Proceedings of the 2025 Winter Simulation Conference
E. Azar, A. Djanatliev, A. Harper, C. Kogler, V. Ramamohan, A. Anagnostou, and S. J. E. Taylor, eds.

PYSIRTEM: AN EFFICIENT MODULAR SIMULATION PLATFORM FOR THE ANALYSIS OF PANDEMIC SCENARIOS

Preetom Biswas[1], Giulia Pedrielli[1], and K. Selçuk Candan[1]

[1]School of Computing and Augmented Intelligence, Arizona State University, Tempe, AZ, USA

ABSTRACT

Conventional population-based ODE models struggle against increased level of resolution since incorporating many states exponentially increases computational costs, and demands robust calibration for numerous hyperparameters. PySIRTEM is a spatiotemporal SEIR-based epidemic simulation platform that provides high resolution analysis of viral disease progression and mitigation. Based on the authors-developed Matlab© simulator SIRTEM, PySIRTEM 's modular design reflects key health processes, including infection, testing, immunity, and hospitalization, enabling flexible manipulation of transition rates. Unlike SIRTEM, PySIRTEM uses a Sequential Monte Carlo (SMC) particle filter to dynamically learn epidemiological parameters using historical COVID-19 data from several U.S. states. The improved accuracy (by orders of magnitude) make PySIRTEM ideal for informed decision-making by detecting outbreaks and fluctuations. We further demonstrate PySIRTEM 's usability performing a factorial analysis to assess the impact of different hyperparameter configurations on the predicted epidemic dynamics. Finally, we analyze containment scenarios with varying trends, showcasing PySIRTEM 's adaptability and effectiveness.

1 INTRODUCTION

The devastating outbreak of COVID-19 disease, caused by the airborne SARS-CoV-2 virus, has had a long-lasting impact around the world, with 7 million deaths to date (Hossain et al. 2021). Various mitigation strategies have been employed by policymakers to curb the onslaught of infection and outbreaks such as lockdown, isolation, optimal vaccine allocation, preemptive quarantine, and random testing (Anderson et al. 2020; Walensky and Del Rio 2020; Roy et al. 2021). Despite recent significant breakthroughs in medical science, there are substantial challenges in impeding the destruction of pandemics. The incredible advancement of transportation, especially the ease of international air travel, allows for rapid dissemination of viruses with a high reproduction number such as COVID-19. The lack of reliable testing data for early detection of a novel virus outbreak poses an added obstacle. Acquiring reliable population-level epidemiological data is another obstacle due to underreporting, misreporting, and testing limitations (Lau et al. 2021). In addition, the adaptive and mutative abilities of certain viruses lead to the inception of newer, often more virulent strains, stressing the need for a flexible and meticulous epidemiological model.

Duan et al. (2015) classify epidemiological models into three categories: complex network models, agent-based models, and analytical models. Network-based models focus on person-to-person interactions, representing disease spread as a space-time graph (Chang et al. 2021). Although limited by computational complexity for large-scale studies, they can be coupled with analytical models to infer inter-zonal relation and optimize containment policies. Here, each node acts as an individual zone with internal mobility and edges denote contagion influence between zones (Roy et al. 2021a). Agent-based models, which simulate individual behaviors and interactions, allow for high-resolution analysis of disease transmission. Various studies have employed agent-based modeling in the scope of COVID-19, Dengue, HIV, both independently and coupled with analytical models (Kerr et al. 2021; Rhee 2006; Miksch et al. 2015). Nevertheless, their use is also limited to smaller case studies due to high computational demands. Traditional analytical models such as differential equation-based SIR (susceptible-infected-recovered), SIS (susceptible-infected-

susceptible), SEIR (susceptible-exposed-infected-recovered), and SEIRS (susceptible-exposed-infected-recovered-susceptible) are most commonly used to study population-level pandemic progression (Jit and Brisson 2011; Li and Muldowney 1995).

To support optimized policy making, these models can incorporate additional processes such as quarantine and hospitalization (He et al. 2020). For example, Chu et al. (2020) integrated mask wearing, eye protection, and social distancing to study infection reduction while Maged et al. (2023) modeled masked and unmasked populations separately within an SEIR framework to analyze infection dynamics. Moreover, testing campaigns are crucial in identifying outbreaks, tracking trends, and steering interventions, particularly in diseases like COVID-19 with significant asymptomatic spread. However, testing can be costly and unreliable. So epidemic models must address constraints such as cost, daily capacity, test accuracy, sensitivity and specificity. As mass testing can be costly, choosing optimal testing strategies is significant in identifying target populations that require quarantine (Omori et al. 2020). Wells et al. (2021) also explored reducing the quarantine period with exit testing as an effective alternative to the 14-day full quarantine. Furthermore, embedding spatiotemporal mobility is another key factor in accurately predicting and simulating viral progression across zones and individuals (Niehus et al. 2020; Roy et al. 2021b).

While effective, any singular model fails to simultaneously incorporate all of these important health processes such as immunity, quarantine, hospitalization, and multiple testing modalities while accounting for population movement and viral evolution. In our earlier work (Azad et al. 2022), we introduced a spatially informed SEIRS model, `SIRTEM` (*Spatially Informed Rapid Testing for Epidemic Modelling*), which couples the spatio-temporal dynamics of the COVID-19 epidemic with above-mentioned compartments to solve the optimal cost-effective multi-modal testing strategy problem while satisfying practical constraints. However, the inclusion of such complex processes pose a great challenge in effectively identifying and optimizing intrinsic model parameters. While Sequential Monte Carlo methodology has been applied for pandemic modeling (Papageorgiou and Tsaklidis 2024), its application to a model of this scale and complexity remains largely unexplored. In this paper, we present `PySIRTEM` which implements and extends the `SIRTEM` model to accommodate the dynamic learning of model parameters and provides flexible manipulation to accurately simulate real-world viral interdiction scenarios.

