

Behavior Model Calibration for Epidemic Simulations

Meghendra Singh, Achla Marathe, Madhav V. Marathe, Samarth Swarup
 Network Dynamics and Simulation Science Lab, Biocomplexity Institute of Virginia Tech
 Blacksburg, VA
 {meghs,amarathe,mmarathe,swarup}@vt.edu

ABSTRACT

Computational epidemiologists frequently employ large-scale agent-based simulations of human populations to study disease outbreaks and assess intervention strategies. The agents used in such simulations rarely capture the real-world decision-making of human beings. An absence of realistic agent behavior can undermine the reliability of insights generated by such simulations and might make them ill-suited for informing public health policies. In this paper, we address this problem by developing a methodology to create and calibrate an agent decision making model for a large multi-agent simulation, using survey data. Our method optimizes a cost vector associated with the various behaviors to match the behavior distributions observed in a detailed survey of human behaviors during influenza outbreaks. Our approach is a data-driven way of incorporating decision making for agents in large-scale epidemic simulations.

KEYWORDS

Human behavior modeling; Markov decision processes; Agent based simulation.

ACM Reference Format:

Meghendra Singh, Achla Marathe, Madhav V. Marathe, Samarth Swarup. 2018. Behavior Model Calibration for Epidemic Simulations. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018)*, Stockholm, Sweden, July 10–15, 2018, IFAAMAS, 9 pages.

1 INTRODUCTION

Behavior is a crucial aspect of infectious disease outbreak control in general, as demonstrated most recently by the Ebola outbreak and the measles outbreak in the USA. The spread of HIV in parts of Africa also was facilitated by social, economic and behavioral factors [12]. SARS rates fell during the epidemic, partly due to behavioral choices made by individuals which led to a reduction in population contact rates and to rapid hospital attendance by symptomatic individuals [31]. Understanding the interaction of self-initiated individual behavior with disease dynamics is essential while studying epidemic spread through human populations [13].

Influenza epidemics occur annually and place a huge cost upon society [21]. Vaccination rates for seasonal influenza tend to be less than 50% in the USA, which means that strategies for mitigation are very important, including self-initiated behavioral interventions such as hygiene practices, staying home when sick, avoiding crowded places, and more. Additionally, influenza can often have mild or no symptoms, so only a small fraction of the infected go

to hospitals, making it hard to assess the true size of the epidemic. Further, the extent to which people change their behaviors during an influenza outbreak to reduce their susceptibility is not well understood so far. Infectious disease outbreaks also have an associated contagion of information, with people adjusting their behavior based on their perceptions of risk, upon receiving information about the outbreak. Modeling self-initiated human behavior is a very important piece of the puzzle in furthering the understanding of infectious disease epidemics, yet most epidemiological models rarely include models of human behavior.

There have been only a few attempts to model behavior in a data-driven way for social simulation (see [30] for a recent example). Part of the reason for this is the relative paucity of data on behaviors during epidemics. Another crucial issue is how to calibrate behavior models when data *are* indeed available. Computational simulations of epidemics have become quite sophisticated at incorporating multiple sources of data and calibrating disease models [11], but the analogous methodology for representing and calibrating behavior models and integrating them with these large-scale simulations is yet to be developed.

In this work we take a phenomenological approach to behavior modeling, in which we are concerned with getting the population-level distributions of behaviors right, but are less concerned with the individual decision-making process and the psychological and cognitive factors that might be involved. This is because the goal of our work is to be able to simulate populations of behaving agents and to extrapolate the population-level consequences of patterns of behavior. Thus we use a Markov Decision Process (MDP) to model agent decision making about disease avoidance behaviors in a large scale influenza simulation. We use a survey of influenza avoidance behaviors to calibrate the MDP model such that the distribution of these behaviors adopted by the simulated agent population closely matches the distribution of behaviors adopted by individuals in the real world. Figure 1 gives a high-level overview of our data-driven approach to modeling agent behavior.

We begin by analyzing the adoption of influenza avoidance behaviors by a representative sample of the US population captured in a survey. We map these behaviors into interventions in the simulation that reduce the simulated individual's (i.e., agent's) susceptibility to influenza contagions. An agent's decision to adopt a set of behaviors is driven by the MDP, and we posit that the right parameters for the MDP would lead to a close match between the patterns of behavior observed in the simulated and real world. To this end we iteratively tune the parameters of the MDP model used in the epidemic simulator in an optimization loop which is the process of behavior model calibration. We experiment with three optimization methods for the calibration process and find that our approach can closely match observed distributions of disease avoidance behaviors from the survey data quite well in the simulation.

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10–15, 2018, Stockholm, Sweden. © 2018 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

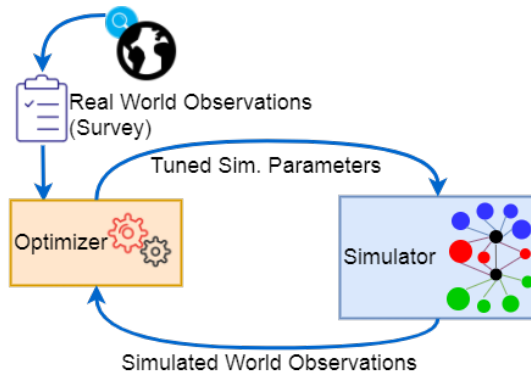


Figure 1: The behavior model calibration approach

The rest of this paper is organized as follows. We begin by discussing related work on behavior modeling from an agent’s perspective and from a public health perspective. Then we describe the survey data which is used to parameterize the behavior model, and the behavior model itself. Next, we describe the epidemic simulator, the synthetic population used in this study and the contagion spread models used by the simulator. Next we detail the agent decision model and the behavior model calibration problem and the optimization approaches used for calibration. We present results of our experiments with different optimization approaches and end with a discussion of limitations and future work.

