

# Differentiable Agent-based Epidemiology

Ayush Chopra\*  
Massachusetts Institute of Technology  
USA

Alexander Rodríguez\*  
Georgia Institute of Technology  
USA

Jayakumar Subramanian  
Adobe MDSR  
India

Arnau Quera-Bofarull  
University of Oxford  
United Kingdom

Balaji Krishnamurthy  
Adobe MDSR  
India

B. Aditya Prakash  
Georgia Institute of Technology  
USA

Ramesh Raskar  
Massachusetts Institute of Technology  
USA

## ABSTRACT

Mechanistic simulators are an indispensable tool for epidemiology to explore the behavior of complex, dynamic infections under varying conditions and navigate uncertain environments. Agent-based models (ABMs) are an increasingly popular simulation paradigm that can represent the heterogeneity of contact interactions with granular detail and agency of individual behavior. However, conventional ABM frameworks not differentiable and present challenges in scalability; due to which it is non-trivial to connect them to auxiliary data sources. In this paper, we introduce GRADABM: a scalable, differentiable design for agent-based modeling that is amenable to gradient-based learning with automatic differentiation. GRADABM can quickly simulate million-size populations in few seconds on commodity hardware, integrate with deep neural networks and ingest heterogeneous data sources. This provides an array of practical benefits for calibration, forecasting, and evaluating policy interventions. We demonstrate the efficacy of GRADABM via extensive experiments with real COVID-19 and influenza datasets.

## CCS CONCEPTS

• **Computing methodologies** → **Machine learning; Agent / discrete models**; • **Applied computing** → **Epidemiology**.

## KEYWORDS

Differentiable Agent-based Modeling, Computational Epidemiology, Automatic Differentiation, Deep Neural Networks

### ACM Reference Format:

Ayush Chopra, Alexander Rodríguez, Jayakumar Subramanian, Arnau Quera-Bofarull, Balaji Krishnamurthy, B. Aditya Prakash, and Ramesh Raskar. 2023. Differentiable Agent-based Epidemiology. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, London, United Kingdom, May 29 – June 2, 2023, IFAAMAS, 13 pages.

## 1 INTRODUCTION

Agent-based models (ABMs) are discrete simulators which comprise a collection of agents which can act and interact within a computational world [19, 37, 50, 61, 72]. They can explicitly represent

\*Both authors contributed equally to this research.

*Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. © 2023 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

heterogeneity of the interacting population via underlying contact networks and model the adaptability of individual behavior to enable more realistic simulations. As a result, ABMs are increasingly popular in epidemiology in place of classical ODE-based models, such as the SIR (susceptible-infected-recovered) model [28, 33]. In recent efforts to contain the spread of COVID-19, ABMs have been used to evaluate the benefit of delaying 2nd dose of the vaccine [61], deploying mobile apps for digital contact tracing [6], and prioritizing test speed over specificity [45]. The utility of such simulators for practical decision making depends upon their ability to recreate the population with great detail, integrate with real-world data streams, and analyze the sensitivity of results.

However, ABMs are conventionally slow, difficult to scale to large population sizes [52] and tough to calibrate with real-world data [30]. Conventional ABM frameworks follow an object-oriented (agent-centered) design where the agents are modeled as objects. While conceptually appealing, object-oriented approaches often represent inter-agent infection transmission (over contact networks) and within-agent disease progression inefficiently, resulting in poor scalability as the agent population and number of interactions grows. Performing a single forward simulation over a large ABM can take several days [12, 18], making calibration results subpar where approaches require multiple forward passes [22, 61, 68]. This is an issue since simulation results (emergent behavior) can be highly sensitive to the scale of the input population and parameters. In addition, incorporating novel sources of data that could inform calibration and other downstream tasks (e.g., forecasting) is often laborious and adds overhead complexity to the ABM (e.g., incorporating digital exposure data to ABMs [34]).

In this paper, we introduce GRADABM to alleviate these concerns and enhance the potential of ABMs for practical decision making in epidemiology (Fig. 1). Our key insight is to make simulations tensorized and differentiable which enables gradient-based learning with automatic differentiation. First, in contrast to conventional ABMs, GRADABM follows a tensorized design [25] where agent states are represented as vectors and their interaction networks as (sparse) adjacency matrices. As a result, forward simulation of agent interaction and disease progression can be run rapidly, in a highly parallelized manner (millions-size networked populations in a few seconds on commodity hardware). Second, ABMs are stochastic simulators which routinely require sampling from discrete distributions for execution (e.g., all interactions with an infected agent may

Simulator	Input (data)		Process (modeling)			Output (utility)	
	Granularity	Heterogeneity	Fast	Scalable	Differentiable	What-if Interventions	Forecasting
Compartmental models <small>[35, 27, 44, 10, 47, 9]</small>	medium	low	yes	no	yes	medium	yes
ABM (conventional) <small>[22, 37, 42, 53]</small>	high	medium	no	yes	no	high	no
Differentiable ABM (GradABM)	high	high	yes	yes	yes	high	yes

**Figure 1:** We introduce GRADABM, an end-to-end differentiable ABM which can: i) ingest multi-granular (citizen, block, county, state) and heterogeneous (clinical, behavioral, policy, census, survey) data sources; ii) simulate realistic populations *quickly and reliably* with gradient-based calibration; iii) facilitate flexible interventions for policy decisions and efficient forecasting with end-to-end DNN-GRADABM pipelines. The key idea of GRADABM is a general-purpose differentiable sparse tensor-calculus design which we validate by demonstrating utility for epidemiology here.

not result in a new infection; COVID-19 tests may return false positives/negatives with some probability). Ensuring differentiability requires gradient estimation through such discrete stochasticity. To achieve this, GRADABM reparameterizes discrete distributions (e.g., Bernoulli) with continuous relaxations using the gumbel-softmax gradient estimator [39, 41]. This allows GRADABM to benefit from gradient-based optimization and tools for automatic differentiation.

Using GRADABM’s scalable and differentiable design, we merge it with deep neural networks (DNNs) to help with calibration, forecasting and policy evaluation. In particular, we demonstrate how our method can be coupled with a calibration DNN allowing seamless integration of novel datasets from heterogeneous sources—several of which were difficult to incorporate in the calibration of previous ABM designs. To illustrate the effectiveness of our proposed framework, we run several experiments using real world datasets. In particular, we treat COVID-19 and Influenza as case studies, highlighting GRADABM’s benefits with respect to robust, time-space personalized forecasting, and evaluation of policy interventions.

**Our Contributions:** 1) Present GRADABM, which is a scalable and differentiable ABM design amenable for gradient-based learning with auto-diff. 2) Demonstrate that GRADABM allows quickly simulating realistic population sizes, integrating with deep neural networks and ingesting heterogeneous data sources which provides an array of practical benefits for calibration, forecasting, and policy evaluation. 3) Validate effectiveness of GRADABM through experiments using real-world datasets of COVID-19 and Influenza.

## 2 BACKGROUND AND RELATED WORK

Before describing GRADABM, we first review relevant works in agent-based models (ABMs) and with relevance to epidemiology. Then, we present our work in context of broad research in differentiable simulators and also briefly mention an insight we find useful when designing GRADABM (permutation invariance). We refer the reader to [28, 33, 46, 59] for a more extensive discussion.

### 2.1 Agent-based Modeling

ABMs represent systems [19, 37] as collections of agents which can act and interact with each other within a computational world.

Owing to their flexibility, practitioners have applied ABMs in a wide variety of problem domains such as modeling cells in a tumor micro-environment to diagnose cancers [50], humans in a physical environment to study economic policies [72], infectious diseases [6, 61] and avatars in a digital world to study misinformation [20].

**ABMs in Epidemiology.** In the context of epidemiology, ABMs are used to understand how disease spreads and evaluate efficacy of health interventions [48]. They simulate transmission of pathogens within multiple contact networks with an epidemiological model [12, 34] used to describe between-host transmission and within-host progression of infection. [29, 52] survey some methods and assumptions for designing such simulators. Our work introduces a tensorized and differentiable epidemiological simulator that can generalize across multiple infections and model specifications.

