization algorithms when applied to stochastic optimization problems. The algorithmic implementation of StoMADS-PB is provided in the supplementary material (Section S.II).

The interested reader is referred to Audet et al. [40] and Dzhahini et al. [41] for further reading on StoMADS. Details on the progressive barrier approach, how the frame \mathcal{P}^k is constructed, and how the mesh and frame sizes are updated during each iteration of the algorithm are also provided in these references [40], [41], [44], [45].

We also compare the performance of StoMADS-PB with its deterministic counterpart, mesh adaptive direct search (MADS) for solving the problem in Equation (10). The NOMAD implementation of MADS is used for the comparative studies in this paper [46]. Another comparison is made with evolutionary algorithms since they are generally easy to understand, implement, and (as a result) widely used for solving optimization problems. We use the MATLAB implementation of a genetic algorithm (GA) with an augmented Lagrangian method for constrained optimization [47]–[50].

We adapt MADS and GAs for solving the stochastic problem in Equation (10) by using the mean of n^k realizations of the stochastic objective and constraints as the deterministic value of the objective f^k and constraint c_j^k functions available to the algorithms. We use parallel processing for computing f^k and c_j^k when $n^k > 1$ to accelerate all three algorithms. The intuition behind the modified deterministic algorithms stems from the promise of combining multiple computational intelligence principles (probabilistic reasoning and evolutionary computation) to overcome the limitations of each method [3].

We explore the effect of several important algorithmic parameters, for each of the three optimization algorithms in Section V. We vary ε_f which controls the accuracy threshold of the estimates for StoMADS-PB. We compare two different configurations of NOMAD denoted as NOMAD-default and NOMAD-basic. Unlike StoMADS-PB and NOMAD-basic, NOMAD-default uses a sophisticated search step based on quadratic models before the poll step is invoked. This improves the algorithm's ability to break free of local optima [46]. In the case of the augmented Lagrangian GA, we investigate the effect of population size \bar{p} on performance of the algorithm as it determines the ability of the algorithm to explore the solution space and identify global optima. A popular choice for population size is $\bar{p} = \max\{50, 10n\}$, where n is the dimensionality of the problem [50], [51].

The settings for the algorithmic parameters of StoMADS-PB, NOMAD, and augmented Lagrangian GA are specified in the supplementary material (Section S.II). We acknowledge that the performance of an algorithm can depend on several parameters. GAs in particular should be tuned for the problem being solved by adjusting the selection, mutation, and crossover methods used to create new individuals during successive generations. Furthermore, the choice of initial penalty and penalty factors associated with the constraints in the Augmented Lagrangian method is important for balancing the importance of minimizing the objective with feasibility. However, since the emphasis of this paper is on solving policy-making problems given limited time and resources, we explore the performance of these algorithms with most of their parameters set to their recommended default values as reported in the literature.

A more accurate estimate of the objective and constraint

value is obtained by taking the average of $N = 100$ realizations every time a successful feasible iteration occurs in the case of StoMADS-PB and NOMAD, or a new best feasible individual with higher fitness is found in the case of the GA:

$$\mathbb{E}_\Theta [f_\Theta(\mathbf{x})] \approx \bar{f}_\Theta = -\frac{1}{N} \sum_{n_r=1}^N M_{n_r}^t \quad \text{and}$$

$$\mathbb{E}_\Theta [c_\Theta(\mathbf{x})] \approx \bar{c}_\Theta = -H_{\max} + \frac{1}{N} \sum_{n_r=1}^N n_{I, \max, n_r}.$$

The estimates \bar{f}_Θ and \bar{c}_Θ are used to compare the performance of all three algorithms.

V. RESULTS AND DISCUSSION

The results of the policy-making optimization problem are reported in this section for both the developed ABM and CovidSim. Comparisons between the solutions obtained by different optimization algorithms are also presented.

A. Stochastic Optimization Problem Results Using the Agent-based Model

Optimization problem 1 (defined in Table III) is solved using different optimization algorithms. Performance profiles of several optimization runs and their algorithmic parameters are shown in Figure 3. Each performance profile is the average of 4 independent optimization runs using the same algorithmic parameters. All algorithms were given a computational budget of 6000 function evaluations. Termination occurs when the default termination criterion is reached or the computational budget is exceeded. A common parameter to all three algorithms was the sampling rate n^k used to construct the estimates of the objective and constraint functions during optimization. It was increased in steps of $n^k = 1$, $n^k = 4$, and $n^k = 20$. It can be observed that increasing n^k improved the ability of all three algorithms to find a feasible solution as shown by Figures 3d, 3e, and 3f, giving merit to the approach of using probabilistic reasoning and evolutionary computation to overcome the stochastic variability of artificial life when attempting to guide its expected behavior.

Several observations can be made regarding the performance of the StoMADS-PB algorithm with respect to NOMAD and GAs. At low sampling rates ($n^k = 1$), StoMADS-PB provided more feasible solutions compared to NOMAD (Figure 3d). Furthermore, StoMADS-PB provided solutions with comparable \bar{f}_Θ and \bar{c}_Θ to those of the GA with fewer function evaluations (Figures 3a and 3d).