Contributions. `PySIRTEM` builds upon the `SIRTEM` model with a novel Sequential Monte Carlo (SMC) particle filter approach that improves the accuracy of transition hyperparameters. Particle filter's black box nature allows for sensitive identification of rapid fluctuations, making `PySIRTEM` a powerful tool to capture intrinsic pandemic behavior. We further provide a detailed analysis of the interaction between various model parameters and their impact on pandemic progression through Design of Experiments. Lastly, we simulate multiple pandemic scenarios to analyze `PySIRTEM`'s effectiveness in informed decision and policy making.

2 APPROACH

2.1 SEIRS Epidemic Model

For `PySIRTEM`, we extend the SEIRS model (Bjørnstad et al. 2021). Susceptible (S) population consists of individuals who are not exposed to infection. Once exposed, they may transition to the Exposed (E) class that denotes an asymptomatic or untested population. If tested positive, exposed individuals transfer to Infected (I) class from where they transition either to Recovered (R) or dead. SEIRS models further allow incorporation of waning immunity in disease modeling, where a recovered individual can again become susceptible over time. A set of differential equations regulate the transition between the health classes using parameters such as infection rate (β), recovery rate (γ), and loss of immunity rate (δ) (see Figure 1a). However, the traditional SEIRS model fails to incorporate important dynamics and processes such as testing, hospitalization, quarantine, and natural and induced immunization.

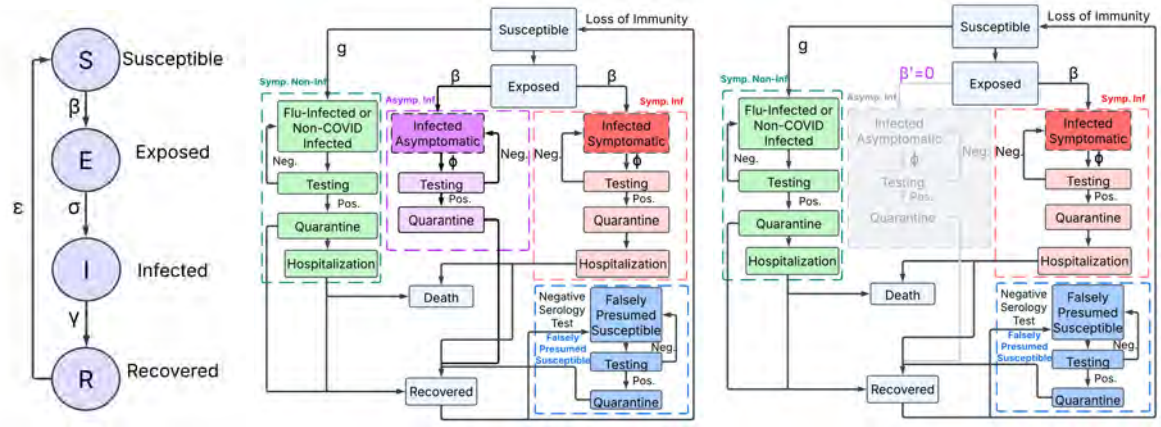

(a) (b) (c)

Figure 1: A brief overview of (a) SEIRS model and (b) `PySIRTEM` model (c) `PySIRTEM` model with Asymptomatic Infected Compartment turned off. Some internal states are not shown in the diagram for simplicity.

2.2 `PySIRTEM`

`PySIRTEM` implements the extended SEIRS model, `SIRTEM` (Azad et al. 2022), to simulate pandemic evolution. `PySIRTEM` builds upon the SEIRS model by incorporating additional states to account for testing, hospitalization, quarantine, and immunity loss. It also makes provisions for optimized symptomatic and asymptomatic testing strategies as well as the accuracy (false positive and false negative rate) of these test results. Furthermore, `PySIRTEM` is spatially informed i.e. accounts for the distribution of population in different zones and associated mixing rates.

`PySIRTEM` models the population using five main compartments: *Susceptible* individuals who are non-infected and display no COVID symptoms; *Symptomatic Infected* individuals showing COVID symptoms; *Asymptomatic Infected* individuals who do not exhibit symptoms; *Symptomatic but not COVID-infected* individuals with flu-like symptoms; and *Falsely Presumed Susceptible* individuals who have natural immunity but are falsely tested negative for COVID antibodies. Each compartment contains multiple internal medical states, resulting in a total of 46 states. The modular nature of PySIRTEM provides high flexibility in adding or removing health processes through simple adjustments to the corresponding rates (Figures 1b and 1c).

`PySIRTEM` also employs Delay Differential Equations (DDEs) instead of traditional Ordinary Differential Equations (ODEs) to accommodate for delays in disease progression and testing results. Analogous to conventional SEIRS models, transition between the compartments are governed by relevant transition rates and parameters such as infection rate, testing rate, hospitalization rate and so on. Comprehensive overview of the subsequent processes and parameters are detailed in the SIRTEM paper (Azad et al. 2022).

2.3 Calibration of Transition Rates using Particle Filters

For a single-zone geographic setup, we calibrate `PySIRTEM` by estimating the three time-varying transition parameters: infection rate (β), testing rate (ϕ), and general non-COVID-19 sickness rate (g) using published reported daily positive and daily negative cases. Note that in a multi-zonal setup, each zone will have its own transition parameters to be estimated. We employ particle filters, also known as Sequential Monte Carlo (SMC) methods, that are highly effective in approximating solutions of non-linear state space systems. Unlike traditional optimization approaches or Kalman filters, particle filters are non-parametric which are ideal for non-linear, non-Gaussian epidemic processes.