2 RELATED WORK

Efforts at modeling human behavior and decision-making have a long history in the public health and psychology literature. These approaches generally focus on the within-agent decision-making process, such as the Health Belief Model (HBM) [32], the Theory of Planned Behavior [2], and the Theory of Reasoned Action [3]. These approaches are now also being implemented in computational models [8, 25]. The main limitations of these approaches are that the models seem hard to calibrate since some of the parameters appear to be confounded (e.g., perceived risk and perceived severity), and further the data required to infer the parameters of these models are hard to obtain.

On the other hand, there are many examples of behavior modeling in computational epidemiology and other domains where behaviors are modeled in a counter-factual or prescriptive way. For example, in a study trying to evaluate whether it is better for people to evacuate or shelter in place after a nuclear detonation [26, 38], the authors only have to model the counter-factual scenarios where everyone shelters or everyone evacuates.

Similarly in large-scale simulation studies of flu interventions, modelers generally restrict themselves to assuming that all individuals in the model do the behavior (with some compliance rate). For example, Halloran et al. [15] studied multiple intervention scenarios in a flu pandemic in the city of Chicago. This work ranked interventions by effectiveness, but did not consider people’s typical behavior during flu epidemics.

In a different context, Singh et al. [34] have used the Belief Desire Intention (BDI) framework to create agents which are embedded




Figure 2: Demographic distributions of the survey sample

into large-scale social simulations. Their approach is more focused on the engineering aspects of the problem, such as the agent architecture, modularity, and the design of the simulation platform, than on calibration of the behaviors. They also focus on explainability of agent decision-making and acceptability to the end-users.

The most relevant recent work is that of Pynadath et al. [30]. They used survey data to parametrize a Partially Observable Markov Decision Process (POMDP) model of stay/leave decision-making in a disaster response scenario. They show how to create the states, actions, environment, and rewards for an agent and then train the agent.

The unique contribution of our work is that we bring together survey data on behavior with a large-scale simulation that is capable of implementing those same behaviors. The combination of the two things leads to a more direct calibration method for the behavior models in our work. The survey provides the relative proportions of a sample of the various behaviors in the population. The simulation allows us to determine the probability of agents acting in particular way due to the infectious disease and information flow dynamics in the population. In combination, we can search the space of costs associated with the behaviors (essentially an inverse reinforcement learning setting [22]).

3 SURVEY OF FLU-RELATED BEHAVIORS

We used data from an epidemiological survey aimed at understanding people’s experience with the influenza illness. The survey was administered to 2168 participants which constituted a nationally representative sample of the US population.

Figure 2 shows the demographic distributions of the sample population. The survey captured respondents’ responses about their health behaviors, demographics, risk perceptions, vaccine uptake,




Figure 3: Distribution of individual avoidance interventions in the survey responses

and information sources for outbreak updates. Of particular relevance to this work were the actions taken by respondents to avoid getting influenza (interventions). These interventions are: Avoid touching eyes (ATE), Avoid touching nose (ATN), Avoid touching mouth (ATM), Wash hands with soap (WHS), Use hand sanitizers (UHS), Clean surfaces at home (CSH), Clean surfaces at work (CSW), Eat nutritious food (ENF), Get adequate rest (GAR), Get recommended vaccine (GRV), Take preventive medicine e.g. antiviral (TPM), Use surgical mask to cover nose and mouth (USM), Avoid contact with people who are sick (ACS) and Avoid crowded places (ACP). For each of these interventions, a respondent chose one of three possible responses: *Never*, *Sometimes*, or *Always*. Figure 3 gives the distribution of survey responses for the 14 interventions. We assume that if a respondent selected *Sometimes* or *Always* for an intervention, they undertake that intervention in their daily lives to avoid getting an influenza infection. Selecting *Never* for an intervention implies that the respondent does not undertake that intervention. This simplifying assumption transforms an individual’s choice of interventions into a binary decision problem. Furthermore, an individual can undertake a combination of these interventions, for example, one might undertake three interventions: avoid touching eyes (ATE), wash hands with soap (WHS) and get recommended vaccine (GRV), so as to avoid getting an influenza infection. We found that there were 351 such intervention combinations in the survey responses. In the rest of this paper, we refer to the interventions as actions and a combination of interventions as behavior. Therefore, an individual choice of actions constitutes their behavior. Figure 4 gives the frequencies of each of the behaviors observed in the survey data plotted on \log_{10} scale. Here, the most likely behavior (marked by 1 on the horizontal axis) corresponds to the individual choosing all actions. The first 10 behaviors are selected by 60% of the individuals. Additionally, 58% of the behaviors (i.e., behaviors 145 to 351) have a frequency of 1. Next, we describe the epidemic simulation approach used in our experiments.

4 EPIDEMIC SIMULATION

Large-scale agent-based epidemic simulators work on social contact networks [14], to simulate the spread of social contagions such as epidemics, social memes, news and information through these networks. In this study we use Episimdemics [6], which uses a synthetic population to simulate the spread of infectious diseases through this network. Synthetic populations are commonly used in




Figure 4: Distribution of influenza avoidance behaviors observed in the survey.

large-scale simulations in multiple domains, including influenza epidemics [11, 28], disaster response [27], and more. Next, we describe the process of generating the synthetic population for Montgomery County, Virginia, which is used for simulating influenza in this study.

4.1 Synthetic population

The synthetic population of Montgomery County, VA, is available online ¹. It represents all the residents of Montgomery County and has 77,820 people grouped into 32,827 households. Each synthetic individual is assigned a daily activity schedule, with multiple activities in the day. The total number of activities is 429,590, and these activities are carried out in 26,941 distinct geographic locations. A location can have different sizes and can contain “sublocations” (e.g., rooms within buildings), which can accommodate 25 to 50 people. The synthetic population is constructed in a series of steps by integrating data from multiple sources like the American Community Survey (ACS) [20], the National Household Travel Survey (NHTS) [33] and HERE (formerly NAVTEQ, for road network data). The process employs algorithms like the iterative proportional fitting (IPF) algorithm [7] and gravity model [10] to generate a synthetic population which is statistically indistinguishable from the real population. A person-person social contact network can be derived from this synthetic population by assuming that people who are at the same location for an overlapping period of time are in contact with each other. Details of the process can be found in the report by Adiga et al. [1]. The resulting social contact network has 77,820 nodes representing people and over 2 million edges. Episimdemics operates on this person-person social contact network, and simulates agents or synthetic individuals as nodes of the network [37]. The spread of contagions through the social contact network is modeled through coupled disease propagation and disease progression models.