Higher sampling rates slowed the convergence of all three algorithms because more function evaluations are needed to construct the estimates of the objective and constraint functions. At both $n^k = 4$ and $n^k = 20$, StoMADS-PB provides more feasible solutions compared to NOMAD (Figures 3e and 3f) and more optimal solutions compared to the GA (Figures 3b and 3c). Although we do not fine tune any of the three algorithms, we conclude that StoMADS-PB provides a good balance between optimality and feasibility at reasonable computational budget compared to the other algorithms with

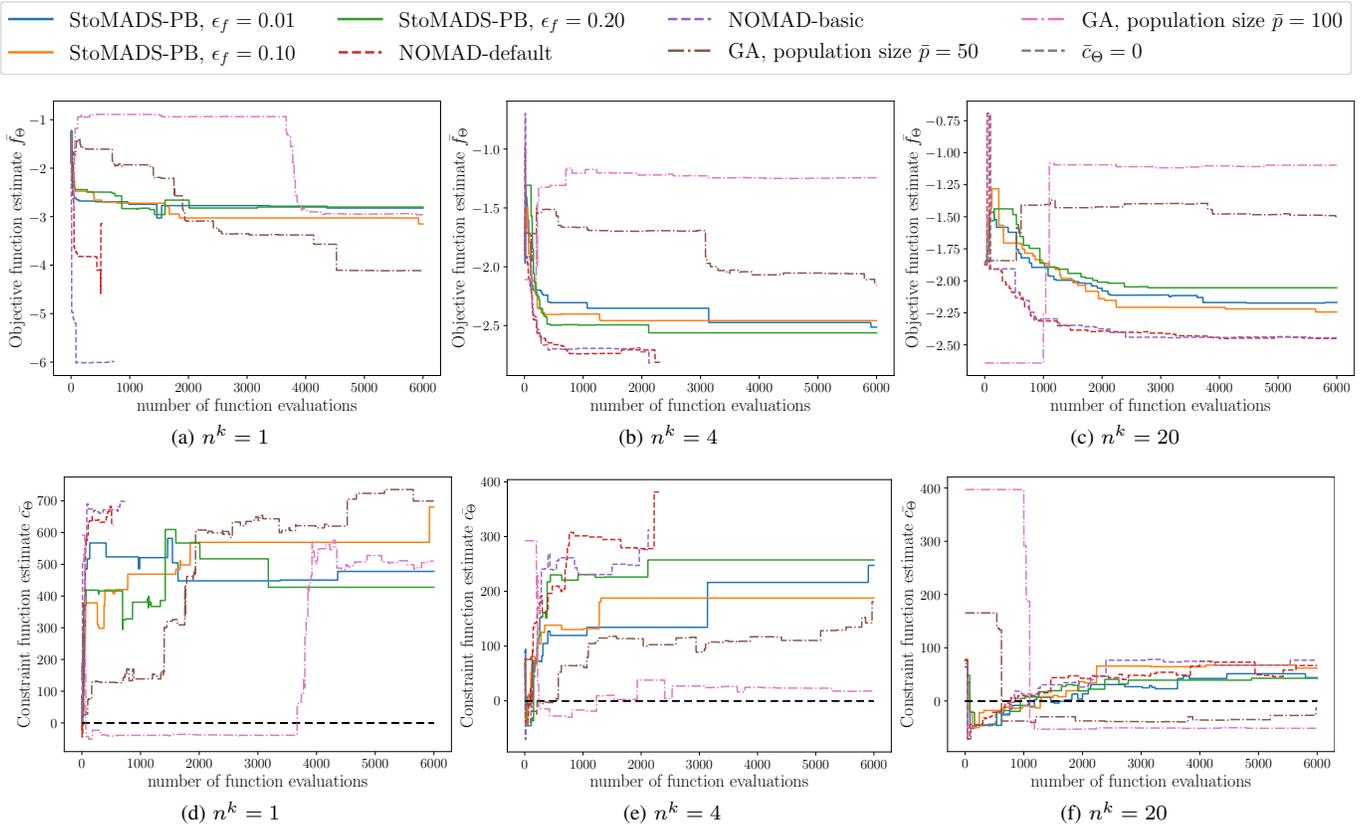


Fig. 3. Performance profiles showing (a),(b),(c) the estimated objective function and (d),(e),(f) the estimated constraint function for the best feasible incumbent solution identified by StoMADS-PB, NOMAD, and augmented Lagrangian GA.

default settings. These results show that StoMADS-PB is a good choice for stochastic problems involving complex and expensive function evaluations such as CovidSim.

We identify the best feasible solution found by each of the three algorithms by looking at the objective and constraint function estimates \bar{c}_{Θ} and \bar{f}_{Θ} . Figures 4a and 4b show the distribution of the $N = 100$ realizations used to compute \bar{c}_{Θ} and \bar{f}_{Θ} , respectively. For each algorithm, the feasible solutions ($\bar{c}_{\Theta} \leq 0$) were ranked by \bar{f}_{Θ} . Only StoMADS-PB and the GA yielded feasible solutions. This is because of the relatively small feasible space of solutions (9%) compared to the rest of the solution space. In the case of NOMAD, the least infeasible solution is chosen.

The trajectories of the epidemic and mobility are shown in Figures 4c and 4d, respectively for each solution. The StoMADS-PB solution ($\bar{f}_{\Theta} = -2.24$, $\bar{c}_{\Theta} = -2.56$) favored high testing rates n_T (close to the upper bound $n_T \leq 51$), a large cohort of essential workers n_E (close to the upper bound $n_E \leq 101$), and moderate social distancing. This implies that our model places more emphasis on large and early testing campaigns compared to social distancing and lockdowns for mitigating pandemics. This result agrees with several observational studies in the literature [12], [14]. The GA provided the most conservative solution ($\bar{c}_{\Theta} = -65.72$) out of all three algorithms but was suboptimal ($\bar{f}_{\Theta} = -0.76$). NOMAD provided the least conservative solution and performed the worst in terms of feasibility ($\bar{c}_{\Theta} = 26.69$). This is because NOMAD relied on few samples (bounded from above by the

number of poll candidates) of the objective and constraint functions during each iteration of the algorithm. StoMADS-PB resamples points within the vicinity of the incumbent solution on unsuccessful and uncertain iterations, while the GA samples a population of individuals that is gradually narrowed down to the vicinity of the fittest feasible individual reducing uncertainty in its estimated fitness and feasibility. The results show that StoMADS-PB strikes a good balance between optimality and feasibility. The numerical statistics of each solution are provided in the supplementary material (Section S.III).