2.3.1 Basics of a Particle Filter

Particle filters leverage a set of weighted particles to derive the posterior distribution of latent states in a state-space model. A state-space model consists of a latent Markov process (X_t) and a series of noisy observations (Y_t) at discrete time. The state transition model defines the evolution of the latent state and follows the following probability,

$$X_t \sim P(X_t \mid X_{t-1}).$$

Observations are conditionally independent given the latent states and each observation is only dependant on the current latent state.

$$Y_t \sim g(Y_t \mid X_t),$$

where $g(\cdot)$ denotes the observation likelihood function. Given a set of observations $Y_{1:t} = \{Y_1, Y_2, \ldots, Y_t\}$, the particle filter estimates the posterior distribution $\pi(X_t|Y_{1:t})$ by using a set of N weighted particles $\{x_t^i, W_t^i\}_{i=1}^N$ evolving over time, where x_t is a sampled particle from the distribution of X_t. The weights are normalized such that $\sum_{i=1}^N W_t^i = 1$. The particle filter recursively computes importance sampling approximations of π_t as follows,

$$\pi_t(X_t \mid Y_{1:t}) \approx \hat{\pi}_t(X_t|Y_{1:t}) = \sum_{i=1}^N W_t^i \delta(X_t - X_t^i).$$

Here $\hat{\pi}_t$ is the empirical approximation π_t using weighted particles and δ represents the Dirac delta function. The value of δ is 0 everywhere except at 0, and integrates to 1. The parameters δ is used to approximate the posterior distribution as the sum of weighted point masses. The algorithm consists of three major steps:

i. **Resample:** Draw $({X^*_{t-1}}^1, \ldots, {X^*_{t-1}}^N)$ (with replacement) from $\hat{\pi}_{t-1}$ according to weights $\{W_t^i\}$ to discard lower weighted particle before propagation. Here ${X^*_{t-1}}^i$ is a resampled particle.

ii. **Propagate:** Draw X_t^i from $P(\cdot \mid {X^*_{t-1}}^i)$, independently for different indices i.

iii. **Reweight:** Set $W_t^i \propto g(Y_t \mid X_t^i)$ and Normalize $W_t^i \leftarrow \frac{W_t^i}{\sum_j W_t^j}$.

Note that at time $t = 0$, particles are drawn from π_0 and W_0^i is set to $1/N$. Further detailed discussions on particle filters can be found in (Künsch 2013; Doucet and Johansen 2009) and related literature.

2.3.2 `PySIRTEM` State Space Model Setup And Particle Filter Iteration

In the context of epidemic modeling and forecasting, our goal is to estimate the following latent state vector:

$$X_t = [\beta_t, \phi_t, g_t],$$

where, β_t is the infection rate for susceptible population, ϕ_t is the diagnostic testing rate for symptomatic individuals, and g_t is the ratio of susceptible population who have fever for non-COVID infections.

We use reported daily positive (y^+) and negative (y^-) case counts for training due to their availability and comprehensiveness. However, the framework can be extended to incorporate datasets such as daily deaths. Given X_t, `PySIRTEM` model produces two predicted time series $\hat{Y} = [\hat{y}^+, \hat{y}^-]$ (predicted daily positives $\hat{y}^+$ and negatives $\hat{y}^-$), which are compared against the actual reported case numbers $Y = [y^+, y^-]$.

We assume X_t evolves stochastically with a two-step memory and define the state transition and observation models as

$$X_t = f(X_{t-1}, X_{t-2}) + \varepsilon_t^X,$$
$$Y_t = F(X_t) + \varepsilon_t^Y,$$

where F represents the `PySIRTEM` simulation function that takes in the transition rates and returns $\hat{Y}_t = [\hat{y}_t^+, \hat{y}_t^-]$. Here, both $\varepsilon_t^x \sim \mathcal{N}(0, \Sigma)$ and $\varepsilon_t^y \sim \mathcal{N}(0, \sigma^2)$ denote small random noise.

`PySIRTEM` utilizes a dynamic 2-step memory process where future updates depend on the two preceding states to model the state transitions. We generate $N = 200$ initial particles $\{x_0^i\}_{i=1}^N$ from a prior uniform distribution $U(0,1)$. In particle filter, approximation accuracy increases with the number of particles. Since large number of particles adds significant computational cost for `PySIRTEM`, we use an importance density

(or proposal distribution) that incorporates the previous two predicted states (Equation 1, 2, 3), limiting search in relevant state space and providing better approximation with fewer particles (Arulampalam et al. 2002). Each particle, x_t^i, the state transition function is called and calculated using x_{t-1}^i and x_{t-2}^i. For x_1^i, only initial x_0^i is used.

The transition rates $x_t = [\beta_t, \phi_t, g_t]$ are only updated weekly (a choice made for simplicity) and kept constant otherwise. For a given week k, the rates are calculated as follows:

$$\beta_k = w_1^\beta \cdot \beta_{k-1} + w_2^\beta \cdot \beta_{k-2} + \varepsilon_k^\beta \tag{1}$$

$$\phi_k = w_1^\phi \cdot \phi_{k-1} + w_2^\phi \cdot \phi_{k-2} + \varepsilon_k^\phi \tag{2}$$

$$g_k = w_1^g \cdot g_{k-1} + w_2^g \cdot g_{k-2} + \varepsilon_k^g \tag{3}$$

Above, the weights $w_1^\beta, w_2^\beta, w_1^\phi, w_2^\phi, w_1^g$, and w_2^g are randomly sampled between $[0.1, 0.8]$ each week so that the model does not overfit by allowing unrealistic fluctuations while providing enough flexibility to reflect real-world changes quickly. A small Gaussian perturbation $\varepsilon_k \sim \mathcal{N}(0, \sigma^2)$ is added for stochasticity. The new rates are bounded between $[0, 1]$ to ensure realistic values.

The `PySIRTEM` model simulation is run using each of the state particles as input:

$$[\hat{y}_t^{+^i}, \hat{y}_t^{-^i}] = F(\beta_t^i, \phi_t^i, g_t^i).$$

We compare the outputs with reported case data obtained from public sources (CTW nd) and update the weights by computing the likelihood defined on line 8 of Algorithm 1, where σ_{obs}^2 represents the noise in observational data. The weights are normalized and N particles are resampled at the next time point proportionate to their updated weights, but are only updated weekly. The weekly update is to ensure the model does not overfit and also to account for the delay periods in epidemic processes. The process is repeated throughout the horizon (T) and the Maximum a Posteriori (MAP) estimation for β_t, ϕ_t, g_t are recorded at each timepoint. The complete list of all the parameter values can be found in Azad et al. (2022).