4.2 Contagion propagation and progression models

The contagion propagation (inter-host) and contagion progression (intra-host) models used in this study were developed by the National Institutes of Health, Models of Infectious Disease Agent Study (MIDAS) project [24]. The inter-host model specifies how an uninfected agent gets exposed to the disease by an infected agent. Such exposures are probabilistic in nature and result from interaction between agents [4]. We can use a probability function to model the probability of a susceptible agent i getting infected based on

¹<http://ndssl.vbi.vt.edu/synthetic-data/>

Table 1: Variable values used in the influenza propagation model

Variable	Value
r in Susceptible states	0.0
r in Exposed state	0.0
r in Infectious state	1.0
r in Recovered state	0.0
s_i in Susceptible states	[0.574649378, 1.0]
s_i in Exposed state	0.0
s_i in Infectious state	0.0
s_i in Recovered state	0.0
ρ	1.0

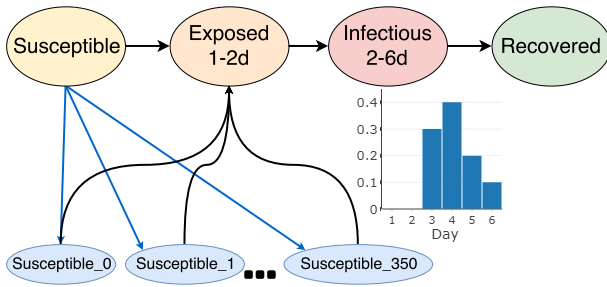


Figure 5: Within-host disease PTTS

its immediate social contact network neighborhood and disease, and agent- and environment-specific parameters. The probability function is defined as,

$$p_i = 1 - \exp\left(-\tau \sum_{r \in R} N_r \ln(1 - r s_i \rho)\right). \quad (1)$$

Here, p_i is the probability of a susceptible agent i getting infected, τ is the duration of exposure in minutes, which is the time spent by agent i at the same location as N_r (collocated) infectious agents, R is the set of infectivities (r s) of each of the N_r infectious agents collocated with the susceptible agent i , s_i is the susceptibility of i and ρ is the probability of a single completely susceptible agent getting infected by a single completely infectious agent through one minute of exposure [5]. The specific values of the disease propagation model parameters used in our influenza simulations are given in Table 1. Upon getting infected the progression of the disease within an agent is modeled using the disease progression or intra-host model. This model is a Probabilistic Timed Transition System (PTTS), which is based on the SEIR model. Figure 5 shows a schematic of the disease PTTS we have used in this study for simulating the progression of influenza within an individual.

Each agent remains in the susceptible (S) state until it comes in contact with one or more infected agents in its social contact




Figure 6: Within-host information PTTS

network neighborhood. Then we calculate p_i based on the susceptibility of agent i (which in turn depends on its behaviors) and the number of infected agents in the neighborhood (N_r). Agent i is then set to be exposed with probability p_i . In this case the agent’s disease PTTS transitions to the exposed (E) state from the Susceptible (S) state. An agent remains in the exposed (E) state for a period of one to two days after which it transitions into the infected (I) state. The exact duration (in hours) for which an agent remains in the exposed state before transitioning to infected state is once again determined by a random draw from the uniform distribution $\mathcal{U}(24, 48)$. However, a transition to the infected state is guaranteed to happen in 48 hours of getting infected. Once an agent transitions to the infected state, it remains infectious for a period of two to six days, after which it transitions to the recovered (R) state. The duration (in days) for which an agent stays in the infected state is determined using the distribution: (3 days/4 days/5 days/6 days), (0.3/0.4/0.2/0.1) as shown in the histogram presented in Figure 5. Agents in the infectious state can spread the infection to other susceptible agents, which come in contact with them through the social contact network as evident by the infectious agent’s infectivity of (1.0). An agent in the recovered state is no longer contagious and cannot get new infections.

The intra-host and inter-host models can also be used to simulate the propagation of other contagions in the social contact network. One such contagion is that of information about the disease outbreak, which can be used by the agents to decide on a disease avoidance behavior. Figure 6 models the spread of information contagion about the influenza outbreak. This within-host information PTTS functions similarly to the disease PTTS, but has only two states (i.e., uninformed and informed). Initially, all agents are in the “uninformed” state, but once an agent gets infected by the influenza contagion, its information state changes to “informed”. Agents in the contact network neighborhood of an infected agent get informed about the outbreak faster than they get infected. This is achieved by having an information contagion susceptibility two times the susceptibility to the disease contagion. This results in the information about the disease outbreak spreading faster than the disease itself. The information contagion model reflects the role of peer influence in an individual’s decision-making. Table 2 lists the specific values of the variables that were used for the information contagion in our simulations. Upon receiving information about the outbreak, an agent can decide to undertake one or more actions so as to avoid getting the disease. In the context of influenza, an agent can undertake any of the 351 combinations of the 14 actions that were discussed in Section 3. For the purpose of this study we assume that adopting a combination of actions (i.e. a behavior) only leads to a reduction in the agent’s susceptibility to influenza. When a susceptible agent decides to undertake an action combination b their disease PTTS transitions to the corresponding Susceptible _{b}

Table 2: Variable values used in the info. propagation model

Variable	Value
r in Uninformed state	0.0
r in Informed state	0.5
s_i in Uninformed state	0.00008
s_i in Informed state	0.0
ρ	1.0

state, which are shown in Figure 5. Each of these 351 susceptible states implies a lower susceptibility towards influenza for the agent than the original Susceptible (S) state. We assume that if an action is easy to take then a lot of people will take it, but it will also have a smaller effect on susceptibility. Hence, the reduction in susceptibility due to an action is an inverse function of the number of survey participants who responded as undertaking that action. Therefore, upon receiving information about the influenza outbreak a susceptible agent can decide to undertake a behavior, which will reduce its susceptibility towards influenza.