**Scalability of ABMs.** It is a key consideration since modeling granular population details is computationally expensive. Conventional frameworks like Mesa [53] follow an object-oriented design which while ease to use (Python-API) is prohibitively slow to scale (require few hours for a single iteration). There have been attempts to reduce this burden and simulate realistic scale with distributed HPC systems and GPU-optimized implementations. EpiFast [18] demonstrated simulating large contact networks in a few minutes on distributed systems. However, the required compute for such executions is expensive and not widely available. OpenABM [34] presents an optimized C++/Cuda API to build an epidemiological ABM for fast execution on commodity hardware but it is non-trivial for epidemiological practitioners to extend, whilst not being amenable to gradient-based learning. Motivated by recent work in molecular dynamics [62], GRADABM is implemented using highly optimized sparse tensor APIs of auto-diff packages (e.g. PyTorch/JAX) to accelerate and scale simulations while preserving ease of use.

**Calibration of ABMs.** This involves identifying appropriate values for latent parameters and is essential to ensure reliability of results. Generally, a hybrid strategy is used where some parameters are sourced from control trials conducted offline by clinical experts and then others are calibrated in-silico by achieving a goodness-of-fit between emergent ABM output and real-world macro data

(e.g. number of deaths). The conventional technique is to use grid search [22, 61], beam search [68] by running several forward simulations and finding the best-fit. This is often slow and tough to scale to many parameters. Recent works have proposed new optimization methods for faster and more accurate calibration. For example, [57] uses an ML model to identify what parameters to calibrate, however, values are identified with random search. [70] use a gaussian process emulator which is computationally more efficient than the ABM. In a similar vein, [6] propose a surrogate strategy by first calibrating a compartmental SEIR-based model using standard ODE solvers, and supplying the resulting parameters to an ABM. However, such surrogate methods increase model miss-specification error as they cannot simulate stochasticity inherent in the ABM. Further, the SEIR-like models have few tunable parameters in comparison to ABMs, limiting the degree of calibration which can be achieved. In contrast, the scalable and differentiable design of GRADABM allows efficiently calibrating large-scale ABMs with gradient-based learning (and using deep neural networks).

**DNNs with ABMs.** Integrating deep neural networks (DNNs) with ABMs is an active direction. One popular approach has focused on multi-agent reinforcement learning (MARL), wherein the ABM is treated as a reinforcement learning environment, and a DNN is used to learn policies for the agents. Such perspective has been employed by [72] to learn equitable economic policies, [64] to analyze societal segregation dynamics, and [55] to learn oil and gas macro strategies. An alternate strategy has focused on using DNNs to emulate agent-based models [9, 15, 57] by using the ABM to generate training dataset for DNNs. Recently, [15] used an epidemiological ABM to generate synthetic datasets and trained DNN models to predict infectiousness for contact tracing. [57] used a similar strategy (of DNN as surrogate) to identify what simulation parameters to calibrate. Both these directions (MARL, Surrogate) are constrained by non-differentiability of the ABM and hence cannot utilize automatic differentiation. In contrast, we introduce differentiable ABMs (GRADABM) which allows joint training of DNNs with the ABM. This makes it possible to integrate heterogeneous data to infer latent micro variables and calibrate simulation parameters through hybrid DNN-ABM pipelines.

## 2.2 Automatic Differentiation for Simulations

Differentiable simulation is a powerful family of techniques that applies gradient-based methods to learning and control of physical systems. With progress in automatic differentiation, as computing gradients become easier, such simulators are emerging across wide range of systems. In combination (and joint training) with deep neural networks, such simulators are being used from computer graphics for looking through scattering media [23] and reconstructing human body pose [14]; to molecular dynamics for analysing the force field of ionic liquids [49, 62]; to robotics for grasping and locomotion [32, 38]. Even in epidemiology, gradient-based optimization with epidemiological simulators has also been recently explored with SIR-like models to learn simulation parameters [10, 54] and for making neural models to learn epidemic dynamics from mechanistic models [58]. GRADABM aims to realize differentiable simulations for stochastic agent-based models. In contrast to relevant recent work

in [8, 11], GRADABM obtains gradient estimates through stochasticity by efficiently reparameterizing with the gumbel-softmax gradient estimator (unlike smoothing in [8] which has exponential cost scaling), shows scalability to million-size populations (unlike 25 agents in [11]) and also demonstrates ability to integrate with deep neural networks to calibrate large epidemiological ABMs.

## 2.3 Invariances, Computation and ABMs

As a result of leveraging the structure of the physical world, DNN architectures are effective in overcoming the curse of dimensionality [21]. This is key to efficient computation on grids with CNNs (translation invariance) and graphs with GNNs via neural message passing (permutation invariance). We posit that while useful for learning DNNs, utilizing these invariances can make computation tractable in physical systems, where they exist naturally. Specifically, we observe that epidemiological models [6, 25, 34, 61] also adhere to permutation invariance wherein the order of infectious interactions (pair-wise message passing) within a step does not matter when estimating the probability of infection from those interactions. Motivated by this observation, we implement the transmission model of GRADABM as a differentiable message passing operation defined with physical equations instead of DNNs

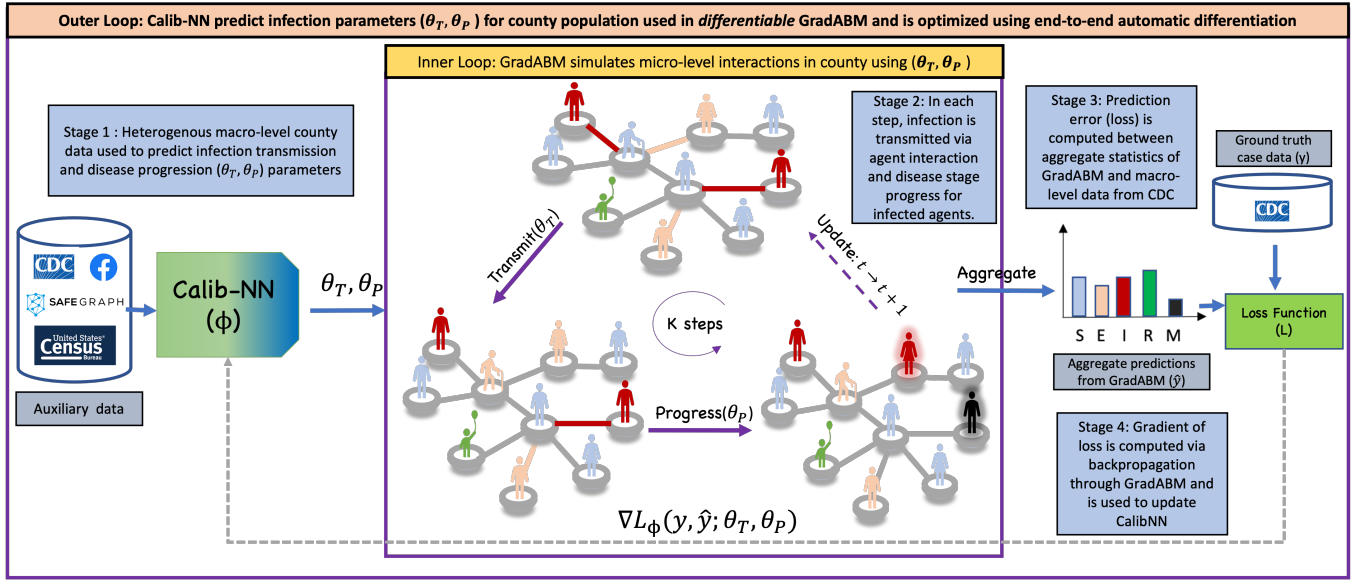
## 3 DIFFERENTIABLE AGENT-BASED MODELING FOR EPIDEMIOLOGY

In what follows, we present a framework for epidemiological modeling with differentiable ABMs. The pipeline is summarized in Fig. 2, where the inner loop is the differentiable epidemiological simulator (GRADABM) and the outer loop is the gradient-based calibration procedure using the neural network CALIBNN. The epidemiological simulator is described in Sec. 3.1, the calibration procedure in Sec. 3.2 and the use of the final model for forecasting and evaluating policy interventions in Sec. 3.3.