B. Stochastic Optimization Problem Results Using Different Epidemiological Models

The optimal interventions obtained by solving the unified optimization problem in Table III are visualized in Figure 5. The optimal solutions have comparable values to those obtained in Section V-A. Several optimization runs with StoMADS-PB and CovidSim were performed. Two distinct feasible solutions emerged. The first solution (shown in in Figures 5a and 5b) favors more social distancing S_D (given by contact rate in CovidSim) for a large cohort of essential workers n_E (low compliance rates $\approx 90\%$ in CovidSim). Our ABM predicts a comparable median trajectory for infections but there is a large variance in the realizations from our model and as a result, this policy would not emerge as a possible solution when considering the number of infections reported by our model. The second solution (shown in Figures 5a and

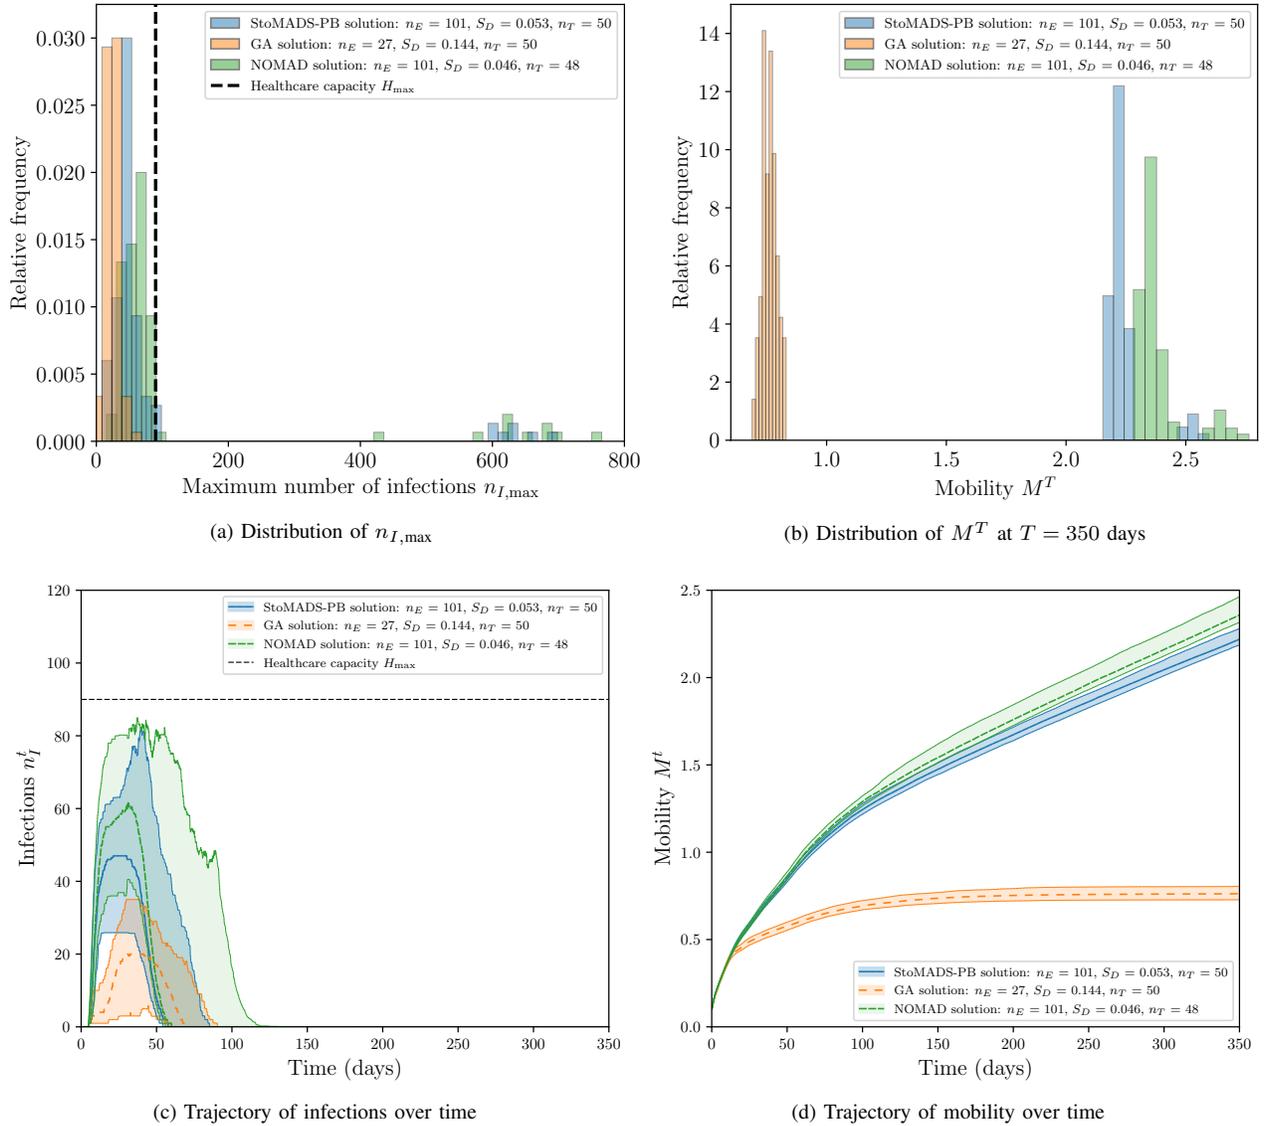


Fig. 4. ABM outcomes and statistics for the best solution obtained from several optimization algorithms. The distribution of $N = 100$ realizations for (a) the maximum number of infections and (b) the mobility. (c),(d) The solid line represents the median of $N = 100$ realizations with shaded areas representing the 10-90 percentile range.