2.4 Factorial Analysis

Factorial analysis (Montgomery 2017) is a powerful approach to identify latent variables and gauge relationships between a set of observed variables. It can either be utilized as a means for identifying underlying structure or a tool for testing a hypothesized structure. In `PySIRTEM`, we apply factorial analysis to study the effects of various trends in transition parameters. Specifically, we analyze how different infection rate (β), testing rate (ϕ), and general sickness rate (g) impact daily case numbers. We explore both the main effects, i.e., how each factor individually influences outcome, and interaction effects, i.e., how the rates jointly impact the case trends.

For `PySIRTEM`, we conduct a full factorial experiment by varying the three rates in three distinct temporal shapes:

i. increasing: the rate gradually increases over time.
ii. decreasing: the rate gradually decreases over time.
iii. single-peak: the rate follows a bell-curve trajectory, peaking at an intermediate timepoint.

Additionally, each shape has different bounds denoting the highest and lowest value. Specifically, we denote the bounds in three classes: high, medium, and low, which result in a total of nine unique setups for each shape. The specific bounds for the rates are listed in Table 1.

Main Effects Analysis. Main effect in factorial design refers to the influence of a singular independent variable on the dependent variable, ignoring the effects of other independent variables. The main effect gives insight into whether changing the level of one factor affects the outcome variable or not. The main effect of a factor A is denoted by,

$$\alpha_i = \overline{R_i} - \overline{R},$$

Algorithm 1: Particle filter iteration overview.

Input: Initial Particles $x_0 = [x_0^1, x_0^2, \ldots, x_0^N]$
Observed cases $Y = [y_1, y_2, ...y_T]$ where $y_t = [y_t^+, y_t^-]$

Initialize MAP state list: S = []
Initialize MAP simulation list: $\hat{Y}$ = []
for $t = 1 \ldots T$ **do**
  if $t \mod 7 == 0$ **then**
    for $i = 1 \ldots |X_t|$ **do**
      $x_t^i \leftarrow$ `StateTransitionFn`(x_{t-1}^i, x_{t-8}^i);
  for $i = 1 \ldots |x_t|$ **do**
    $[\hat{y}_t^{+^i}, \hat{y}_t^{-^i}] \leftarrow$ `SimulationFnPySIRTEM`(x_t^i);
    Compute likelihood weight:

$$w_t^i \propto \exp\left(-\frac{(y_t^+ - \hat{y}_t^{+^i})^2}{2\sigma_{\text{obs}}^2} - \frac{(y_t^- - \hat{y}_t^{-^i})^2}{2\sigma_{\text{obs}}^2}\right)$$

  Normalize weights: $w_t^i \leftarrow \frac{w_t^i}{\sum_j w_t^j}$;
  $X_t \leftarrow$ `Resample`(x_t, w_t);
  Compute MAP estimates:

$$x_t^{MAP} = \arg\max_{x_t} p(x_t | y_{1..t})$$

$$\hat{y}_t^{MAP} = \arg\max_{y_t} p(y_t | x_t^{MAP})$$

  Append $\hat{y}_t^{MAP}$ to S;
  Append $\hat{\theta}_t$ to $\hat{Y}$;

Output: MAP state history S
MAP simulation history $\hat{Y}$

Table 1: Range of rates for different shapes and orders for factorial analysis.

Rate	**Shape**	**Order**	**Range**
Infection Rate (β)	Increasing, Decreasing, Single-peak	High	[0.05, 0.99]
		Medium	[0.05, 0.7]
		Low	[0.05, 0.5]
Testing Rate (ϕ)	Increasing, Decreasing, Single-peak	High	[0.05, 0.4]
		Medium	[0.05, 0.3]
		Low	[0.05, 0.2]
General Sickness Rate (g)	Increasing, Decreasing, Single-peak	High	[0.001, 0.05]
		Medium	[0.001, 0.03]
		Low	[0.001, 0.02]

where $\overline{R_i}$ is the mean response for level i of factor A and $\overline{R}$ is the grand mean i.e the average response of all levels. In our experiment, we denote the levels of each factor by mapping the trends to 1 and -1 if they are used in the simulation or not respectively. For example, if a simulation is run with high increasing β, high increasing ϕ, and low single-peak g, we map these 3 factors to level 1 and all the other factors to level -1.

Interaction Plots Analysis. Interaction plots provide a visual report on how the relationship between one factor and the outcome depends on the levels of another factor and reveal additional insight of combined interactions possibly overlooked by main effects analysis. The interaction between level i of factor A and level j of factor B is denoted by

$$(\alpha\beta)_{ij} = \overline{R_{ij}} - \overline{R_{i.}} - \overline{R_{.j}} + \overline{R},$$

where $\overline{R_{ij}}$ is the mean response for the combination of level i of A and level j of B, $\overline{R_{i.}}$ is the mean response for level i of A (ignoring B), $\overline{R_{.j}}$ is the mean response for level j of B (ignoring A), and $\overline{R}$ is the grand mean. A large positive or negative $(\alpha\beta)_{ij}$ value indicates strong interaction while $(\alpha\beta)_{ij} \approx 0$ means no interaction.

Parallel lines in the plot denote no interaction while non-parallel lines suggest one factor's effect depends on the level of another factor. Crossing lines indicate strong interaction where the direction of the effect reverses based on the level of the other factor. Strong interaction between two factors imply that the effect of the factors cannot be evaluated in isolation and main effect analysis should be interpreted with caution.