The assumed inverse relationship between ease of undertaking an action and the reduction in susceptibility associated with it may or may not hold in the real world. This assumption and the simplified disease and information models used in this study are not central to the behavior calibration methodology and merely serve the purpose of demonstrating the behavior calibration approach, which is the main contribution of this paper. Given the availability of quantitative data linking actions with reduction in susceptibility, one can easily eschew this assumption, without affecting the overall behavior calibration approach. Additionally, the disease and information models can be substituted with more complex variants for modeling any infectious disease as well as peer influence. These aspects are outside the scope of the current study. Next, we define the agent behavior selection process modeled as a MDP.

5 AGENT DECISION MAKING MODEL

We define the Markov decision process (MDP) driving agent decision making by $\langle S, \mathcal{B}, \mathcal{P}, \mathcal{R} \rangle$. The state space S consists of all the possible health states for an agent, which are the states of the intra-host disease PTTS shown in Figure 5. The behavior space \mathcal{B} consists of the 351 behaviors that an agent can adopt. The transition model \mathcal{P} determines the transitions in the state space, given an action $b \in \mathcal{B}$. In our case, \mathcal{P} is intensively specified by the simulation model and cannot be computed extensively. The reward function \mathcal{R} in an MDP determines the expected value of reward received by an agent upon transitioning to a state in S . In the context of the influenza simulation, we associate not getting infected with a positive reward for an individual and taking an action with a negative reward (or cost). Thus, on a particular day in the simulation, if an agent remains in one of the susceptible states they receive a positive reward. An agent can select a behavior, leading to lower susceptibility and increased chance of remaining in the susceptible states. Although this will lead to accumulation of more rewards, selecting a behavior has a cost associated with it which reduces the reward. Hence, the behavioral decision for a susceptible and informed agent

Table 3: Variable values used in the ODE model

Variable	Value
ODE simulation duration	100 days
μ	0.0
β	[0.3, 0.6]
γ	0.125
σ	0.5
Initial proportion of exposed agents	0.0001
Initial proportion of infectious agents	0.0001
Initial proportion of susceptible agents	0.9998
Initial proportion of recovered agents	0.0

is to select an optimal behavior such that its accumulated reward over the simulation duration is maximized. On a particular day in the simulation, this behavioral decision depends on the agent's probability of getting exposed and the cost associated with the 351 behaviors. Taking these aspects into account, we define a value function $V_d(b)$ for a behavior b on day d as,

$$V_d(b) = \sum_{i=d}^D \{(1 - cost_b) * (1 - P_i(S \rightarrow E|b))\}. \quad (2)$$

Here, $cost_b$ is the cost associated with the behavior b and $P_i(S \rightarrow E|b)$ is the probability of any agent in the population transitioning from the susceptible state (S) to the exposed state (E) on the i^{th} day of the simulation, given that the agent undertakes the behavior b . D is the total number of days for which the simulation is being executed. The first term of the value function represents the net reward received by the agent on a particular day, which is the +1 reward for not transitioning to the exposed state minus the cost associated with the behavior b . The second term represents the probability of not transitioning to the exposed state on a particular day. The optimal behavior for the day d will be the one that maximizes the value function. On each day of the simulation, each agent has to choose a behavior based on its estimate of its probability of getting infected. Since an agent does not have access to the full state of the simulated population, it uses a simple differential equation model of the SEIR process [16] to estimate $P_i(S \rightarrow E|b)$:

$$\begin{aligned} \frac{dS}{dt} &= \mu - (\beta I + \mu)S, \\ \frac{dE}{dt} &= \beta SI - (\mu + \sigma)E, \\ \frac{dI}{dt} &= \sigma E - (\mu + \gamma)I, \\ \frac{dR}{dt} &= \gamma I - \mu R. \end{aligned}$$

Here, S, E, I and R represent the number of susceptible, exposed, infectious and recovered agents in a population. μ and β represent the natural mortality and transmission rates respectively. γ and σ represent the recovery and infection rate respectively. For the purpose of this study, the values of variables used in the ODE model are listed in Table 3. We assume the natural mortality rate to be zero and the values of other variables are taken from a standard ODE

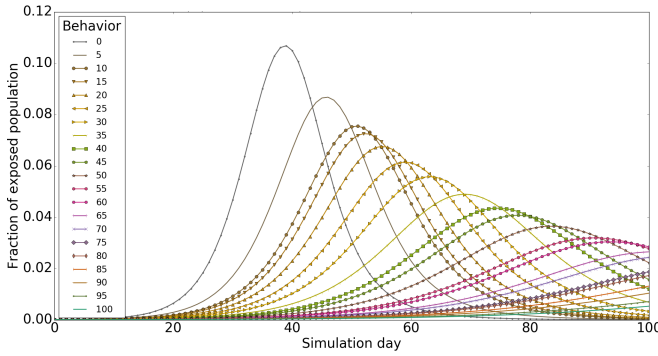


Figure 7: Epicurves generated by the ODE model for 21 of the 351 behaviors

based SEIR model [17]. Also, repeated exposure is not modeled because once an agent gets infected and recovers, they cannot be infected again by the same strain of influenza virus. The ODE model generates epidemic curves or “epicurves”, which represent the fraction of exposed individuals in the population over time. We use these fractions to estimate $P_i(S \rightarrow E|b)$. We assume that adopting one of the 351 behaviors would lead to a change in the transmission rate (β), while the recovery and infection rates (γ and σ) would remain unchanged. This assumption leads to a range of β values as specified in Table 3. This results in different epicurves for each of the 351 behaviors. Figure 7 shows 21 of the 351 epicurves generated for the avoidance behaviors. Given $P_i(S \rightarrow E|b)$, i.e., the probability of any agent in the population transitioning from the susceptible to the exposed state on day i for each of the 351 behaviors and $cost_b$, i.e., the cost associated with each behavior b , we can compute the optimal behavior for any day d in the simulation as $\max_b V_d(b)$, where d refers to the day in the simulation for which the optimal behavior is being computed. During each simulation day any susceptible agent who has received information about the outbreak would choose to undertake the optimal behavior.