5b) favors more testing n_T for a modest cohort of essential workers n_E (compliance rate $\approx 95\%$ in CovidSim). Our ABM agrees with this CovidSim solution since the interquartile range of the realizations for both models is within the healthcare capacity.

Policies suggested by using our ABM are in general more conservative compared to those suggested by using CovidSim. Our ABM focuses on relatively small populations and geographic areas. When compliance and testing rates are low, the effects of disease superspreaders (i.e., the essential workers or non-compliant individuals) are more pronounced in the small geographic location of our model. We also assume a uniform population density in our model (Figure 5d) which explains the higher contact rates and superspreader effects observed in our model. CovidSim covers a much larger geographic location and the population is heterogeneously distributed based on census data as shown in Figure 5c. The geographic location

studied in this paper (the Dominion of Canada) is sparsely populated. Nonetheless, the intervention policies obtained by using CovidSim can guide policy makers to more optimal interventions as given by the socio-economic impact predicted by our model (Figure 5b).

Figure 5a suggests that the results obtained using our model have a larger variance compared to those obtained using CovidSim. This is because the random interactions in the larger population in CovidSim ($\bar{n} = 36,460,098$) average out to more closely resemble the underlying trend. Furthermore, the random seed associated with the setup of the social contact network in CovidSim was fixed. The model was randomly seeded at runtime only; this may contribute to lower variance. This was done because the initial setup for assigning individual geographic, demographic, and behavioral characteristics is an expensive operation. For this reason, the developers of CovidSim recommend fixing the social contact network's seed