3 EXPERIMENTAL RESULTS

3.1 `PySIRTEM` Validation

We confirmed the validity and accuracy of the `PySIRTEM` parameters through particle filter approach by comparing the simulation results from the model against published confirmed case data from three U.S. states - Arizona, Florida, and Minnesota. `PySIRTEM` is implemented in `Python 3.12.9` using libraries such as `NumPy`, `SciPy`, `ddeint`, and `pfilter` and run on Sol Supercomputer at Arizona State University, an environment with 12 CPUs and 8 GiB memory. The particle filter algorithm is initialized with $N = 200$ particles and horizon $T = 343$ days, with (β), (ϕ), and (g) rates for each state calibrated weekly. Weights were updated based on the equation in line 8 of Algorithm 1 and importance sampling was employed to discard lower weighted particles. A complete list of `SIRTEM` simulation parameters and their sources are provided in Azad et al. (2022). `PySIRTEM` is then simulated with the calibrated rates to generate estimated daily cases and compared against reported confirmed cases. The confirmed daily positive (y^+) and daily negative (y^-) were obtained from public sources (CTW nd) between March 22, 2020 to March 1, 2021. Forecasts were made for a 7-day period after each weekly calibration of the rates.

Figures 2a and 2b show that `PySIRTEM` is quite accurate in capturing the trends of infection for Arizona. Figure 2c displays the intrinsic parametric trends which explain the daily pandemic progression i.e., the peak in β at week 10, 30, and 40 correspond with the following rise in positive cases. In Figure 2d, we compare the mean daily error in estimation for 20 iterations from `SIRTEM`'s AR(2) calibration method (Azad et al. 2022) and `PySIRTEM`'s particle filter approach, with `PySIRTEM` providing more accurate results. Runtime evaluations further display that PySIRTEM is approximately twice as fast as SIRTEM in learning and verifying optimal transition rates. All experiments were conducted in the same computational environment. Comprehensive error and runtime comparisons for all the states are listed in Table 2.

Table 2: Normalized mean-squared error and avg. runtime for selected U.S. States from `PySIRTEM` vs `SIRTEM` calibration results.

State	Population	Avg. Reported Positives	Avg. Reported Negatives	Normalized Pos. MSE	Normalized Neg. MSE	Mean Normalized MSE		Avg. Runtime (hrs)	
						`PySIRTEM`	`SIRTEM`	`PySIRTEM`	`SIRTEM`
Arizona	7.3 million	1774.8	7876.7	298.1	52.7	175.4	2226.38	9.11	15.8
Florida	21.5 million	5144.3	25434.5	136.8	868.8	502.8	3408.86	9.02	16.1
Minnesota	5.6 million	1387.9	9026.2	745.2	54.3	399.7	1667.72	9.15	16.2

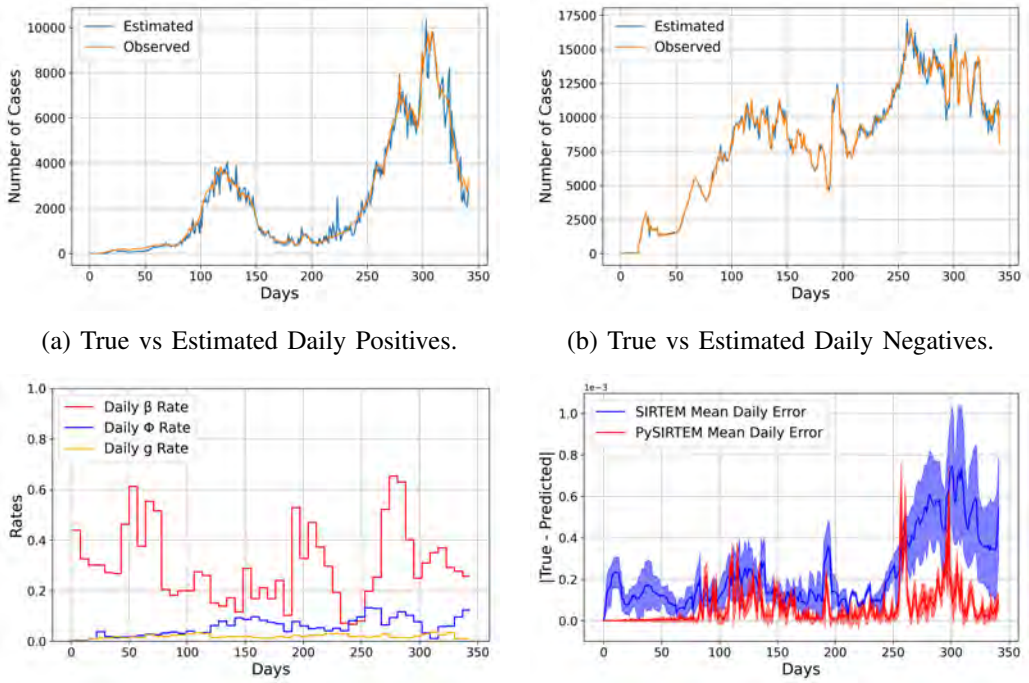

(a) True vs Estimated Daily Positives.

(b) True vs Estimated Daily Negatives.

(c) Weekly Calibrated Transition Rates.

(d) Mean Daily Prediction Error.

Figure 2: Model prediction results for the State of Arizona.

3.2 Factorial Analysis Outcomes

In this section, we analyze how different infection rate (β), testing rate (ϕ), and general sickness rate (g) impact daily case numbers; in particular, we aggregate the findings from factorial analysis using *interaction effects matrix plots* (NIST 2012) in Figure 3. The diagonal plots in the matrices denote the main effects while the off-diagonal plots indicate interaction. At each cellplot at row r and column c, the red line refers to relation between factor c and the peak value (maximum daily positive cases reported) when factor r is not used in the simulation and the blue line refers to the interaction between factor c and the peak value in the presence of factor r. For each of the experiments, we keep the infection rate trends constant (but vary the order or intensity i.e. high, medium, low), only changing the testing rate and general sickness rate trends to compare the pairwise interaction between the factors.

For space scarcity, factor names are abbreviated in Figure 3. For example, *i_incr_H* refers to high increasing infection rate trend. The values for each trend are again generated from Table 1. Since no two trends for the same factor are possible for an independent run of the simulation, interaction plots for same factors have been grayed out.