### 3.1 GRADABM: Epidemiological Simulator

Here, we describe the inner loop shown in Fig. 2. We consider a  $K$ -step discrete-event simulation with a population of  $n$  interaction agents. The key contribution of this work is GRADABM: an agent-based modeling design amenable for gradient-based learning with automatic differentiation. To achieve this, GRADABM follows a tensorized implementation where each agent state is represented as a vector, the interaction networks as adjacency matrices and all discrete distributions (e.g., Bernoulli) are reparameterized with continuous approximations (Gumbel-Softmax [41]), for end-to-end differentiability. As we discuss later, this design allows scaling agent-based models to realistic population sizes and integrating them with deep neural networks, which provides an array of practical benefits. Before proceeding, we first describe the epidemiological model used in our implementation which follows from [25, 34, 61]

The state of each agent  $j \in \{1, \dots, n\}$ , at time step  $t \in \{0, \dots, K\}$ , is denoted by a 3-D vector  $X_j^t = \{a_j, d_j^t, e_j^t\}$  where  $a_j \in \{0 - 10, 11 - 20, 21 - 30, \dots, 71 - 80, 80+\}$  is the age;  $d_j^t$  is the current disease stage; and  $e_j^t \in \{-1, \dots, t - 1\}$  denotes the time step of last exposure. For example, in COVID-19,  $d_j^t$  can take values in  $\{S, E, I, R, M\}$  which denote susceptible (S), exposed (E), infected (I), recovered (R)



**Figure 2: Differentiable Agent-based Epidemiology involves 4 stages** i) heterogeneous macro-level population data (CDC, census, behavioral, survey) is input to a calibration model (CALIBNN) to predict epidemiological parameters  $(\theta_T, \theta_P)$ . ii)  $(\theta_T, \theta_P)$  are used run  $K$  forward steps of the fully-differentiable epidemiological model (GRADABM) which simulates micro-level infection transmission (Transmit) and disease progression (Progress) over individual contact networks. Disease statistics are aggregated (Aggregate) at end of  $K$  steps to obtain the macro-level simulation output  $(\hat{y})$ . iii) Error between predicted  $\hat{y}$  and real-world case statistics  $(y)$  is used to define a loss  $(L(\hat{y}; y))$ . iv) Gradient of this loss is computed by automatic differentiation through the micro-level GRADABM to update weights of macro-level GRADABM using gradient descent.

and mortality (M) stages. Agent state may update due to change in disease stage resulting from new infectious interaction with an exposed or infected agent (captured by the transmission function) or the natural progression of a previously incubated infection (captured by the progression function). This simulator is parameterized with time-dependent parameters which govern the transmission function  $(\theta_T^t = [R^t, S_a, T_s, i^0])$  and progression function  $(\theta_P^t = [\tau_{EI}, \tau_{IR}, \tau_{IM}, m])$ . Here  $R^t$  is the time-dependent disease reproduction number (measure of infection rate),  $S_a$  is the age-stratified susceptibility,  $T_s$  is the disease stage-stratified infection transmissibility,  $i^0$  is the percentage of infections at time  $t = 0$ ,  $\tau_{EI}$ ,  $\tau_{IR}$ , and  $\tau_{IM}$  are the pairwise stage transition times from E to I, I to R and I to M, respectively, and  $m$  is the mortality rate (fraction of expired agents amongst all agents in the  $\{R, M\}$  disease stage). At each step  $t$ , every agent  $i$  interacts with a set of neighbors  $\{j : j \in \mathcal{N}_i\}$  from which she/he may accumulate or transmit infection. This disease stage evolution, at any step  $t$ , is given by:

$$d_i^{t+1} = \text{Update}(X_i^t, \mathcal{N}_i, (X_j^t)_{j \in \mathcal{N}_i}, \theta_T^t, \theta_P^t), \quad (1)$$

where  $\text{Update}(X_i^t, \mathcal{N}_i, (X_j^t)_{j \in \mathcal{N}_i}, \theta_T^t, \theta_P^t)$  is

$$= \begin{cases} \text{Transmit}(X_i^t, \mathcal{N}_i, (X_j^t)_{j \in \mathcal{N}_i}, \theta_T^t), & \text{if } d_i^t = S, \\ \text{Progress}(X_i^t, \theta_P^t), & \text{if } d_i^t \in \{E, I\}. \end{cases} \quad (2)$$

The transmission function  $\text{Transmit}(X_i^t, \mathcal{N}_i, (X_j^t)_{j \in \mathcal{N}_i}, \theta_T)$  computes the probability of infection transmission as a result of interactions between susceptible and infected (or exposed) agents. This probability is proportional to an aggregate rate of transmission

accumulated over the multiple interactions and is defined as

$$q(t) = 1 - e^{-\lambda_A} \quad (3)$$

where  $\lambda_A = \bigcup_{j \in \mathcal{N}_i} (\lambda(R^t, S_i, T_j, e_j^t))$ . Here,  $\lambda$  denotes the rate of transmission from a single interaction and  $\bigcup$  is an aggregation function which accumulates transmission over multiple interactions. In our implementation,  $\lambda$  is a linear function as defined in [6] and  $\bigcup$  is the summation ( $\sum$ ) function following the invariance of infection transmission noted in Sec. 2 (more details in appendix). Before proceeding, the probability  $q(t)$  is used to sample the agent's new disease stage from a Bernoulli random variable as  $\hat{Q} = \text{Bernoulli}(q(t), 1 - q(t))$ . Once infected, the agent can now infect other agents on subsequent steps and also enters a hierarchy of disease stage progression which triggers subsequent changes in the state of the agent. The progression function  $\text{Progress}(X_i^t, \theta_P)$  updates the disease stage from  $E \rightarrow I$  or from  $I \rightarrow \{R, M\}$  at times determined by the stage transition time parameters, which may also be stochastic and require sampling from discrete distributions. This sequence of transmission and progression is repeated for  $K$  steps of the simulation. Finally, the aggregate cumulative deaths is determined as:  $\hat{y} = \text{Aggregate}(d_i^T) = m \times (d_i^T \in \{R, M\})_{i \in \{1, \dots, n\}}$ , where  $m \in \theta_P$ . Ensuring differentiability of GRADABM requires gradient estimation of discrete stochasticity. For this we replace the non-differentiable sample from the Bernoulli distribution with the Gumbel-Softmax gradient estimator [39, 42].

### 3.2 Training and Calibration of GRADABM

In order to use GRADABM for real-world applications, we need to be able to calibrate its parameters  $(\theta_T^t, \theta_P^t)$  in such a way that the aggregate predicted quantities from the model such as cumulative deaths match the observed values from real-world data. Here, the differentiability of GRADABM enables the possibility of utilizing gradient-based learning to calibrate the simulator by integrating with deep neural networks. Further, we note that the tensorized specification of GRADABM is also critical since fast forward simulations are required to this iterative optimization where the simulator is executed at each optimization step. The calibration protocol is visualized in the outer loop in Fig. 2 and described below. This outer training and calibration loop runs for  $W$  optimization steps, with each step  $w \in \{1, \dots, W\}$  comprising of the following four stages:

*Stage 1.* : At the beginning of each outer loop, candidate parameters  $(\theta_T^t, \theta_P^t)_{t \in \{0, \dots, K\}}^w$  are generated by a calibration neural network (CALIBNN), whose weights at time step  $w$  we denote by  $\phi^w$ . To aid the generation of parameters, the calibration network is passed heterogeneous data as input. This may include data sources that were not considered during the initial construction of the underlying ABM being calibrated. For example, the calibration neural network may be passed social network data (e.g., self-reported symptom surveys conducted by Facebook) that was not integrated into the ABM, but that may still be of value when predicting disease prevalence. Intuitively, one may view this calibration network as a domain expert who integrates new data that the ABM designer was previously unaware of.

*Stage 2.* : After parameters have been generated, they are passed to GRADABM in order to perform a  $K$  step simulation, as described previously as the inner loop in Fig. 2. This outputs the aggregate predictions from the simulator which includes cumulative infections, deaths etc.

*Stage 3.* : The cumulative deaths across the  $K$  steps of simulation are then compared with corresponding ground truth values. More specifically, a mean-squared loss between the simulated and actual cumulative deaths is computed. This is defined as:  $\mathcal{L}(\hat{y}^w, y; (\theta_T^t, \theta_P^t)^w) = \text{MSE}(\hat{y}^w, y)$ , where  $y$  is the ground truth data representing cumulative deaths,  $\hat{y}^w$  is the cumulative deaths predicted by GRADABM at the  $w^{\text{th}}$  training step with input parameters  $(\theta_T^t, \theta_P^t)^w$  and MSE denotes the mean-squared error function.