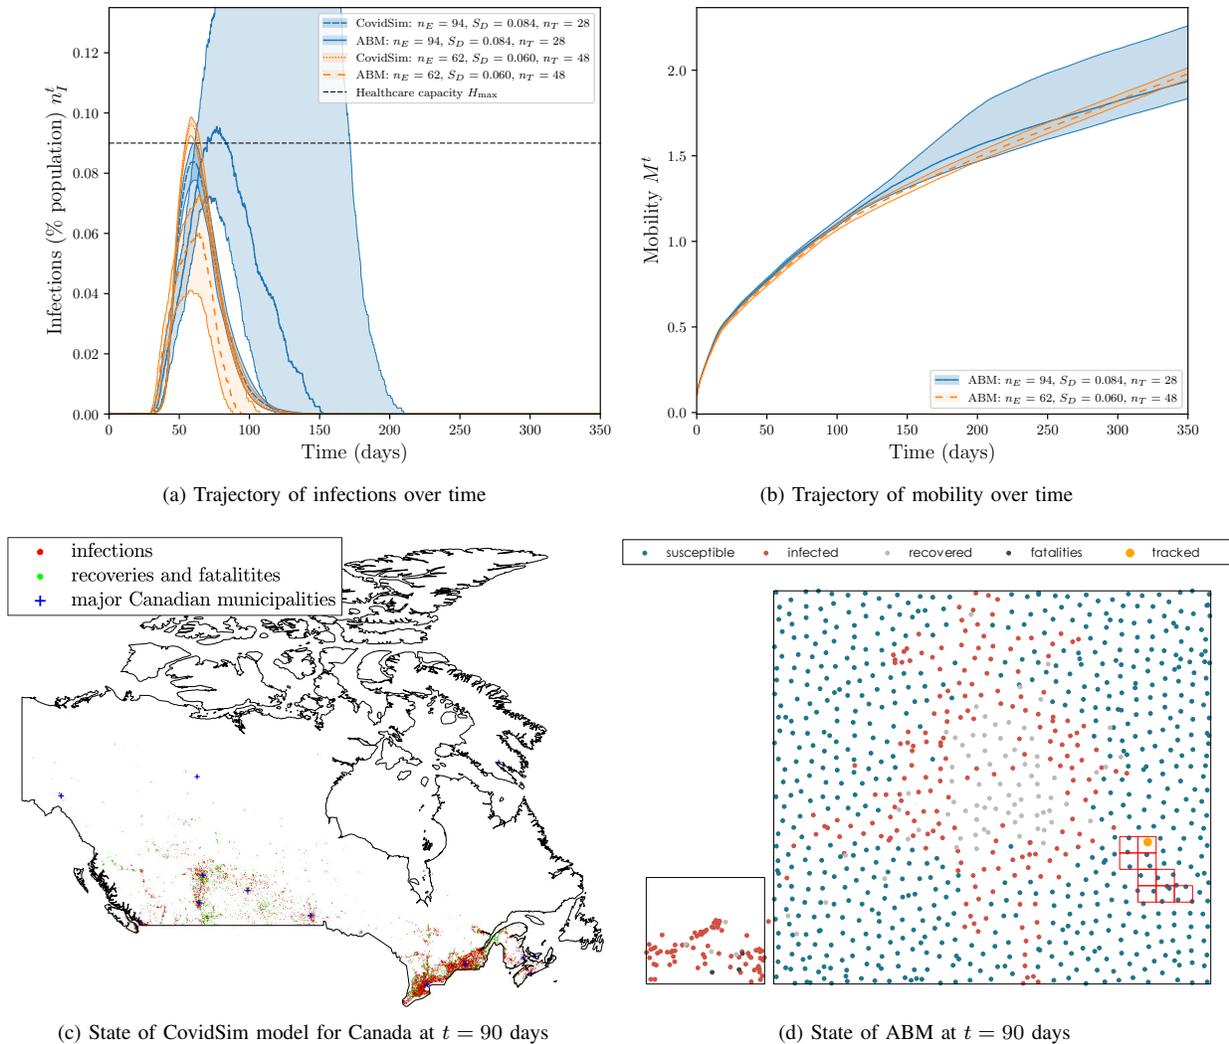


Fig. 5. Optimal results obtained from unified CovidSim and ABM optimization problem. The median trajectory of the epidemic for $N = 100$ realizations is shown in (a),(b) as predicted by either model at the optimizer with shaded areas representing the interquartile range. A snapshot of a single realization of both models is shown during the height of the epidemic for the intervention policy $n_E = 62$, $S_D = 0.06$, and $n_T = 48$ applied to (c) the CovidSim model and (d) the developed ABM with a single agent's mobility traced (shown by the red grids). See supplementary videos for an animation of these realizations.

during parametric studies [21].

We also assumed a linear mapping between the variables of our ABM and the CovidSim parameters in Table III. The actual function that maps our ABM parameters to those of CovidSim could be identified via supervised learning approaches such that the mean error between the predicted trajectories of both models is minimized [52].

Our model provides a conservative bound on the critical amount of intervention needed to control a local outbreak while CovidSim provides interventions that are more applicable at the national level. Our model is less computationally expensive than CovidSim and yields optimal policies that are comparable to those obtained using CovidSim.

VI. CONCLUSION

We present an abstract ABM for modeling human social systems and infectious diseases with several elements of novelty. They are listed as follows

- Adaptive behavior of agents modeled by the means of repulsive forces between agents. These repulsive forces are representative of intervention measures such as social distancing and change in magnitude depending on the distance between agents.
- Socio-economic cost of intervention measures on the population is modeled by means of a mobility metric which quantifies the relative freedom that agents possess.
- Efficient computation of social interaction among agents is accelerated using GPU programming.

In addition to the elements of novelty of the proposed ABM, a stochastic optimization algorithm (StoMADS-PB) with rigorous convergence properties based on martingale theory, evolutionary computation, and deterministic direct search are used to design intervention measures that maximize a mobility metric representative of socio-economic impact subject to several disease spread constraints such as healthcare capacity [1].

The obtained results in this paper provide lessons about infectious disease control. The abstract ABM developed in this work showed that the trajectory of a pandemic becomes highly unpredictable when the disease spirals out of control. The policy-making methods ensure that this does not happen by enforcing a stochastic constraint on the maximum number of infections. Stochastic constraint handling methods such as the ones used in this paper are very important for generating policies to control a novel pandemic such as COVID-19 and keep it in check. The most effective intervention measures provided by the optimization algorithms involve a large testing campaign and relatively moderate social distancing given that the compliance rate is 90% or more. This result qualitatively agrees with several observational studies in the literature [12], [14]. The trade-off between maximum number of infections and socio-economic impact can be investigated using a parametric study of different healthcare capacities.

Comparative studies with other optimization algorithms showed that only the StoMADS-PB and the augmented Lagrangian genetic algorithm discovered feasible solutions with respect to the healthcare capacity constraints. Furthermore, solutions of StoMADS-PB provided the best objective function values with fewer model evaluations when compared to other deterministic optimization algorithms. This study combined computational intelligence principles of probabilistic reasoning and evolutionary computation to guide the emergent behavior of an artificial life system despite its stochastic variability. The identification of optimal public health policies for a pandemic demonstrates the benefit of utilizing multiple computational intelligence principles to overcome their individual limitations (e.g., the assumption of deterministic fitness functions when using evolutionary algorithms).

The presented stochastic optimization approach does not require the use of our ABM. It can be coupled to any stochastic model, epidemiological or otherwise. A plethora of empirically grounded epidemiological models are available publicly and can be used to formulate policies. A comparative study with CovidSim shows that our model is complementary and useful for estimating socio-economic impact of different interventions. Different models can be combined in a similar fashion to enhance problem-solving formulations or conduct exploratory pilot studies [53].

The ABM can be extended to include more detailed social interaction networks, environments, and socio-economic impact metrics and continued development of the ABM is planned on a publicly available code repository (GitHub, https://github.com/khbalhandawi/COVID_SIM_GPU).

Finally, this work aims at demonstrating that the use of computational models and rigorous algorithms can contribute to determining policies that protect public health without suffocating socio-economic activity for ongoing and future pandemics.

ACKNOWLEDGMENTS

The authors are grateful for the partial financial support of their work through NSERC discovery grant 436193-2018. The authors would like to acknowledge the help and support of Dr.

Joseph Dzahini and Prof. Sébastien Le Digabel (Department of Applied Mathematics and Industrial Engineering, Ecole Polytechnique de Montreal) for their help and support with stochastic blackbox optimization. The authors would like to thank Dr. Dzahini for providing the MATLAB implementation of StoMADS-PB and coding the wrapper for the agent-based model. The first author would like to thank Martin Roa Villescás (Department of Electrical Engineering, Eindhoven University of Technology) for providing a minimum working example for interfacing scientific computing programs (e.g., agent-based models) with Qt user interfaces. The authors would like to acknowledge Marieh Al-Handawi (Division of Science and Mathematics, New York University), Saif Alhazaimeh (Worley Parsons), and Mohammad Sohayeb Akiel (Department of Civil Engineering and Applied Mechanics, McGill University) for providing some of the computational resources used to run the simulations in this paper.