In the plots in Figure 3a, we see non-parallel lines between increasing β and both increasing and decreasing ϕ, implying strong levels of interaction among these. Specifically, when ϕ increases, more positive cases are detected, suggesting higher testing is effective in identifying infections (*t_incr_M*, *i_incr_M*). However, decreasing ϕ leads to under-reporting (*t_decr_H*, *i_incr_H*) as the lines point downward (lower positive cases) even when β is increasing. We can observe that a testing policy where majority of testing is concentrated on a single timepoint (*t_incr_M*) is less effective in containing spread, as implied by the parallel lines (*t_1p_M*, *i_incr_M*).

The plots in Figure 3a further show that with increasing β, non-COVID g rate has trivial impact on the peak of infection because the number of false positives due to g rate is still greatly outnumbered by true positives (*g_incr_M*, *i_incr_M*). The g rate, however, shows strong interaction with a single-peak testing

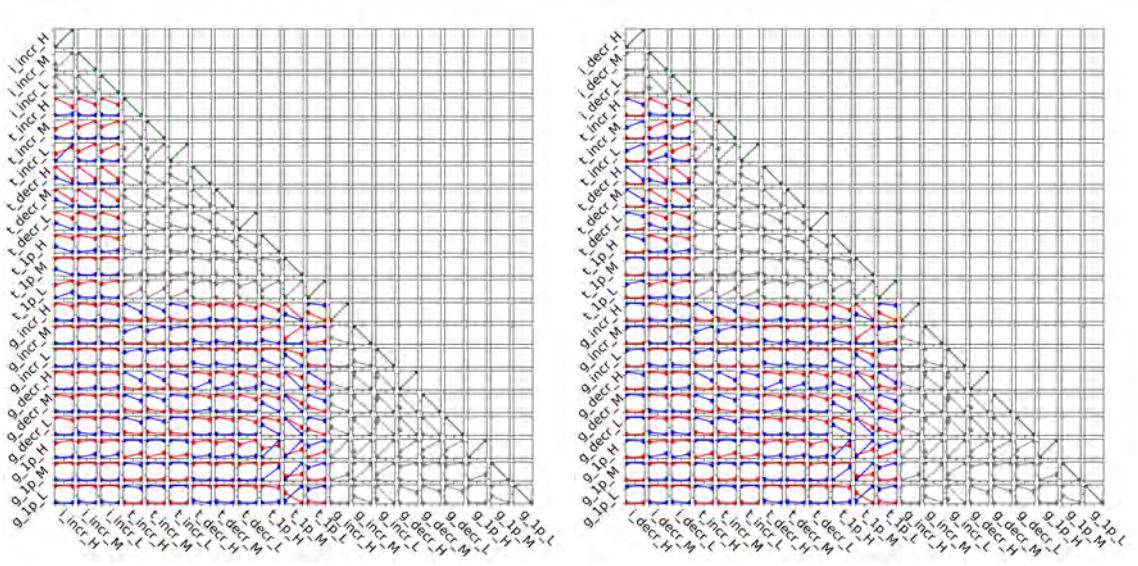

(a) Constant Infection Rate with Increasing Trend (b) Constant Infection Rate with Decreasing Trend

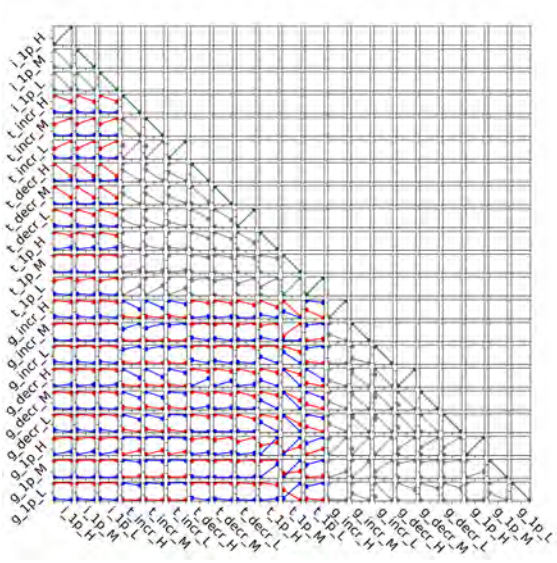

(c) Constant Infection Rate with Single-Peak Trend

Figure 3: Factorial Analysis for different trends in `PySIRTEM` transition parameters (Outcome variable: maximum daily positive cases reported over horizon T). Factors in x and y-axis are abbreviated as follows: i, t, g for infection, testing, and general sickness rate respectively; *incr*, *decr*, and $1p$ for increasing, decreasing, and single-peak trends; H, M, and L for the rate bounds (high, medium, low).

rate policy. Particularly, when a high number of individuals are sick with non-COVID illness, high testing rate may lead to more false positives stimulating higher positive reporting (g_1p_H, t_1p_H). This is evident from the blue line having a steeper positive slope than the red even when β trend is constant in both cases.

Figure 3b displays that if β is decreasing rapidly, a high test rate can quickly flatten the infection curve resulting in lower positive cases (t_incr_H, i_decr_H). However, a low increasing ϕ may not be

sufficient to curb a slowly decreasing infection trend (*t_incr_L*, *i_decr_H*), stressing the significance of adaptive testing policies. Here the blue line displays positive slope implying insufficient testing which leads to higher spread and higher reporting of positive cases.

Figure 3c reveals that for bell-curve β trend, a similar single-peak testing policy is mostly ineffective as indicated by parallel interaction plots for (t_1p_M, i_1p_M). However, increasing and decreasing test rates are capable of containing positive cases. Moreover, plot for (*t_decr_M*, i_1p_M) shows that decreasing ϕ (i.e. strong testing at the start) is more effective in keeping the reported cases low (the downward slope suggests lower spread and reporting) than increasing ϕ (*t_incr_M*, i_1p_M) where we see a positive slope which demonstrates that testing policy failed to regulate infection.