*Stage 4.* : Since GRADABM is differentiable (as established earlier), we may calculate the gradient of the computed loss with respect to each parameter of CALIBNN, via backpropagation. Moreover, due to GRADABM's network-centric tensorized design, auto-diff packages from libraries such as PyTorch can be leveraged to compute these gradients efficiently in a highly parallelized manner. The final step of this gradient computation involves back-propagation through the weights of the CALIBNN neural network ( $\phi$ ) as well.

Once, the gradient of the loss is computed, we use the classical gradient descent algorithm to update the weights of the calibration neural network at the end of each iteration as follows:

$$\phi^{w+1} = \phi^w - \alpha \frac{\partial \mathcal{L}(\hat{y}^w, y; (\theta_T^t, \theta_P^t)^w)}{\partial \phi}, \quad (4)$$

where  $\alpha$  is the learning rate,  $(\theta_T^t, \theta_P^t)^w = f(D; \phi^w)$  and  $\hat{y}^w$  is computed by calling Update  $K$  times with same parameters  $(\theta_T^t, \theta_P^t)^w$ . In essence, calibration here involves optimizing weights of CALIBNN ( $\phi$ ) which predicts the simulation parameters  $(\theta_T, \theta_P)$  instead of the parameters directly. We are essentially training the calibration neural network to understand heterogeneous data as a means of producing good parameterizations for GRADABM. We call this calibration pipeline involving CALIBNN as deep calibrated GRADABM (DC-GRADABM).

Consider for instance the scenario with multiple simulators in an experiment. This is realistic since a distinct GRADABM may be defined for each county in a state. We now have the opportunity to jointly calibrate parameters for all counties using a shared CALIBNN. Given  $T$  counties, the weights of CALIBNN  $\phi$  may be trained on the average of the loss for all counties (Stage 2). The same weights of CALIBNN can be used to predict personalized parameters for each county-specific GRADABM. We call such a multi-task calibration procedure as joint deep calibrated GRADABM (JDC-GRADABM).

Note that each step of gradient descent requires a full forward simulation of GRADABM. As a result, GRADABM's network-centric design and tensorized implementation are key in ensuring that a sufficient number of gradient steps can be quickly computed to effectively train the calibration network. Additionally, observe that this optimization strategy is not possible with classical object-oriented ABMs, as they do not allow for scalable gradient computation.

*REMARK 1.* We note that the CALIBNN based approach presented in this paper is fundamentally distinct from emulation or surrogate models. Emulation models take the same input as the ABM's input and predict the output of the ABM, without actually simulating any agent behavior. In contrast, the output of CALIBNN serves as the input to GRADABM, which then simulates the behavior of the agents. Thus, CALIBNN extends simulation pipeline, by enabling us to learn the correct inputs for GRADABM.

### 3.3 Forecasting and policy evaluation

Once trained the hybrid framework we propose, consisting of CALIBNN and GRADABM, it is ready to make forecasts in new, previously unseen scenarios as well as help decide on intervention policies. These unseen scenarios may include predicting future the evolution of multiple infections for a county/region, or it may involve making predictions even when the data (inputs to CALIBNN) for a county/region is noisy. In each of these scenarios, the data available for the county/region is fed to CALIBNN, which then outputs personalized time-varying parameters  $\theta_T^t$  and  $\theta_P^t$  for  $t \in \{0, \dots, K + H\}$ , where  $H$  is the forecasting horizon. These parameters are used in GRADABM, which is run for  $K+H$  steps, where the last  $H$  simulation steps will serve as our forecasts.

## 4 EXPERIMENTAL SETUP

To illustrate the effectiveness of our framework, we conduct experiments for COVID-19 and influenza on multiple counties of the state of Massachusetts, USA, learning personalized parameters for each county. Before presenting our analyses, we provide details regarding our experimental setup.

## 4.1 Heterogeneous Data Sources

We describe the data sources passed as input to CalibNN, used to construct the population for GRADABM and parameterize disease transmission and progression. (i) To start, we outline the features used as input to CalibNN during our experiments. In the case of COVID-19, CalibNN receives 5 input signals<sup>1</sup> including insurance claims data, online symptoms surveys from Facebook, and line-list data. Since all of these datasets contain granular information regarding each specific county, the data points input to CalibNN varies based on the county being modeled. Meanwhile, for influenza, CalibNN receives 14 signals originating from the Google symptoms dataset [2]. This data is reported at the state level and hence we use the same data for all counties. (ii) In order to construct GradABM, several datasets are used. Demographic information from the US Census [3] is used to generate a synthetic agent population. Agent interaction graphs are created using census block level mobility data from Safegraph and leveraging demographic information (such as age) inside each census block. (iii) With respect to parameters in the transmission and progression models, we utilize clinical data from reliable sources such as the CDC [4] and clinical papers [31, 34]. The target variable for COVID-19 is the COVID-associated mortality, whilst for influenza we use influenza-like-illness (ILI) counts collected by the CDC. Ground truth data for both target variables is obtained from the JHU and the CDC.

## 4.2 Baselines

To showcase the performance of our calibration procedure (the outer loop described in Fig. 2), we compare GRADABM with numerous benchmarks inspired by popular approaches for calibrating an ABM. **ExpertSearch-ABM:** Following [34, 61], it uses parameters set by a combination of expert advice and randomized search. For COVID-19, we get  $R_0$  and case-fatality rate from [5, 17] and  $R_0$  for flu from [26]. **SurrogateODE-ABM:** Following [6], it calibrates parameters by using a compartmental ODE model to emulate the ABM. We use SEIRM and SIRS models for COVID-19 and influenza respectively. Lastly, we compare three versions using GRADABM that leverage gradients during the calibration process, namely calibrated GRADABM (C-GRADABM), deep calibrated GRADABM (DC-GRADABM) and jointly deep calibrated GRADABM (JDC-GRADABM) respectively. **C-GRADABM:** directly optimizes the simulation parameters via gradient descent, without a calibration neural network (CalibNN). To clarify, the gradient update (in Stage 4) is given as:  $[\theta_T^t, \theta_P^t]^{w+1} = [\theta_T^t, \theta_P^t]^w - \alpha \frac{\partial \mathcal{L}(\hat{y}^w, y; (\theta_T^t, \theta_P^t)^w)}{\partial [\theta_T^t, \theta_P^t]^w}$ . **DC-GRADABM:** calibrates the parameters by optimizing weights of CalibNN using automatic differentiation, as described in Sec. 3.2. **JDC-GRADABM:** jointly calibrates parameters for all counties using a shared CalibNN (multi-task learning), as in Sec. 3.2. Here, instead of a new CalibNN for each county, a single CalibNN is shared. Intuitively, one may hope that such a calibration network learns to transfer useful information about disease spread from one county when reasoning about another.

<sup>1</sup>Data links: delphi.cmu.edu; goo.gl/covid19symptomdataset; safegraph.com; coronavirus.jhu.edu; gis.cdc.gov/grasp/fluview/fluportaldashboard.html

## 4.3 Metrics

In order to provide a rigorous evaluation for the aforementioned calibration procedures, we use several standard metrics for evaluating epidemic predictions [7, 67]. Specifically, we use normal deviation (ND), root mean squared error (RMSE) and mean absolute error (MAE). We provide further details in our appendix. Following CDC forecasting guidelines [16, 27], we make weekly predictions for 1 to 4 weeks ahead in the future. Our evaluation for both diseases is of at least 4 months in 10 counties. More details regarding the periods of evaluation and counties investigated are in the appendix.