However, the costs associated with the behaviors are unknown. The survey described in section 3 gives information about the behaviors that people do, but does not tell the perceived cost of each behavior. Note that cost does not refer to actual dollar costs, but the implicit “behavioral” cost. For example, people may not prefer to use surgical masks, even though they provide the best protection against infection, because of perceptions of a lack of social acceptability to their daily use. To determine these implicit costs that can lead to the observed behavior, we calibrate the model by treating it as an inverse reinforcement learning problem where we have to use the forward simulator in a loop with an optimizer to estimate the cost vector. This is described in the next section.

5.1 Behavior model calibration

We consider the distribution of behaviors selected by participants of the outbreak survey as our real world observations. We calibrate the $cost_b$ parameter of the MDP so as to minimize the mean squared error (MSE) between the simulation and survey distributions of behaviors. Thus, the objective of behavior model calibration is to compute an optimal cost, $cost_b^*$, associated with behavior b ,

$\forall b \in \mathcal{B}$. This can be viewed as an inverse reinforcement learning problem [23] where given the States, Actions and Transitions (i.e., $(\mathcal{S}, \mathcal{A}, \mathcal{P})$) of an MDP, along with some estimate of an optimal policy π^* , an optimal reward function \mathcal{R} is estimated. We specify the costs associated with the behaviors in a vector C , such that each element $c_i \in C$ corresponds to $cost_i \forall 0 \leq i \leq 350$. The objective function $J(C)$ is defined as follows:

$$J(C) = \frac{1}{2 * |C|} \sum_{b \in \mathcal{B}} (N_{C_b} - N_b)^2 \quad (3)$$

$$C^* = \arg \min_C J(C) \quad (4)$$

Here, C is a vector of costs associated with the 351 behaviors. N_{C_b} is the proportion of agents which decide to follow the behavior b in the simulation, for the costs vector C . N_b is the proportion of survey respondents that selected behavior b . In equation (4), C^* is the optimal cost vector, which minimizes the objective function $J(C)$, thereby minimizing the difference between the distribution of behaviors produced by the agent decision model used in the simulation and those observed in the survey data. We experiment with three optimization methods for behavior model calibration, viz. Numerical Gradient Descent (NGD), Cross Entropy (CE) method and Smoothed-Cross Entropy (SCE) method:

- (1) *NGD for behavior model calibration:* In this approach, we begin with a random cost vector \hat{C} and generate $K - 1$ cost perturbations uniformly around \hat{C} . This can be achieved by sampling $K - 1$ points uniformly from the surface of an n -sphere centered at \hat{C} , where $n = 351$ is the dimensionality of \hat{C} . The K cost vectors (i.e. \hat{C} along with $K - 1$ perturbations) form the set of candidate optimal cost vectors which are used to parametrize the agent decision making model in K Epidemics influenza simulations. In each of the K simulations, agents employ the corresponding cost vector along with the ODE SEIR model to decide on behaviors as influenza and information propagate through the social contact network. Each of the K simulations result in a distribution of behaviors. Thus, we can compute the objective function defined in equation (3) for each of the K cost vectors and compute the optimal cost vector C^* out of the K cost vectors which has the minimum value of $J(C)$. We then generate $K - 1$ new cost vectors around C^* and repeat this process over n iterations to get successively better optimal cost vectors, or until the value of the objective function $J(C)$ falls below a particular threshold. Figure 8 shows the flow chart of the NGD algorithm applied to behavior model calibration.
- (2) *CE method for behavior model calibration:* An issue with NGD is slow convergence. One way to achieve faster convergence is to use a fast Monte Carlo-based combinatorial optimization algorithm like the Cross Entropy (CE) method [18, 35]. The CE method operates by generating a random data sample using a particular mechanism (e.g., sampling from a Gaussian distribution), followed by updating the parameters of the mechanism (e.g., updating the mean and variance of the Gaussian distribution) to produce a “better” sample in the next iteration. In the case of using CE method for behavior model calibration, once again we begin with a random cost

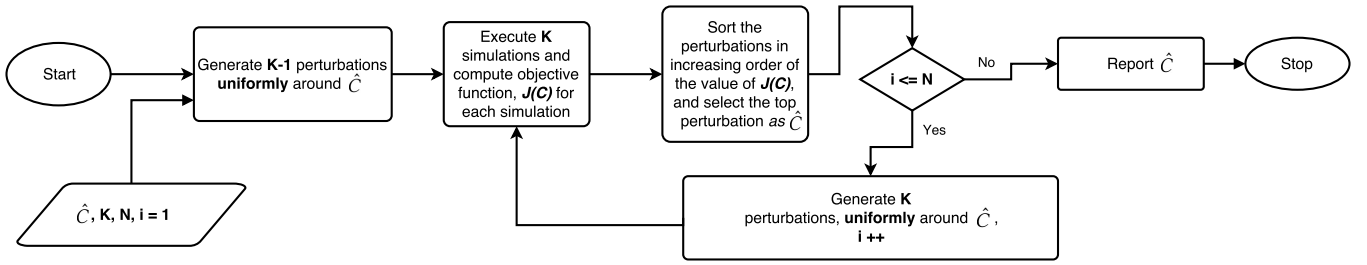


Figure 8: Flow chart of numerical gradient descent for behavior model calibration