REFERENCES

- [1] N. Hoertel, M. Blachier, C. Blanco, M. Olsson, M. Massetti, M. S. Rico, F. Limosin, and H. Leleu, "A stochastic agent-based model of the SARS-CoV-2 epidemic in France," *Nature Medicine*, vol. 26, no. 9, pp. 1417–1421, 2020.
- [2] S. Venkatramanan, B. Lewis, J. Chen, D. Higdon, A. Vullikanti, and M. Marathe, "Using data-driven agent-based models for forecasting emerging infectious diseases," *Epidemics*, vol. 22, pp. 43–49, 2018.
- [3] V. S. Tseng, J. Jia-Ching Ying, S. T. Wong, D. J. Cook, and J. Liu, "Computational intelligence techniques for combating COVID-19: A survey," pp. 10–22, nov 2020.
- [4] D. Q. Quach, D. P. Playne, and C. J. Scogings, "GPUAnimats—Simulating animats, an agent-based, artificial life model on graphical processing units," *Concurrency and Computation: Practice and Experience*, vol. 32, no. 24, p. e5968, dec 2020. [Online]. Available: <https://onlinelibrary.wiley.com/doi/full/10.1002/cpe.5968><https://onlinelibrary.wiley.com/doi/abs/10.1002/cpe.5968><https://onlinelibrary.wiley.com/doi/10.1002/cpe.5968>
- [5] T. McAttee and C. Szabo, "Complex systems and artificial life: A decade's overview," *The 2019 Conference on Artificial Life*, pp. 263–270, jul 2019. [Online]. Available: https://www.mitpressjournals.org/doi/abs/10.1162/fisal_a_00172
- [6] S. Abar, G. K. Theodoropoulos, P. Lemariniere, and G. M. O'Hare, "Agent based modelling and simulation tools: A review of the state-of-art software," pp. 13–33, may 2017.
- [7] W. O. Kermack and A. G. McKendrick, "A contribution to the mathematical theory of epidemics," in *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, vol. 115, no. 772. The Royal Society, 1927, pp. 700–721.
- [8] A. L. Bertozzi, E. Franco, G. Mohler, M. B. Short, and D. Sledge, "The challenges of modeling and forecasting the spread of COVID-19," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 117, no. 29, pp. 16732–16738, 2020.
- [9] P. Van den Driessche, "Reproduction numbers of infectious disease models," *Infectious Disease Modelling*, vol. 2, no. 3, pp. 288–303, 2017.
- [10] T. S. Brett and P. Rohani, "Transmission dynamics reveal the impracticality of COVID-19 herd immunity strategies," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 117, no. 41, pp. 25 897–25 903, 2020.
- [11] F. Balabdaoui and D. Mohr, "Age-stratified discrete compartment model of the COVID-19 epidemic with application to Switzerland," *Scientific Reports*, vol. 10, no. 1, 2020.
- [12] G. Pullano, L. Di Domenico, C. E. Sabbatini, E. Valdano, C. Turbelin, M. Debin, C. Guerrisi, C. Kengne-Kuetche, C. Souty, T. Hanslik, T. Blanchon, P. Y. Boëlle, J. Figoni, S. Vaux, C. Campese, S. Bernard-Stoecklin, and V. Colizza, "Underdetection of COVID-19 cases in France threatens epidemic control," *Nature*, vol. 590, no. 7844, pp. 134–139, 2020.
- [13] S. Y. Del Valle, J. M. Hyman, and N. Chitnis, "Mathematical models of contact patterns between age groups for predicting the spread of infectious diseases," *Mathematical Biosciences and Engineering*, vol. 10, no. 5-6, pp. 1475–1497, 2013.