## 4.4 Implementation Details

In each of our experiments, we allow ABM parameters to change at each time step. That is, quantities such as the disease reproduction number and age-stratified susceptibilities are free to vary as the epidemic progresses. Consider the following intuitive argument for this design choice. It is reasonable to believe that individuals take greater precautions as an epidemic progresses. For example, many people began washing their hands more frequently during the COVID-19 pandemic, reducing the risk of disease transmission. Whilst GRADABM does not explicitly model hand washing, parameters such as age-stratified susceptibility can be varied through time to account for its effect, especially if time series data related to sanitary practices is passed as input to CalibNN. Put differently, by allowing for ABM parameters to vary through time, we may account for environment dynamics or changes in agent behavior that aren't explicitly modeled by the simulator. We stress that our calibration framework does not depend upon CalibNN's neural architecture and that end-to-end training via a combination of backpropagation and gradient descent is possible regardless of CalibNN's implementation. In our own experiments, we design CalibNN so that it may take advantage of the heterogeneous time series data it is passed as input. More specifically, we employ an encoder-decoder architecture, based on GRUs and self-attention [69]. More details regarding our choice of CalibNN architecture can be found in the appendix.<sup>2</sup>

## 5 RESULTS

*First*, we show that calibrating with CALIBNN can improve GRADABM's forecasting accuracy significantly. Moreover, we show that GRADABM scales efficiently as the number of interactions between agents grows. Next, we also show that our calibration procedure is robust to noisy data, in the sense that GRADABM can be well calibrated even with observational error in ground truth; and that GRADABM can be used to conduct sensitivity analyses for evaluating the sensitivity of policy interventions.

### 5.1 Forecasting multiple infectious diseases

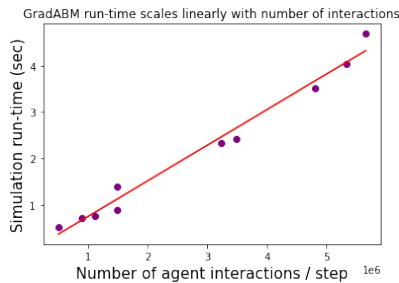
Table 1 displays the results of calibration for each of the benchmarks discussed in Section 4. Observe that JDC-GRADABM outperforms other benchmarks across all metrics for both Influenza and COVID-19 (rows 1-3 in Table 1)<sup>3</sup>. For instance, JDC-GRADABM improves ND to 0.97 from 8.75 in ExpertSearch-ABM and 2.21 in Emulator-ABM on COVID-19; and to 0.41 from 0.57 in ExpertSearch-ABM and

<sup>2</sup>Our code and data is publicly available: <https://github.com/AdityaLab/GradABM>

<sup>3</sup>We performed the unpaired t-test test ( $\alpha = 0.05$ ) over 5 runs to verify JDC-GRADABM performance gains w.r.t other ABM calibration methods are statistically significant

**Table 1: Forecasting results for COVID-19 and influenza over 5 runs. JDC-GRADABM is the only one consistently among the best performing for all (lower error metrics is better).**

Model	COVID-19			Influenza		
	ND	RMSE	MAE	ND	RMSE	MAE
ExpertSearch-ABM [61]	8.75	689.92	270.13	0.57	2.03	1.72
SurrogateODE-ABM [6]	2.21 ± 1.36	121.87 ± 63.97	68.20 ± 41.84	0.59 ± 0.02	2.17 ± 0.05	1.77 ± 0.05
<b>JDC-GRADABM</b>	<b>0.97 ± 0.18</b>	<b>50.99 ± 12.12</b>	<b>30.02 ± 5.60</b>	<b>0.41 ± 0.02</b>	<b>1.47 ± 0.06</b>	<b>1.22 ± 0.06</b>
<b>DC-GRADABM</b>	1.15 ± 0.24	67.09 ± 23.89	35.50 ± 7.36	0.50 ± 0.19	1.78 ± 0.62	1.50 ± 0.57
<b>C-GRADABM</b>	2.39 ± 0.35	205.14 ± 42.56	73.66 ± 10.88	0.88 ± 0.14	2.97 ± 0.44	2.64 ± 0.43

**Figure 3: GRADABM run-time scales linearly with the number of interactions and is roughly 300x faster than prior-art. This benefit is due to the sparse-tensor calculus based design.**

0.59 in Emulator-ABM on Influenza. More generally, we observe that JDC-GRADABM achieves a consistent gain of 8x to 12x over ExpertSearch-ABM on all metrics in COVID-19 experiments. This performance benefit can be attributed to three key features which allow learning personalized time-varying parameters: gradient-based calibration, integration of heterogeneous data sources and joint county training via shared weights of CALIBNN.

Meanwhile, by comparing JDC-GRADABM to both DC-GRADABM and C-GRADABM (rows 3-5 in Table 1), we make some immediate observations. Firstly, using CALIBNN in tandem with GRADABM (JDC-GRADABM, DC-GRADABM) yields meaningful performance improvements over training GRADABM directly via gradient descent (C-GRADABM). Specifically, there is significant improvement in performance in DC-GRADABM over C-GRADABM across all metrics (eg: RMSE goes from 205.99 to 67.09). We conjecture that this is due to the ability of CALIBNN’s to alleviate overfitting by representing parameters in a high-dim space of DNN (instead of scalars in C-GRADABM), inducing a joint distribution over all parameters and incorporating heterogeneous data sources when suggesting parameters. Moreover, we observe that joint training in JDC-GRADABM leads to marginal improvement in performance over DC-GRADABM (eg: RMSE goes from 67.09 to 50.99 in COVID-19 experiments). This indicates that CALIBNN can implicitly use data from one county in order to improve calibration for others (through shared weights)

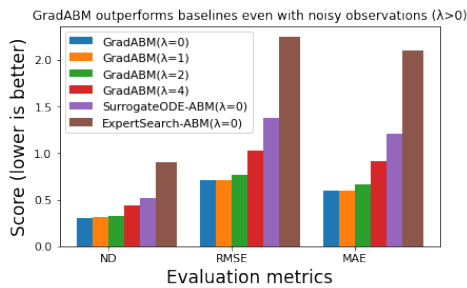
## 5.2 Robustness, Scalability and Decision Making

Next, we study the practical applicability of our framework by focusing on three key considerations: i) scalability to large populations, ii) robustness to observation error in data and iii) utility for practical decision making via sensitivity analyses of policy interventions.

**5.2.1 Scalability of GRADABM simulations.** To investigate scalability, we examine how the simulation time of GRADABM scales as the number of agent interactions increases. As is evident from Figure 3, the run-time for GRADABM scales linearly with number of interactions in the population and executes very quickly (even as the adjacency matrix scales quadratically). For instance, GRADABM executes a simulation with 800,000 agents (5 million interactions) over 133 steps in 4 seconds on a GPU (and 60 seconds on CPU). This is roughly 300x faster than the equivalent Mesa [53] implementation. This performance improvement can be attributed to GRADABM’s sparse tensor-calculus implementation. Since forward simulation is efficient, sensitivity analysis of GRADABM can be conducted rapidly via repeated simulations with different parameter settings. Moreover, iterative algorithms which perform a forward simulation at each step, such as gradient descent, can be used for calibration (as is demonstrated in our work).

**5.2.2 GRADABM robustness to observational error.** To investigate the robustness of our proposed calibration procedure, we run experiments using ground truth data distorted by gaussian noise. More specifically, we add gaussian noise to each ground truth target with mean  $\mu = 0$  and varying scales of standard deviation  $s$ . To set the standard deviation of noise for each county, we first compute the standard deviation of the ground truth data and multiply it by a  $\lambda$  factor. In our experiments, we test four different values for  $\lambda$ . Results are presented in Figure 4. Even for a large degree of noise ( $\lambda = 4$ ), we observe that JDC-ABM outperforms both Emulator-ABM and ExpertSearch-ABM on noiseless data. We attribute this to CALIBNN. More specifically, we conjecture that CALIBNN alleviates overfitting by representing parameters in high-dim space of the neural network (instead of scalars) and allows for integration of heterogeneous data that may assist in selecting appropriate parameters even when the ground truth is noisy.