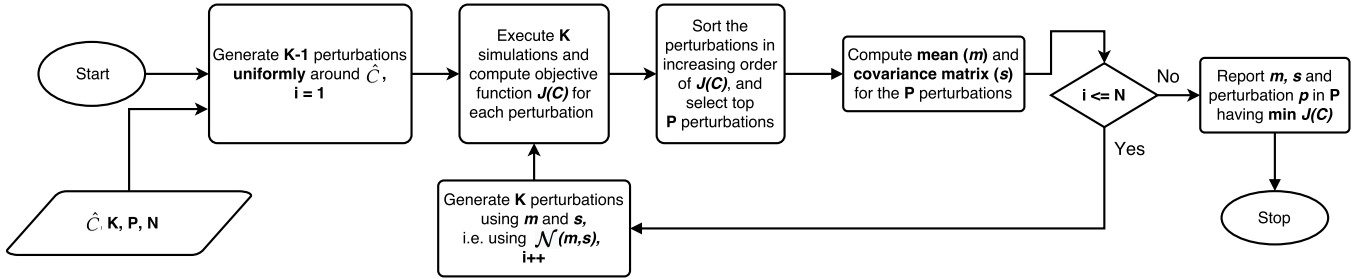


Figure 9: Flow chart of CE method for behavior model calibration

vector \hat{C} and generate $K - 1$ perturbations uniformly around \hat{C} from the surface of an n -sphere centered at \hat{C} . The K perturbations are used to parametrize the agent decision making model in K Epidemics influenza simulations. Similar to NGD, we then compute the value of objective function $J(C)$ for each of the K perturbations. However, instead of selecting one perturbation with the minimum value of $J(C)$, we select P perturbations that have the least values for the objective function $J(C)$ out of the K perturbations. Next, we compute the mean m and covariance matrix S for these top P perturbations and use the multi-variate Gaussian distribution $\mathcal{N}(m, S)$ to once again generate K perturbations for use in the next iteration. We continue this process for n iterations before reporting the results. Figure 9 shows the flow chart for the CE method applied to behavior model calibration.

- (3) *Smoothed-CE method for behavior model calibration*: In our experiments we observed that the covariance matrix S computed by the CE method quickly converged to zeros at an early stage of the optimization. This is analogous to the algorithm getting stuck in a local minimum. In order to prevent this behavior of the CE method, we use a smoothing parameter α to update the mean m_t and covariance matrix S_t on the t^{th} iteration [19].

$$m_t = \alpha \hat{m} + (1 - \alpha)m_{t-1} \quad (5)$$

$$S_t = \alpha \hat{S} + (1 - \alpha)S_{t-1} \quad (6)$$

Here, \hat{m} is the mean and \hat{S} the covariance matrix of the top P perturbations computed for iteration t , m_{t-1} , S_{t-1} is the mean and covariance matrix computed for the last iteration (i.e. $t - 1$) and α is the smoothing parameter such that $0 < \alpha < 1$. Next, we present the calibration results obtained using the NGD, CE and SCE methods.

6 RESULTS

In order to compare the calibration results obtained using the three approaches, we initialized each of them with a common \hat{C} vector, with each element in \hat{C} being equal to a constant c (we kept $c = 0.5$ in our experiments). We initialized the NGD algorithm with $K = 15$ and executed it for $N = 2000$ iterations, which resulted in a $J(C)$ value of 0.0001651. The calibration resulted in a mean error of 0.009268 for the action choice distribution as shown in Figure 11. In order to obtain faster optimization, we experimented with the CE method, initializing it with $K = 30$, $P = 15$ and executed it for $N = 100$ iterations. This resulted in a $J(C)$ value of 0.00020380 and a mean error of 0.013574 for the action choice distributions. Although the final objective function value and the mean error obtained using CE method were worse than those obtained using NGD, the CE method substantially reduced the number of iterations (from 2000 to 100), which amounts to a 20 times speedup, since each iteration takes the same amount of time in both the cases.

We realized that the covariance matrix S computed in the CE method converged to zeros on the 30th iteration, which stopped the optimization process at an early stage. To address this problem we experimented with the Smoothed-CE (SCE) method, initialized with a smoothing factor $\alpha = 0.5$ and executed it for $N = 100$ iterations. This time the calibration process resulted in a $J(C)$ value of 0.00016363 and a mean error of 0.008447, which are an improvement over the CE method results and are much closer to the results obtained using NGD. Figure 10 compares the distribution of the action choices observed in the survey responses with those generated by an uncalibrated agent decision model and those generated after calibration using NGD, CE method and SCE method.




Figure 10: Distribution of influenza avoidance action choices observed in the survey, those generated by an uncalibrated agent decision model and those generated by the three calibration approaches




Figure 11: Mean error between the avoidance action distribution observed in the survey and those generated by the three calibration approaches

7 CONCLUSIONS

In this work, we have shown how to use survey data to calibrate an agent decision making model for a large-scale flu epidemic simulation. The technique is not specific to infectious diseases, so it applies to surveys and simulations in general, though it requires domain-specific decision-making model for each agent akin to the SEIR model used here. The objective of this work was to simulate populations of behaving agents which match the behaving individuals in real populations and exhaustive parameter sweeps for the optimization techniques used in the calibration might improve the results reported here. Additionally, gradient free optimization techniques like Nelder-Mead Simplex or CMA-ES might be better suited for the optimization process. However, all of these aspects are independent of the basic behavior model calibration approach discussed in this work.

While we could construct regression models to predict the probability of an agent adopting a particular behavior using the survey, this does not give us a model of agent decision-making. Our approach gives a model of agent decision-making, which can be interpreted, inspected and compared with other models in the literature. We have developed a very general and practical representation of behaviors for simulations. The behaviors, like options [36], are higher level descriptions that have initiation and termination conditions, and a policy which can alter the actions (in the form of the activity schedule) and the states (in the form of the FSMs) of each agent. This general representation can be adapted to most social simulation scenarios. Of course, there are scenarios where FSMs

will not be powerful enough to represent the agent’s state, but in that case the behavior model will just have to specify how to work with the more complex representation of state.