- [14] A. Riccardi, J. Gemignani, F. Fernandez-Navarro, and A. Heffernan, "Optimisation of Non-Pharmaceutical Measures in COVID-19 Growth via Neural Networks," *IEEE Transactions on Emerging Topics in Computational Intelligence*, vol. 5, no. 1, pp. 79–91, feb 2021.
- [15] V. A. Karatayev, M. Anand, and C. T. Bauch, "Local lockdowns outperform global lockdown on the far side of the COVID-19 epidemic curve," *Proceedings of the National Academy of Sciences*, vol. 117, no. 39, p. 202014385, 2020.
- [16] C. A. O'Neil and L. Sattenspiel, "Agent-based modeling of the spread of the 1918-1919 flu in three Canadian fur trading communities," *American Journal of Human Biology*, vol. 22, no. 6, pp. 757–767, 2010.
- [17] J. Parker and J. M. Epstein, "A distributed platform for global-scale agent-based models of disease transmission," *ACM Transactions on Modeling and Computer Simulation*, vol. 22, no. 1, pp. 1–25, 2011.
- [18] L. Willem, S. Stijven, E. Tijssens, P. Beutels, N. Hens, and J. Broeckhove, "Optimizing agent-based transmission models for infectious diseases," *BMC Bioinformatics*, vol. 16, no. 1, p. 183, 2015.
- [19] W. Duan, X. Qiu, Z. Cao, X. Zheng, and K. Cui, "Heterogeneous and stochastic agent-based models for analyzing infectious diseases' super spreaders," *IEEE Intelligent Systems*, vol. 28, no. 4, pp. 18–25, 2013.
- [20] P. Kasaie and W. D. Kelton, "Simulation optimization for allocation of epidemic-control resources," *IIE Transactions on Healthcare Systems Engineering*, vol. 3, no. 2, pp. 78–93, 2013.
- [21] N. M. Ferguson, D. Cummings, S. Cauchemez, C. Fraser, S. Riley, A. Meeyai, S. Iamsirithaworn, and D. S. Burke, "Strategies for containing an emerging influenza pandemic in Southeast Asia," *Nature*, vol. 437, no. 7056, pp. 209–214, sep 2005. [Online]. Available: <https://www.nature.com/articles/nature04017>
- [22] B. Marshall and S. Galea, "Formalizing the role of agent-based modeling in causal inference and epidemiology," *American Journal of Epidemiology*, vol. 181, no. 2, pp. 92–99, 2015.
- [23] E. Bruch and J. Atwell, "Agent-Based Models in Empirical Social Research," *Sociological Methods and Research*, vol. 44, no. 2, pp. 186–221, 2015.
- [24] C. M. Macal, *Agent Based Modeling and Artificial Life*. New York, NY: Springer New York, 2009, pp. 112–131. [Online]. Available: https://doi.org/10.1007/978-0-387-30440-3_7
- [25] Z. D. Jian, H. J. Chang, T. sheng Hsu, and D. W. Wang, "Applying deep learning for surrogate construction of simulation systems," in *Advances in Intelligent Systems and Computing*, vol. 873. Springer Verlag, jul 2019, pp. 335–350. [Online]. Available: https://doi.org/10.1007/978-3-030-01470-4_18
- [26] G. Huang, G. B. Huang, S. Song, and K. You, "Trends in extreme learning machines: A review," *Neural Networks*, vol. 61, pp. 32–48, jan 2015. [Online]. Available: <http://dx.doi.org/10.1016/j.neunet.2014.10.001>
- [27] N. Ferguson et al., "Report 9: Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand," *Imperial College COVID-19 Response Team*, no. March, pp. 1–20, 2020. [Online]. Available: <https://www.imperial.ac.uk/media/imperial-college/medicine/sph/ide/gida-fellowships/Imperial-College-COVID19-NPI-modelling-16-03-2020.pdf>
- [28] P. Kasaie and W. D. Kelton, "Resource allocation for controlling epidemics: Calibrating, analyzing, and optimizing an agent-based simulation," *IIE Transactions on Healthcare Systems Engineering*, vol. 3, no. 2, pp. 94–109, 2013.
- [29] P. Sambaturu, B. Adhikari, B. A. Prakash, S. Venkatramanan, and A. Vullikanti, *Designing Effective and Practical Interventions to Contain Epidemics*. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, 2020, p. 1187–1195.
- [30] Van Gent, P., "An agent-based simulation of corona and other viruses in python," *GitHub*, [Accessed 2021-02-18]. [Online]. Available: https://github.com/paulvangentcom/python_corona_simulation
- [31] P. DSouza, "Epidemic Simulation," *GitHub*, [Accessed 2021-02-18]. [Online]. Available: <https://prajwalsouza.github.io/Experiments/Epidemic-Simulation.html>
- [32] G. Sanderson, "Simulating an epidemic," *Youtube*, [Accessed 2021-02-18]. [Online]. Available: <https://www.youtube.com/watch?v=gxAaO2rsdIs>
- [33] H. Stevens, "Why outbreaks like coronavirus spread exponentially, and how to 'flatten the curve'," *Washington Post*, pp. 1–10, 2020. [Online]. Available: <https://www.washingtonpost.com/graphics/2020/world/corona-simulator/>
- [34] J. H. Bergstrand, "The gravity equation in international trade: Some microeconomic foundations and empirical evidence," *The Review of Economics and Statistics*, vol. 67, no. 3, p. 474, 1985.
- [35] M. L. Brandeau, *Allocating Resources to Control Infectious Diseases*. Boston, MA: Springer US, 2004, pp. 443–464.
- [36] P. M. Davidson and S. L. Szanton, "Nursing homes and COVID-19: We can and should do better," *Journal of Clinical Nursing*, vol. 29, no. 15-16, pp. 2758–2759, 2020.
- [37] J. J. Cavallo, D. A. Donoho, and H. P. Forman, "Hospital capacity and operations in the coronavirus disease 2019 (COVID-19) pandemic — Planning for the Nth patient," *JAMA Health Forum*, vol. 1, no. 3, pp. e200345–e200345, 03 2020. [Online]. Available: <https://doi.org/10.1001/jamahealthforum.2020.0345>
- [38] N. M. Ferguson, D. Cummings, C. Fraser, J. C. Cajka, P. C. Cooley, and D. S. Burke, "Strategies for mitigating an influenza pandemic," *Nature*, vol. 442, no. 7101, pp. 448–452, jul 2006. [Online]. Available: <https://www.nature.com/articles/nature04795>
- [39] G. S. Zaric and M. L. Brandeau, "Resource allocation for epidemic control over short time horizons," *Mathematical Biosciences*, vol. 171, no. 1, pp. 33–58, 2001.
- [40] C. Audet, K. J. Dzahini, M. Kokkolaras, and S. Le Digabel, "Stochastic mesh adaptive direct search for blackbox optimization using probabilistic estimates," *Computational Optimization and Applications*, vol. 79, no. 1, pp. 1–34, may 2021. [Online]. Available: <https://doi.org/10.1007/s10589-020-00249-0>
- [41] K. J. Dzahini, M. Kokkolaras, and S. Le Digabel, "Constrained stochastic blackbox optimization using a progressive barrier and probabilistic estimates," 2020. [Online]. Available: <https://arxiv.org/pdf/2011.04225>
- [42] C. Audet, "A survey on direct search methods for blackbox optimization and their applications," in *Mathematics Without Boundaries: Surveys in Interdisciplinary Research*. Springer New York, 2014, pp. 31–56.
- [43] R. Durrett, *Probability: Theory and Examples*, 5th ed. Cambridge University Press, 2019.
- [44] C. Audet and J. E. Dennis, "A progressive barrier for derivative-free nonlinear programming," *SIAM Journal on Optimization*, vol. 20, no. 1, pp. 445–472, 2009. [Online]. Available: <http://www.siam.org/journals/ojsa.phphttp://epubs.siam.org/doi/10.1137/070692662>
- [45] C. Audet, A. R. Conn, S. Le Digabel, and M. Peyrega, "A progressive barrier derivative-free trust-region algorithm for constrained optimization," *Computational Optimization and Applications*, vol. 71, no. 2, pp. 307–329, 2018.
- [46] S. Le Digabel, "Algorithm 909," *ACM transactions on mathematical software*, vol. 37, no. 4, pp. 1–15, 2011.
- [47] R. M. Lewis and V. Torczon, "A globally convergent augmented Lagrangian pattern search algorithm for optimization with general constraints and simple bounds," *SIAM Journal on Optimization*, vol. 12, no. 4, pp. 1075–1089, 2002. [Online]. Available: <https://epubs.siam.org/doi/abs/10.1137/S1052623498339727>
- [48] A. Conn, N. Gould, and P. Toint, "Globally convergent augmented Lagrangian barrier algorithm for optimization with general inequality constraints and simple bounds," *Mathematics of Computation*, vol. 217, no. 97, pp. 261–288, 1997. [Online]. Available: <https://www.ams.org/journal-terms-of-use>
- [49] A. R. Conn, N. Gould, and P. L. Toint, "Globally convergent augmented Lagrangian algorithm for optimization with general constraints and simple bounds," *SIAM Journal on Numerical Analysis*, vol. 28, no. 2, pp. 545–572, jul 1991.
- [50] K. Deb and S. Srivastava, "A genetic algorithm based augmented Lagrangian method for constrained optimization," *Computational Optimization and Applications*, vol. 53, no. 3, pp. 869–902, 2012.
- [51] M. Georgioudakis and V. Plevris, "A Comparative Study of Differential Evolution Variants in Constrained Structural Optimization," *Frontiers in Built Environment*, vol. 6, p. 102, jul 2020. [Online]. Available: <https://www.frontiersin.org/articles/10.3389/fbuil.2020.00102/full>
- [52] G. Luo, "A review of automatic selection methods for machine learning algorithms and hyper-parameter values," *Network Modeling Analysis in Health Informatics and Bioinformatics*, vol. 5, no. 1, pp. 1–16, dec 2016. [Online]. Available: <https://link.springer.com/article/10.1007/s13721-016-0125-6>
- [53] W. Edeling, H. Arabnejad, R. Sinclair, D. Suleimenova, K. Gopalakrishnan, B. Bosak, D. Groen, I. Mahmood, D. Crommelin, and P. V. Coveney, "The impact of uncertainty on predictions of the CovidSim epidemiological code," *Nature Computational Science*, vol. 1, no. 2, pp. 128–135, feb 2021. [Online]. Available: <https://doi.org/10.1038/s43588-021-00028-9>