**5.2.3 Evaluating sensitivity of policy interventions with GRADABM.** Early in the COVID-19 pandemic, uncertainty in vaccine supplies



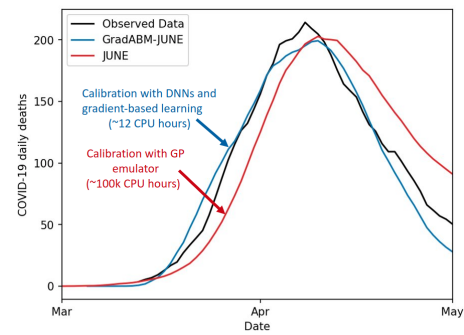
**Figure 4: GRADABM is robust to noisy data. GRADABM achieves lower forecasting error than all baselines even when it is trained with noisy data ( $\lambda > 0$ ) while the baselines receive original data. This is achieved due to differentiable design which allows encoding simulation parameters with CALIBNN and learning using heterogeneous data sources.**

forced policy makers to ask this question: *what-if we delay administration of second dose of the COVID-19 vaccine?*. This decision is sensitive to several variables, including the protection offered by the first vaccine dose, which was unknown at the time [36, 47, 61]. Here, we elucidate how GRADABM can help with decision making under uncertainty by enabling analysis of the sensitivity of policy interventions. More specifically, we consider a scenario with two alternative policies - P1: second dose is administered under standard schedule and P2: second dose is administered with a delay (more details in the appendix). Following the exact experimental setup from clinical works on this question [61], we vary the protection of the first COVID-19 dose from 50% to 80% in ten percentile increments and run simulations for each of the four configurations under both policies. Then, we compare both policies by computing the ratio of cumulative deaths of P2 by P1, which we denote as relative mortality. Basically, if the relative mortality is less than 1, then policy P2 is better (can delay the second COVID-19 dose); while relative mortality greater than 1 implies that policy P1 is better (don't delay the second COVID-19 dose). Results for Franklin County, MA are shown in the appendix with additional details. We observe that once the protection of the first COVID-19 vaccine dose is greater than 60%, GRADABM recommends policy P2 (to delay the second dose). These results are consistent with past clinical works that evaluated this policy intervention [40, 61].

## 6 CONCLUSION AND FUTURE WORK

We introduced GRADABM, a design for agent-based modeling that is amenable to gradient-based learning with automatic differentiation. GRADABM achieves this via a tensorized implementation where each agent state is represented as a vector, the interaction networks as adjacency matrices and all discrete distributions (e.g., Bernoulli) are reparameterized for end-to-end differentiability. Experiments demonstrate that GRADABM can quickly simulate realistic population sizes, integrate with deep neural networks and ingest heterogeneous data sources. This provides an array of practical benefits for calibration, forecasting, and evaluating policy interventions.

The compatibility of GRADABM with automatic differentiation opens up multiple directions of future work both in terms of utility



**Figure 5: Daily deaths in London during the first wave of COVID-19 for the calibrated models JUNE [12] and GRADABM-JUNE and observed data. Note we obtain the similar quality of fit with much less computing time.**

to epidemiology and for advancing the science of agent-based modeling. First, in our work, disease progression in an individual was modeled using a linear deterministic model, but future work could explore more complex and stochastic models (e.g., [61]). Future work could also explore how to use gradient-based explainability methods [63] to have a better understanding of the underlying drivers of epidemic predictions. In this work, calibration with CALIBNN predicts the maximum likelihood estimates but our design allows us to also estimate the posterior distribution of the parameters using techniques from generative DNNs such as normalizing flows [51]. In addition, ABMs may be mis-specified for the actual disease dynamics for which we may have to change the model specification. In this work, we focused on GRADABM using the OpenABM epidemiological model [35] designed for the US. As we explain next, GRADABM generalizes to other model specifications.

The JUNE epidemiological model [12] is a large-scale agent-based model for epidemiology designed to support the National Health Service in England. It was also adopted by the UN Global Pulse to study refugee settlements [13]. In the England setting, JUNE simulates the movement and interactions of 55 million agents and features significantly more complex contact graphs than OpenABM, since it is based on the highest resolution of English census data. We were able to successfully translate it to the fully tensorized design by converting the contact structure in a heterogeneous graph by representing agents and locations as different node types. In this new implementation, which we refer to as GRADABM-JUNE, we achieve a 10,000x speed increase in simulation time. Due to the high computational cost of JUNE, the original model was calibrated using a surrogate model in the form of a Gaussian Process (GP) emulator [70]. The differentiability of the new implementation coupled with the CALIBNN pipeline allows us to quickly find parameter sets (reducing calibration time from 100,000 CPU hours to 12 CPU hours) that match the observed data well (Fig. 5) and are comparable in quality to the fit of the GP emulator. This case study is an encouraging evidence on the utility of GRADABM for creating real impact and also is a step towards our vision of making GRADABM a domain-agnostic design for scalable and differentiable agent-based modeling.