There are many opportunities for extending this work. In general, the scientific process of gathering data through a survey and then developing a model with it proceeds in an abductive loop in the sense of Peirce [29]. Thus one of the most useful purposes of such simulations is in the “context of discovery”, i.e., to generate new hypotheses. By integrating behavior models, we can now generate detailed forecasts of behavior adoption, which can then be confirmed (or disconfirmed) through new surveys. The goal is to bring rapidity and rigor into the study of human behavior in context. This need to be done by making the behavior modeling process as data driven as possible and then using the models to drive further hypothesis generation and data collection.

In general, human behavior, disease dynamics, and interventions co-evolve. So there is not necessarily a static model of human behavior that can be inferred once and used thereafter. The feedback between these three facets of the system needs to be accounted for. This will require modeling how behaviors change over time and under different circumstances. For example, a challenge would be to be able to predict the level of worry that was observed over Ebola in the US even though the number of actual cases was very small. This requires, as a first step, a more careful modeling of the information contagion process, and extending it to include emotion or fear contagion as well (e.g., see [9]).

ACKNOWLEDGMENTS

We thank our external collaborators and members of the Network Dynamics and Simulation Science Laboratory (NDSSL) for their suggestions and comments. This work has been partially supported by DTRA CNIMS Contract HDTRA1-17-0118, NIH Grant 1R01GM109718, NSF IBSS Grant SMA-1520359, NSF BIG DATA Grant IIS-1633028, NSF DIBBS Grant ACI-1443054, NSF NRT-DESE Grant DGE-154362, and NIH MIDAS Cooperative Agreement U01GM070694.