3.3 Containment Scenario Analysis

In order to demonstrate the capabilities of `PySIRTEM` as a simulation tool, we conduct a series of experiments comparing results generated from parameter setups displaying low and strong interaction derived from factor analysis. We gauge the rate of impact when different rates change their trends. The outcomes are compared against the simulation results obtained from the setup where the rates display only one trend i.e. increasing or decreasing. Finally, we evaluate the rate of change in infection and positive cases in order for better policy-making.

We run the simulation for $N = 150$ days and the trends are kept constant after day 100 to observe the effect of the trends. At $n = 75$, we toggle the trends of one of the rates. For example, testing rate (ϕ) follows a decreasing trend up to day 75 and then it gradually starts increasing in Figure 4a. In each of the experiments, we keep the other rates' trends unchanged i.e., the infection (β) and general sickness rate (g) follow an increasing trend throughout. Similarly, when β changes trends, ϕ and g are kept as increasing.

The positive (y^{+*}) and infected (I^*) case numbers are compared against data obtained from `PySIRTEM` where the rates have unchanged trends i.e β always increasing and ϕ always decreasing. The rates of change are calculated as follows,

$$\Delta y_t^+ = \frac{y_t^{+*} - y_t^+}{y_t^+},$$
$$\Delta I_t = \frac{I_t^* - I_t}{I_t}.$$

As we see in Figure 4a, when ϕ starts to increase (after the vertical red dashed line), there is a quick surge in reported cases; however the infected numbers gradually decrease and the higher testing not only limits the peak of infection but also delays it. On the other hand, Figure 4b shows that, when β suddenly starts to increase, the infected hill is wider even though the testing is able to identify the peak; this suggests that infection is lasting and stronger contingency plans are required. These experiments, along with others involving various hyperparameters, give valuable insights in managing resources and policy making during different scenarios during the course of a pandemic.

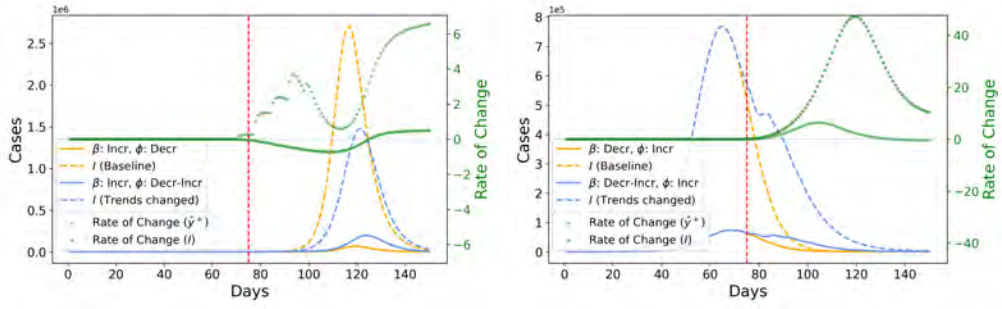

(a) Test Rate ϕ (Decreasing to Increasing). (b) Inf Rate β (Decreasing to Increasing).

Figure 4: Comparison of Model prediction results for two different scenarios.

4 CONCLUSION

`PySIRTEM` provides a powerful and flexible tool for simulating pandemic dynamics via the addition of a particle filter for dynamic estimation of parameters. Our factorial design analysis showcases that `PySIRTEM` effectively captures critical factor interactions, including infection and test rates, and facilitates informed decision-making. The model is versatile in simulating various real world pandemic scenarios. Future efforts could see us explore additional variables like vaccination techniques, mutated viral strains, and demographic traits in the context of `PySIRTEM`. To further enhance granularity, an agent-based model (ABM) can supplement `PySIRTEM`'s high-level dynamics by simulating individual behavior and interaction.

5 ACKNOWLEDGEMENTS

This work was partially supported by NSF CNS 2125246 and NSF DBI 2412115.

REFERENCES

Anderson, R. M., H. Heesterbeek, D. Klinkenberg, and T. D. Hollingsworth. 2020. "How will Country-based Mitigation Measures Influence The Course of The COVID-19 Epidemic?". *The Lancet* 395:931–934.

Arulampalam, M., S. Maskell, N. Gordon, and T. Clapp. 2002. "A Tutorial on Particle Filters for Online Nonlinear/Non-Gaussian Bayesian Tracking". *IEEE Transactions on Signal Processing* 50(2):174–188.

Azad, F. T., R. W. Dodge, A. M. Varghese, J. Lee, G. Pedrielli, K. S. Candan *et al*. 2022. "SIRTEM: Spatially Informed Rapid Testing for Epidemic Modeling and Response to COVID-19". *ACM Transactions on Spatial Algorithms and Systems* 8(4):1–43.

Bjørnstad, O. N., K. Shea, M. Krzywinski, and N. Altman. 2021. "Author Correction: The SEIRS Model for Infectious Disease Dynamics". *Nature Methods* 18:321.

Chang, S., E. Pierson, P. W. Koh, J. Gerardin, B. Redbird, D. Grusky *et al*. 2021. "Mobility Network Models of COVID-19 Explain Inequities and Inform Reopening". *Nature* 589(7840):82–87.

Chu, D. K., E. A. Akl, S. Duda, K. Solo, S. Yaacoub, H. J. Schünemann, *et al*. 2020. "Physical Distancing, Face Masks, and Eye Protection to Prevent Person-to-Person Transmission of SARS-CoV-2 and COVID-19: A Systematic Review and Meta-Analysis". *The Lancet* 395(10242):1973–1987.

CTW n.d.. "The Covid Tracking Project". https://covidtracking.com/data. accessed January 7, 2024.

Doucet, A., and A. M. Johansen. 2009. "A Tutorial on Particle Filtering and Smoothing: Fifteen Years Later". *Handbook of Nonlinear Filtering* 12(656-704):3.

Duan, W., Z. Fan, P. Zhang, G. Guo, and X. Qiu. 2015. "Mathematical and Computational Approaches to Epidemic Modeling: A Comprehensive Review". *Frontiers of Computer Science* 9:806–826.