## REFERENCES

- [1] 2020. CDC - U.S. Influenza Surveillance System: Purpose and Methods. <https://www.cdc.gov/flu/weekly/overview.htm>
- [2] 2021. Google LLC "Google COVID-19 Search Trends symptoms dataset". <http://goo.gl/covid19symptomdataset> Accessed: 04.28.2022.
- [3] 2021. U.S. Census Bureau - Data U.S. Census. <https://www.census.gov/data.html> Accessed: 04.28.2022.
- [4] 2022. CDC – Key Facts About Influenza (Flu). <https://www.cdc.gov/flu/about/keyfacts.htm> Accessed: 04.28.2022.
- [5] Elaheh Abdollahi, David Champredon, Joanne M Langley, Alison P Galvani, and Seyed M Moghadas. 2020. Temporal estimates of case-fatality rate for COVID-19 outbreaks in Canada and the United States. *Cmaj* 192, 25 (2020), E666–E670.
- [6] Matthew Abueg, Robert Hinch, Neo Wu, Luyang Liu, William Probert, Austin Wu, Paul Eastham, Yusef Shafi, Matt Rosencrantz, Michael Dikovskiy, et al. 2021. Modeling the effect of exposure notification and non-pharmaceutical interventions on COVID-19 transmission in Washington state. *NPJ digital medicine* 4, 1 (2021), 1–10.
- [7] Bijaya Adhikari, Xinfeng Xu, Naren Ramakrishnan, and B. Aditya Prakash. 2019. EpiDeep: Exploiting Embeddings for Epidemic Forecasting. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining - KDD '19*. ACM Press, Anchorage, AK, USA, 577–586. <https://doi.org/10.1145/3292500.3330917>
- [8] Philipp Andelfinger. 2021. Differentiable Agent-Based Simulation for Gradient-Guided Simulation-Based Optimization. In *Proceedings of the 2021 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation*. 27–38.
- [9] Claudio Angione, Eric Silverman, and Elisabeth Yaneske. 2022. Using machine learning as a surrogate model for agent-based simulations. *Plos one* 17, 2 (2022), e0263150.
- [10] Sercan Arik, Chun-Liang Li, Jinsung Yoon, Rajarishi Sinha, Arkady Epshteyn, Long Le, Vikas Menon, Shashank Singh, Leyou Zhang, Martin Nikoltshev, et al. 2020. Interpretable sequence learning for COVID-19 forecasting. *Advances in Neural Information Processing Systems* 33 (2020), 18807–18818.
- [11] Gaurav Arya, Moritz Schauer, Frank Schäfer, and Chris Rackauckas. 2022. Automatic Differentiation of Programs with Discrete Randomness. *arXiv preprint arXiv:2210.08572* (2022).
- [12] Joseph Aylett-Bullock, Carolina Cuesta-Lazaro, Arnau Quera-Bofarull, Miguel Icaza-Lizaola, Aidan Sedgewick, Henry Truong, Aoife Curran, Edward Elliott, Tristan Caulfield, Kevin Fong, et al. 2021. June: open-source individual-based epidemiology simulation. *Royal Society open science* 8, 7 (2021), 210506.
- [13] Joseph Aylett-Bullock, Carolina Cuesta-Lazaro, Arnau Quera-Bofarull, Anjali Katta, Katherine Hoffmann Pham, Benjamin Hoover, Hendrik Strobel, Rebecca Moreno Jimenez, Aidan Sedgewick, Egmond Samir Evers, David Kennedy, Sandra Harlass, Allen Gidraf Kahindo Maina, Ahmad Hussien, and Miguel Luengo-Oroz. 2021. Operational response simulation tool for epidemics within refugee and IDP settlements: A scenario-based case study of the Cox's Bazar settlement. *PLoS Computational Biology* 17, 10 (Oct. 2021), e1009360. <https://doi.org/10.1371/journal.pcbi.1009360> Publisher: Public Library of Science.
- [14] Munkhtulga Battogtokh and Rita Borgo. 2022. Simple Techniques for a Novel Human Body Pose Optimisation Using Differentiable Inverse Rendering. *Eurographics 2022-Short Papers* (2022), 65–684.
- [15] Yoshua Bengio, Prateek Gupta, Tegan Maharaj, Nasim Rahaman, Martin Weiss, Tristan Deleu, Eilif Muller, Meng Qu, Victor Schmidt, Pierre-Luc St-Charles, et al. 2020. Predicting infectiousness for proactive contact tracing. *arXiv preprint arXiv:2010.12536* (2020).
- [16] Matthew Biggerstaff, David Alper, Mark Dredze, Spencer Fox, Isaac Chun-Hai Fung, Kyle S. Hickmann, Bryan Lewis, Roni Rosenfeld, Jeffrey Shaman, Ming-Hsiang Tsou, Paola Velardi, Alessandro Vespignani, and Lyn Finelli. 2016. Results from the centers for disease control and prevention's predict the 2013–2014 Influenza Season Challenge. *BMC Infectious Diseases* 16, 1 (Dec. 2016), 357. <https://doi.org/10.1186/s12879-016-1669-x>
- [17] Md Arif Billah, Md Mamun Miah, and Md Nuruzzaman Khan. 2020. Reproductive number of coronavirus: A systematic review and meta-analysis based on global level evidence. *PloS one* 15, 11 (2020), e0242128.
- [18] Keith R Bisset, Jiangzhuo Chen, Xizhou Feng, VS Anil Kumar, and Madhav V Marathe. 2009. EpiFast: a fast algorithm for large scale realistic epidemic simulations on distributed memory systems. In *Proceedings of the 23rd international conference on Supercomputing*. 430–439.
- [19] Eric Bonabeau. 2002. Agent-based modeling: Methods and techniques for simulating human systems. *Proceedings of the national academy of sciences* 99, suppl 3 (2002), 7280–7287.
- [20] Julii Brainard, PR Hunter, and Ian R Hall. 2020. An agent-based model about the effects of fake news on a norovirus outbreak. *Revue d'épidémiologie et de sante publique* 68, 2 (2020), 99–107.
- [21] Michael M. Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. 2017. Geometric Deep Learning: Going beyond Euclidean data. *IEEE Signal Processing Magazine* 34, 4 (July 2017), 18–42. <https://doi.org/10.1109/MSP.2017.2693418>
- [22] Serina Chang, Mandy L Wilson, Bryan Lewis, Zakaria Mehrab, Komal K Dudakiya, Emma Pierson, Pang Wei Koh, Jaline Gerardin, Beth Redbird, David Grusky, et al. 2021. Supporting covid-19 policy response with large-scale mobility-based modeling. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 2632–2642.
- [23] Chengqian Che, Fujun Luan, Shuang Zhao, Kavita Bala, and Ioannis Kgioulekas. 2020. Towards learning-based inverse subsurface scattering. In *2020 IEEE International Conference on Computational Photography (ICCP)*. IEEE, 1–12.
- [24] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning Phrase Representations using RNN Encoder–Decoder for Statistical Machine Translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 1724–1734.
- [25] Ayush Chopra, Ramesh Raskar, Jayakumar Subramanian, Balaji Krishnamurthy, Esma S Gel, Santiago Romero-Brufau, Kalyan S Pasupathy, and Thomas C Kingsley. 2021. DeepABM: scalable and efficient agent-based simulations via geometric learning frameworks—a case study for COVID-19 spread and interventions. In *2021 Winter Simulation Conference (WSC)*. IEEE, 1–12.
- [26] GMAM Chowell, MA Miller, and C Viboud. 2008. Seasonal influenza in the United States, France, and Australia: transmission and prospects for control. *Epidemiology & Infection* 136, 6 (2008), 852–864.
- [27] Estee Y Cramer, Evan L Ray, Velma K Lopez, Johannes Bracher, Andrea Brennen, Alvaro J Castro Rivadeneira, Aaron Gerding, Tilmann Gneiting, Katie H House, Yuxin Huang, et al. 2022. Evaluation of individual and ensemble probabilistic forecasts of COVID-19 mortality in the United States. *Proceedings of the National Academy of Sciences* 119, 15 (2022), e2113561119.
- [28] Nediako B Dimitrov and Lauren Anzel Meyers. 2010. Mathematical approaches to infectious disease prediction and control. In *Risk and optimization in an uncertain world*. INFORMS, 1–25.
- [29] Ken Eames, Shweta Bansal, Simon Frost, and Steven Riley. 2015. Six challenges in measuring contact networks for use in modelling. *Epidemics* 10 (2015), 72–77.
- [30] Wouter Edeling, Hamid Arabnejad, Robert C Sinclair, Diana Suleimenova, Krishnakumar Gopalakrishnan, Bartosz Bosak, Derek Groen, Imran Mahmood, Qureshi Hashmi, Daan Crommelin, et al. 2020. Model uncertainty and decision making: Predicting the Impact of COVID-19 Using the CovidSim Epidemiological Code. (2020).
- [31] Luca Ferretti, Chris Wymant, Michelle Kendall, Lele Zhao, Anel Nurtay, Lucie Abeler-Dörner, Michael Parker, David Bonsall, and Christophe Fraser. 2020. Quantifying SARS-CoV-2 transmission suggests epidemic control with digital contact tracing. *Science* 368, 6491 (2020), eabb6936.
- [32] Zipeng Fu, Ashish Kumar, Ananye Agarwal, Haozhi Qi, Jitendra Malik, and Deepak Pathak. 2022. Coupling vision and proprioception for navigation of legged robots. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 17273–17283.
- [33] Herbert W Hethcote. 2000. The mathematics of infectious diseases. *SIAM review* 42, 4 (2000), 599–653.
- [34] Robert Hinch, William JM Probert, Anel Nurtay, Michelle Kendall, Chris Wymant, Matthew Hall, Katrina Lythgoe, Ana Bulas Cruz, Lele Zhao, Andrea Stewart, et al. 2021. OpenABM-Covid19—An agent-based model for non-pharmaceutical interventions against COVID-19 including contact tracing. *PLoS computational biology* 17, 7 (2021), e1009146.
- [35] Robert Hinch, William J. M. Probert, Anel Nurtay, Michelle Kendall, Chris Wymant, Matthew Hall, Katrina Lythgoe, Ana Bulas Cruz, Lele Zhao, Andrea Stewart, Luca Ferretti, Daniel Montero, James Warren, Nicole Mather, Matthew Abueg, Neo Wu, Olivier Legat, Katie Bentley, Thomas Mead, Kelvin Van-Vuuren, Dylan Feldner-Busztin, Tommaso Ristori, Anthony Finkelstein, David G. Bonsall, Lucie Abeler-Dörner, and Christophe Fraser. 2021. OpenABM-Covid19—An agent-based model for non-pharmaceutical interventions against COVID-19 including contact tracing. *PLoS Computational Biology* 17, 7 (July 2021), e1009146. <https://doi.org/10.1371/journal.pcbi.1009146>
- [36] A Hogan, P Winskill, O Watson, P Walker, C Whittaker, M Baguelin, D Haw, A Lochen, K Gaythorpe, K Ainslie, et al. 2020. Report 33: Modelling the allocation and impact of a COVID-19 vaccine. (2020).
- [37] John H Holland and John H Miller. 1991. Artificial adaptive agents in economic theory. *The American economic review* 81, 2 (1991), 365–370.
- [38] Taylor A Howell, Simon Le Cleac'h, J Zico Kolter, Mac Schwager, and Zachary Manchester. 2022. Dojo: A Differentiable Simulator for Robotics. *arXiv preprint arXiv:2203.00806* (2022).
- [39] Iris AM Huijben, Wouter Kool, Max Benedikt Paulus, and Ruud JG Van Sloun. 2022. A Review of the Gumbel-max Trick and its Extensions for Discrete Stochasticity in Machine Learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2022).
- [40] Natsuko Imai, Thomas Rawson, Edward S Knock, and Raphael Sonabend. 2021. Quantifying the impact of delaying the second COVID-19 vaccine dose in England: a. *Lancet* 397 (2021), 881–91.
- [41] Eric Jang, Shixiang Gu, and Ben Poole. 2017. Categorical Reparametrization with Gumbel-Softmax. In *International Conference on Learning Representations (ICLR 2017)*.