Khalil Al Handawi Khalil Al Handawi is a post-doctoral researcher at the Systems Optimization Lab in the Department of Mechanical Engineering at McGill University. His research focuses on the optimization of complex systems and uncertainty modeling with an emphasis on healthcare applications. His research resulted in public health policy-making tools for the COVID-19 pandemic based on agent-based and machine learning models. He received his PhD in mechanical engineering from McGill University. His dissertation focused on the optimization of

aerospace design for flexibility and robustness when presented with uncertain requirements. His research resulted in a technology transfer at GKN Aerospace Engine Systems, Trollhättan, Sweden. He received his bachelors and masters degrees in mechanical engineering from Khalifa University with an emphasis on instrumentation, photonics, asset integrity management, and corrosion detection in oil and gas structures. He is an ASME and design society member and is the recipient of the McGill Engineering Doctoral Award and the Fonds de recherche du Québec Nature et Technologies (FRQNT) doctoral fellowship.



Michael Kokkolaras Dr. Kokkolaras is the founding director of the Systems Optimization Laboratory in the Department of Mechanical Engineering at McGill University. He joined McGill in 2012 after spending 12 years at the University of Michigan (UM) in Ann Arbor, where he held research faculty appointments at the Department of Mechanical Engineering (primary) and the UM Transportation Research Institute (joint/courtesy); he was the recipient of the 2008 UM College of Engineering Outstanding Research Scientist Award. Dr. Kokkolaras is also

Associate Director of the McGill Institute for Aerospace Engineering (MIAE) and a full member of the Fonds de recherche du Québec Nature et Technologies (FRQNT)-funded multi-university Group for Research in Decision Analysis (GERAD). Since January 2016, he is a Guest Researcher at the Department of Industrial and Materials Science (Division of Product Development) of Chalmers University of Technology in Gothenburg, Sweden. He also spent the 2018-19 academic year as a Visiting Professor at the Department of Industrial Engineering (LGI) of CentraleSupélec / Université Paris-Saclay in France. Dr. Kokkolaras' research interests include multidisciplinary optimization, simulation-based engineering design, uncertainty quantification, modeling and validation, systems of systems, product families and optimization applications in transportation, energy, and healthcare engineering systems. He is an ASME Fellow, an Associate Fellow of the AIAA (serving on its Multidisciplinary Design Optimization Technical Committee since 2006), and elected member of the Advisory Board of the Design Society. Dr. Kokkolaras has served as Associate Editor of the ASME Journal of Mechanical Design from 2008 to 2014, the Transactions of the Canadian Society for Mechanical Engineering from 2013 to 2018, and Optimization and Engineering (Springer) from 2013 to 2018). He is currently serving as Associate Editor of Structural and Multidisciplinary Optimization (Springer) and Aerospace (MDPI). He has also served as Chair of the Design Automation Executive Committee of the ASME Design Engineering Division, Program and Conference Chair of the ASME Design Automation Conference, and Program Co-Chair of the 2017 International Conference on Engineering Design.