REFERENCES

- [1] Abhijin Adiga, Aditya Agashe, Shaikh Arifuzzaman, Christopher L. Barrett, Richard J. Beckman, Keith R. Bisset, Jiangzhuo Chen, Youngyun Chungbaek, Stephen G. Eubank, Sandeep Gupta, Maleq Khan, Christopher J. Kuhlman, Eric Lofgren, Bryan L. Lewis, Achla Marathe, Madhav V. Marathe, Henning S. Mortveit, Eric Nordberg, Caitlin Rivers, Paula Stretz, Samarth Swarup, Amanda Wilson, and Dawen Xie. 2015. *Generating a Synthetic Population of the United States*. Technical Report NDSSL 15-009. Network Dynamics and Simulation Science Laboratory.
- [2] Icek Ajzen. 1991. The theory of planned behavior. *Organizational Behavior and Human Decision Processes* 50 (1991), 179–211.
- [3] I. Ajzen and M. Fishbein. 1980. *Understanding attitudes and predicting social behavior*. Prentice-Hall, Englewood Cliffs, NJ.
- [4] Chris Barrett, Stephen Eubank, and Madhav Marathe. 2006. Modeling and simulation of large biological, information and socio-technical systems: an interaction based approach. In *Interactive computation*. Springer, 353–392.
- [5] Christopher L Barrett, Keith Bisset, Stephen Eubank, Madhav V Marathe, VS Anil Kumar, and Henning S Mortveit. 2007. Modeling and simulation of large biological, information and socio-technical systems: An interaction-based approach. In *Proceedings of symposia in applied mathematics*, Vol. 64. 101.
- [6] Christopher L Barrett, Keith R Bisset, Stephen G Eubank, Xizhou Feng, and Madhav V Marathe. 2008. EpiSimdemics: an efficient algorithm for simulating the spread of infectious disease over large realistic social networks. In *Proceedings of the 2008 ACM/IEEE conference on Supercomputing*. IEEE Press, 37.
- [7] R. J. Beckman, K. A. Baggerly, and M. D. McKay. 1996. Creating Synthetic Base-Line Populations. *Transportation Research A – Policy and Practice* 30 (1996), 415–429.
- [8] David P. Durham, Elizabeth A. Casman, and Steven M. Albert. 2012. Deriving Behavior Model Parameters from Survey Data: Self-Protective Behavior Adoption During the 2009–if;2010 Influenza A(H1N1) Pandemic. *Risk Analysis* 32, 12 (2012).
- [9] Joshua M. Epstein, Jon Parker, Derek Cummings, and Ross A. Hammond. 2008. Coupled Contagion Dynamics of Fear and Disease: Mathematical and Computational Explorations. *PLoS ONE* 3, 12 (2008), e3955.
- [10] S. Erlander and N. F. Stewart. 1990. *The Gravity Model in Transportation Analysis: Theory and Extensions*. VSP, Utrecht, The Netherlands.
- [11] S. Eubank, H. Guclu, V. S. Anil Kumar, M. Marathe, A. Srinivasan, Z. Toroczka, and N. Wang. 2004. Modelling Disease Outbreaks in Realistic Urban Social Networks. *Nature* 429 (May 2004), 180–184.
- [12] Carol Ezzell. 2000. Care for a dying continent. *Scientific American* 282, 5 (2000), 96–105.
- [13] Neil Ferguson. 2007. Capturing human behaviour. *Nature* 446, 7137 (2007), 733–733.
- [14] Laura M. Glass and Robert J. Glass. 2008. Social contact networks for the spread of pandemic influenza in children and teenagers. *BMC Public Health* 8, 1 (2008), 61. DOI : <http://dx.doi.org/10.1186/1471-2458-8-61>
- [15] M. Elizabeth Halloran, Neil M. Ferguson, Stephen Eubank, Ira M. Longini, Jr., Derek A. T. Cummings, Bryan Lewis, Shufu Xu, Christophe Fraser, Anil Vullikanti, Timothy C. Germann, Diane Wagener, Richard Beckman, Kai Kadau, Chris Barrett, Catherine A. Macken, Donald S. Burke, and Philip Cooley. 2008. Modeling Targeted Layered Containment of an Influenza Pandemic in the United States. *PNAS* 105, 12 (2008), 4639–4644.
- [16] Herbert W Hethcote. 1976. Qualitative analyses of communicable disease models. *Mathematical Biosciences* 28, 3-4 (1976), 335–356.
- [17] Matt J Keeling and Pejman Rohani. 2008. *Modeling infectious diseases in humans and animals*. Princeton University Press.
- [18] Dirk P Kroese, Reuven Y Rubinstein, Izack Cohen, Sergey Porotsky, and Thomas Taimre. 2013. Cross-entropy method. In *Encyclopedia of Operations Research and Management Science*. Springer, 326–333.
- [19] Shie Mannor, Reuven Y Rubinstein, and Yohai Gat. 2003. The cross entropy method for fast policy search. In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*. 512–519.
- [20] Mark Mather, Kerri L Rivers, and Linda A Jacobsen. 2005. The American Community Survey. *Population Bulletin* 60, 3 (2005), 1–20. <http://www.census.gov/acs/www/>
- [21] Martin I. Meltzer, Nancy J. Cox, and Keiji Fukuda. 1999. The Economic Impact of Pandemic Influenza in the United States: Priorities for Intervention. *Emerging Infectious Diseases* 5, 5 (1999), 659–671.
- [22] Andrew Y. Ng and Stuart Russell. 2000. Algorithms for Inverse Reinforcement Learning. In *Proc. ICML*.
- [23] Andrew Y Ng, Stuart J Russell, and others. 2000. Algorithms for inverse reinforcement learning. In *Icml*. 663–670.
- [24] NIH. 2009. National Institutes of Health 2009, <http://www.nigms.nih.gov/Initiatives/MIDAS/> (2009). <http://www.nigms.nih.gov/Initiatives/MIDAS/>
- [25] Mark G. Orr, Roxanne Thrush, and David C. Plaut. 2013. The Theory of Reasoned Action as Parallel Constraint Satisfaction: Towards a Dynamic Computational Model of Health Behavior. *PLoS ONE* 8, 5 (2013), e62490.
- [26] Nidhi Parikh, Harshal G. Hayatnagarkar, Richard J. Beckman, Madhav V. Marathe, and Samarth Swarup. 2016. A Comparison of Multiple Behavior Models in a Simulation of the Aftermath of an Improvised Nuclear Detonation. *Autonomous Agents and Multi-Agent Systems, Special Issue on Autonomous Agents for Agent-Based Modeling* 30, 6 (2016), 1148–1174. DOI : <http://dx.doi.org/10.1007/s10458-016-9331-y>
- [27] Nidhi Parikh, Samarth Swarup, Paula E. Stretz, Caitlin M. Rivers, Bryan L. Lewis, Madhav V. Marathe, Stephen G. Eubank, Christopher L. Barrett, Kristian Lum, and Youngyun Chungbaek. 2013. Modeling Human Behavior in the Aftermath of a Hypothetical Improvised Nuclear Detonation. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. Saint Paul, MN, USA.
- [28] Nidhi Parikh, Mina Youssef, Samarth Swarup, Stephen Eubank, and Youngyun Chungbaek. 2014. Cover Your Cough! Quantifying the Benefits of a Localized Healthy Behavior Intervention on Flu Epidemics in Washington DC. In *Proceedings of The International Conference on Social Computing, Behavioral-Cultural Modeling, and Prediction (SBP)*. Washington DC, USA.
- [29] Charles Sander Peirce. 1957. The Logic of Abduction. In *Peirce's Essays in the Philosophy of Science*, Thomas V (Ed.). Liberal Arts Press, 195–205.
- [30] David V. Pynadath, Heather Rosoff, and Richard S. John. 2016. Semi-Automated Construction of Decision-Theoretic Models of Human Behavior. In *Proc. AAMAS*.
- [31] Steven Riley, Christophe Fraser, Christl A Donnelly, Azra C Ghani, Laith J Abu-Raddad, Anthony J Hedley, Gabriel M Leung, Lai-Ming Ho, Tai-Hing Lam, Thuan Q Thach, and others. 2003. Transmission dynamics of the etiological agent of SARS in Hong Kong: impact of public health interventions. *Science* 300, 5627 (2003), 1961–1966.
- [32] Irwin M. Rosenstock. 1966. Why People Use Health Services. *Milbank Memorial Fund Quarterly* 44, 3 Pt 2 (1966), 94–124.
- [33] A. Santos, N. McGuckin, H.Y. Nakamoto, D. Gray, and S. Liss. 2011. *Summary of Travel Trends: 2009 National Household Travel Survey*. Technical Report FHW A-PL-11-022. U.S. Department of Transportation Federal Highway Administration.
- [34] Dharendra Singh, Lin Padgham, and Brian Logan. 2016. Integrating BDI Agents with Agent-Based Simulation Platforms. *Auton Agent Multi-Agent Syst Online-First* (2016). DOI : <http://dx.doi.org/10.1007/s10458-016-9332-x>
- [35] Freek Stulp and Olivier Sigaud. 2012. Path Integral Policy Improvement with Covariance Matrix Adaptation. In *Proceedings of the 29th International Conference on Machine Learning (ICML)*. 0–0.
- [36] Rich Sutton, Doina Precup, and Satinder Singh. 1999. Between MDPs and Semi-MDPs: A Framework for Temporal Abstraction in Reinforcement Learning. *Artificial Intelligence* 112, 1-2 (1999), 181–211.
- [37] Samarth Swarup, Stephen Eubank, and Madhav Marathe. 2014. Computational Epidemiology as a Challenge Domain for Multiagent Systems. In *Proceedings of the Thirtieth International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. Paris, France.
- [38] Lawrence M. Wein, Youngsoo Choi, and Sylvie Denuit. 2010. Analyzing Evacuation Versus Shelter-in-Place Strategies After a Terrorist Nuclear Detonation. *Risk Analysis* 30, 9 (2010), 1315–1327.