He, S., Y. Peng, and K. Sun. 2020. "SEIR Modeling of the COVID-19 and its Dynamics". *Nonlinear Dynamics* 101:1667–1680.

Hossain, M., M. Hassanzadeganroudsari, and V. Apostolopoulos. 2021. "The Emergence of New Strains of SARS-CoV-2. What Does It Mean for COVID-19 Vaccines?". *Expert Review of Vaccines* 20(6):635–638.

Jit, M., and M. Brisson. 2011. "Modelling the Epidemiology of Infectious Diseases for Decision Analysis". *PharmacoEconomics* 29:371–386.

Kerr, C. C., R. M. Stuart, D. Mistry, R. G. Abeysuriya, K. Rosenfeld, G. R. Hart, *et al*. 2021. "Covasim: An Agent-Based Model of COVID-19 Dynamics and Interventions". *PLOS Computational Biology* 17(7):1–32.

Künsch, H. R. 2013. "Particle filters". *Bernoulli* 19(4):1391–1403.

Lau, H., T. Khosrawipour, P. Kocbach, H. Ichii, J. Bania, and V. Khosrawipour. 2021. "Evaluating the Massive Underreporting and Undertesting of COVID-19 Cases in Multiple Global Epicenters". *Pulmonology* 27(2):110–115.

Li, M. Y., and J. S. Muldowney. 1995. "Global Stability for The SEIR Model in Epidemiology". *Mathematical Biosciences* 125:155–164.

Maged, A., A. Ahmed, S. Haridy, A. W. Baker, and M. Xie. 2023. "SEIR Model to Address the Impact of Face Masks Amid COVID-19 Pandemic". *Risk Analysis* 43(1):129–143.

Miksch, F., P. Pichler, K. J. Espinosa, K. S. Casera, A. N. Navarro, and M. Bicher. 2015. "An Agent-Based Epidemic Model for Dengue Simulation in The Philippines". In *2015 Winter Simulation Conference (WSC)*, 3202–3203. IEEE https://doi.org/10.1109/WSC.2015.7408470.

Montgomery, D. C. 2017. *Design and Analysis of Experiments*. John Wiley & Sons.

Niehus, R., P. M. De Salazar, A. R. Taylor, and M. Lipsitch. 2020. "Using Observational Data to Quantify Bias of Traveller-Derived COVID-19 Prevalence Estimates in Wuhan, China". *The Lancet Infectious Diseases* 20(7):803–808.

NIST 2012. "NIST/SEMATECH e-Handbook of Statistical Methods, Section 5.4.6.1: Interaction Effects Matrix Plot". https://www.itl.nist.gov/div898/handbook/pri/section5/pri594.htm. accessed June 1, 2024.

Omori, R., K. Mizumoto, and G. Chowell. 2020. "Changes in Testing Rates Could Mask The Novel Coronavirus Disease (COVID-19) Growth Rate". *International Journal of Infectious Diseases* 94:4.

Papageorgiou, V. E., and G. Tsaklidis. 2024. "A Stochastic Particle Extended SEIRS Model with Repeated Vaccination: Application to Real Data of COVID-19 in Italy". *Mathematical Methods in the Applied Sciences* 47(7):6504–6538.

Rhee, A. J. 2006. *An Agent-Based Approach to HIV/AIDS Epidemic Modeling: A Case Study of Papua New Guinea*. Ph. D. thesis, Massachusetts Institute of Technology. http://hdl.handle.net/1721.1/34528.

Roy, S., P. Biswas, and P. Ghosh. 2021a. "Effectiveness of Network Interdiction Strategies to Limit Contagion During a Pandemic". *IEEE Access* 9:95862–95871.

Roy, S., P. Biswas, and P. Ghosh. 2021b. "Quantifying Mobility and Mixing Propensity in the Spatiotemporal Context of a Pandemic Spread". *IEEE Transactions on Emerging Topics in Computational Intelligence* 5(3):321–331.

Roy, S., R. Dutta, and P. Ghosh. 2021. "Optimal Time-Varying Vaccine Allocation Amid Pandemics With Uncertain Immunity Ratios". *IEEE Access* 9:15110–15121.

Walensky, R. P., and C. Del Rio. 2020. "From Mitigation to Containment of the COVID-19 Pandemic: Putting the SARS-CoV-2 Genie Back in the Bottle". *Jama* 323(19):1889–1890.

Wells, C. R., J. P. Townsend, A. Pandey, S. M. Moghadas, G. Krieger, B. Singer, *et al*. 2021. "Optimal COVID-19 Quarantine and Testing Strategies". *Nature Communications* 12(1):356.

AUTHOR BIOGRAPHIES

PREETOM BISWAS is a Computer Science student in the accelerated 4+1 B.S./M.S. program at Arizona State University for the School of Computing and Augmented Intelligence (SCAI) working with Dr. Giulia Pedrielli and Dr. K.Selçuk Candan on Pandemics modeling. His research focuses on stochastic simulation and optimization within with epidemiology applications. His email address is pbiswa11@asu.edu.

GIULIA PEDRIELLI is currently an Associate professor for the School of Computing and Augmented Intelligence (SCAI) at Arizona State University. She develops her research in design and analysis of random algorithms for global optimization, with focus on improving finite time performance and scalability of these approaches. Applications of her work are in individualized cancer care, bio-manufacturing, design and control of self-assembled RNA structures, verification of cyber-physical systems. Her email address is giulia.pedrielli@asu.edu and her website is https://www.gpedriel.com/.

K. SELÇUK CANDAN is a professor of computer science and engineering at Arizona State University (ASU) and the director of ASU's Center for Assured and Scalable Data Engineering (CASCADE). His primary research interest is in the area of management and analysis of non-traditional, heterogeneous, and imprecise (such as multimedia, web, and scientific) data, with applications in public heath and sustainability, among others. He is an ACM Distinguished Scientist. His email address is candan@asu.edu and his website is https://kscandan.site.