- [42] Eric Jang, Shixiang Gu, and Ben Poole. 2017. Categorical Reparametrization with Gumbel-Softmax. In *International Conference on Learning Representations*.
- [43] Harshavardhan Kamarthi, Ling kai Kong, Alexander Rodríguez, Chao Zhang, and B Aditya Prakash. 2022. CAMul: Calibrated and f=Accurate Multi-view Time-Series Forecasting. In *Proceedings of the ACM Web Conference 2022*. 3174–3185.
- [44] Harshavardhan Kamarthi, Alexander Rodríguez, and B Aditya Prakash. 2022. Back2Future: Leveraging Backfill Dynamics for Improving Real-time Predictions in Future. *International Conference on Learning Representations* (2022).
- [45] Daniel B Larremore, Bryan Wilder, Evan Lester, Soraya Shehata, James M Burke, James A Hay, Milind Tambe, Michael J Mina, and Roy Parker. 2021. Test sensitivity is secondary to frequency and turnaround time for COVID-19 screening. *Science advances* 7, 1 (2021), eabd5393.
- [46] Alexander Lavin, Hector Zenil, Brooks Paige, David Krakauer, Justin Gottschlich, Tim Mattson, Anima Anandkumar, Sanjay Choudry, Kamil Rocki, Atılım Güneş Baydin, et al. 2021. Simulation Intelligence: Towards a New Generation of Scientific Methods. *arXiv preprint arXiv:2112.03235* (2021).
- [47] Azeem Majeed, Katrina Pollock, Simon Hodes, and Marisa Papaluca. 2022. Implementation of covid-19 vaccination in the United Kingdom. *bmj* 378 (2022).
- [48] Madhav Marathe and Anil Kumar S Vullikanti. 2013. Computational epidemiology. *Commun. ACM* 56, 7 (2013), 88–96.
- [49] Hadrián Montes-Campos, Jesús Carrete, Sebastian Bichelmaier, Luis M Varela, and Georg KH Madsen. 2021. A Differentiable Neural-Network Force Field for Ionic Liquids. *Journal of chemical information and modeling* 62, 1 (2021), 88–101.
- [50] Kerri-Ann Norton, Chang Gong, Samira Jamalian, and Aleksander S Popel. 2019. Multiscale agent-based and hybrid modeling of the tumor immune microenvironment. *Processes* 7, 1 (2019), 37.
- [51] George Papamakarios, Eric T Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. 2021. Normalizing Flows for Probabilistic Modeling and Inference. *J. Mach. Learn. Res.* 22, 57 (2021), 1–64.
- [52] Lorenzo Pellis, Frank Ball, Shweta Bansal, Ken Eames, Thomas House, Valerie Isham, and Pieter Trapman. 2015. Eight challenges for network epidemic models. *Epidemics* 10 (2015), 58–62.
- [53] Thomas Pike. 2019. Multi-Level Mesa. *arXiv preprint arXiv:1904.08315* (2019).
- [54] Zhaozhi Qian, Ahmed M Alaa, and Mihaela van der Schaar. 2020. When and How to Lift the Lockdown? Global COVID-19 Scenario Analysis and Policy Assessment using Compartmental Gaussian Processes. *Advances in Neural Information Processing Systems* 33 (2020).
- [55] Dylan Radovic, Lucas Kruitwagen, Christian Schroeder de Witt, Ben Caldecott, Shane Tomlinson, and Mark Workman. 2021. Revealing Robust Oil and Gas Company Macro-Strategies Using Deep Multi-Agent Reinforcement Learning. Available at SSRN 3933996 (2021).
- [56] Nicholas G Reich, Logan C Brooks, Spencer J Fox, Sasikiran Kandula, Craig J McGowan, Evan Moore, Dave Osthus, Evan L Ray, Abhinav Tushar, Teresa K Yamana, et al. 2019. A collaborative multiyear, multimodel assessment of seasonal influenza forecasting in the United States. *Proceedings of the National Academy of Sciences* 116, 8 (2019), 3146–3154.
- [57] Theresa Reiker, Monica Columbeau, Andrew Shattock, Lydia Burgert, Thomas A Smith, Sarah Filippi, Ewan Cameron, and Melissa A Penny. 2021. Machine learning approaches to calibrate individual-based infectious disease models. *medRxiv* (2021).
- [58] Alexander Rodríguez, Jiaming Cui, Naren Ramakrishnan, Bijaya Adhikari, and B Aditya Prakash. 2022. EINNs: Epidemiologically-Informed Neural Networks. *arXiv preprint arXiv:2202.10446* (2022).
- [59] Alexander Rodríguez, Harshavardhan Kamarthi, Pulak Agarwal, Javen Ho, Mira Patel, Suchet Sapre, and B. Aditya Prakash. 2022. Data-Centric Epidemic Forecasting: A Survey. *arXiv preprint arXiv:2207.09370* (2022).
- [60] Alexander Rodríguez, Nikhil Muralidhar, Bijaya Adhikari, Anika Tabassum, Naren Ramakrishnan, and B Aditya Prakash. 2021. Steering a Historical Disease Forecasting Model Under a Pandemic: Case of Flu and COVID-19. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 4855–4863. Issue: 6.
- [61] Santiago Romero-Brufau, Ayush Chopra, Alex J Ryu, Esma Gel, Ramesh Raskar, Walter Kremers, Karen S Anderson, Jayakumar Subramanian, Balaji Krishnamurthy, Abhishek Singh, Kalyan Pasupathy, Yue Dong, John C O'Horo, Walter R Wilson, Oscar Mitchell, and Thomas C Kingsley. 2021. Public health impact of delaying second dose of BNT162b2 or mRNA-1273 covid-19 vaccine: simulation agent based modeling study. *BMJ* (May 2021), n1087. <https://doi.org/10.1136/bmj.n1087>
- [62] Samuel Schoenholz and Ekin Dogus Cubuk. 2020. Jax md: a framework for differentiable physics. *Advances in Neural Information Processing Systems* 33 (2020), 11428–11441.
- [63] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*. 618–626.
- [64] Egemen Sert, Yaneer Bar-Yam, and Alfredo J Morales. 2020. Segregation dynamics with reinforcement learning and agent based modeling. *Scientific reports* 10, 1 (2020), 1–12.
- [65] Jeffrey Shaman, Alicia Karspeck, Wan Yang, James Tamerius, and Marc Lipsitch. 2013. Real-time influenza forecasts during the 2012–2013 season. *NATURE COMMUNICATIONS* (2013), 10.
- [66] Jeffrey Shaman, Virginia E Pitzer, Cécile Viboud, Bryan T Grenfell, and Marc Lipsitch. 2010. Absolute humidity and the seasonal onset of influenza in the continental United States. *PLoS biology* 8, 2 (2010), e1000316.
- [67] Farzaneh Sadat Tabataba, Prithwish Chakraborty, Naren Ramakrishnan, Srinivasan Venkatramanan, Jiangzhuo Chen, Bryan Lewis, and Madhav Marathe. 2017. A framework for evaluating epidemic forecasts. *BMC Infectious Diseases* 17, 1 (Dec. 2017), 345. <https://doi.org/10.1186/s12879-017-2365-1>
- [68] Farzaneh S Tabataba, Bryan Lewis, Milad Hosseinpour, Foroogh S Tabataba, Srinivasan Venkatramanan, Jiangzhuo Chen, Dave Higdon, and Madhav Marathe. 2017. Epidemic forecasting framework combining agent-based models and smart beam particle filtering. In *2017 IEEE international conference on data mining (ICDM)*. IEEE, 1099–1104.
- [69] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems* 30 (2017).
- [70] I. Vernon, J. Owen, J. Aylett-Bullock, C. Cuesta-Lazaro, J. Frawley, A. Quera-Bofarull, A. Sedgewick, D. Shi, H. Truong, M. Turner, J. Walker, T. Caulfield, K. Fong, and F. Krauss. 2022. Bayesian emulation and history matching of JUNE. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 380, 2233 (Oct. 2022), 20220039. <https://doi.org/10.1098/rsta.2022.0039> Publisher: Royal Society.
- [71] Joseph T Wu, Kathy Leung, and Gabriel M Leung. 2020. Nowcasting and forecasting the potential domestic and international spread of the 2019-nCoV outbreak originating in Wuhan, China: a modelling study. *The Lancet* 395, 10225 (2020), 689–697.
- [72] Stephan Zheng, Alexander Trott, Sunil Srinivasa, David C Parkes, and Richard Socher. 2022. The AI Economist: Taxation policy design via two-level deep multiagent reinforcement learning. *Science Advances* 8, 18 (2022), eabk2607.