

INFERRING EPIDEMIC DYNAMICS USING GAUSSIAN PROCESS EMULATION OF AGENT-BASED SIMULATIONS

Abdulrahman A. Ahmed
M. Amin Rahimian

Mark S. Roberts

Department of Industrial Engineering
University of Pittsburgh
3700 O'Hara Street
Pittsburgh, PA 15261, USA

Department of Health Policy and Management
University of Pittsburgh
130 Desoto Street
Pittsburgh, PA 15261, USA

ABSTRACT

Computational models help decision makers understand epidemic dynamics to optimize public health interventions. Agent-based simulation of disease spread in synthetic populations allows us to compare and contrast different effects across identical populations or to investigate the effect of interventions keeping every other factor constant between “digital twins.” FRED (A Framework for Reconstructing Epidemiological Dynamics) is an agent-based modeling system with a geo-spatial perspective using a synthetic population that is constructed based on the U.S. Census data. In this paper, we show how Gaussian process regression can be used on FRED-synthesized data to infer the differing spatial dispersion of the epidemic dynamics for two disease conditions that start from the same initial conditions and spread among identical populations. Our results showcase the utility of agent-based simulation frameworks such as FRED for inferring differences between conditions where controlling for all confounding factors for such comparisons is next to impossible without synthetic data.

1 INTRODUCTION

Agent-based modeling (ABM) is a helpful tool in studying epidemic dynamics. Cajka et al. (2010) create a synthetic population and develop an ABM where they assign the agents in the synthetic population to households, workplaces, schools, and other assignments. Ferguson et al. (2006) develops a large-scale simulation to study different interventions for a novel epidemic outbreak. Chao et al. (2010) develops a simulation for the spread of influenza in a large population, the model is calibrated with historical influenza data. Similarly, Tuite et al. (2010) develops a Markov chain Monte-Carlo simulation model based on historical data of symptoms that are obtained for laboratory-confirmed cases of H1N1 influenza. Parker and Epstein (2011) develops a simulation model that can simulate epidemic outbreaks on a global scale (i.e., several billion agents). Grefenstette et al. (2013) build on these concepts and develop them into a generic framework for epidemic dynamics modeling called FRED (A Framework for Reconstructing Epidemiological Dynamics). Lukens et al. (2014) develop an ABM linked with an equation-based, within-host model for influenza in FRED. Potter et al. (2012) develop an ABM to study local school closure policies during an influenza epidemic. Liu et al. (2015) build a model of Measles transmission using FRED where they study the role of vaccination in preventing Measles outbreaks; this paper has been influential in stressing the importance of vaccination and impacting public policy. Beyond influenza and measles models, Krauland et al. (2020) study cardiovascular disease death and its social determinants using the FRED synthetic population. The latter study demonstrates the possibility of understanding disease risk and consequences of large-scale interventions using FRED synthetic population.

The computational cost of population-scale, agent-based simulations could be high, and modelers resort to inferential tools to estimate surrogate models that summarize relevant information in the simulations and circumvent having to rely on simulation samples for upstream inference and optimization tasks. Gaussian process regression (GPR) is one such inferential tool. Fearnhead et al. (2014) use Gaussian processes in inferring reaction rates of networks which is usually a type of network used in applications of epidemics and biology. Senanayake et al. (2016) and Zimmer and Yaesoubi (2020) develop a GPR model to understand the dynamics of influenza. Zimmer et al. (2017) use Gaussian processes to calibrate and estimate epidemic parameters and point out advantages over other comparable methods. Ball and House (2017) considered inference for a SIR epidemic model on a network using GPR based on the covariance function of a branching process. Buckingham-Jeffery et al. (2018) use GPR for SIR (Susceptible-Infected-Recovered) and SEIR (Susceptible-Exposed-Infected-Recovered) models to estimate underlying parameters.

We show an application of the Gaussian process regression to FRED-synthesized data from two epidemic models, an influenza model (INF) and an opioid use disorder mode (OUD). Subsequently, we can compare the spatial autocorrelation values for the two diseases and compare their spatial dependencies, keeping everything the same across the two synthetic populations. Both models were constructed over the same geographic location with the same agent's information in the area and starting from the same initial conditions. Thus, we can attribute the observed difference in spatial clustering to the differing disease dynamics, because the remaining conditions are the same for both models using the same synthetic population.

This paper is structured into six sections. In Section 2, we briefly introduce the FRED simulation software and detail how FRED works. In Section 3, we provide preliminaries about GPR and introduce our notation for GPR with spatial, FRED-synthesized data. In Section 4, we discuss our modeling assumptions inference and methodology. In Section 5, we discuss the results of our study and its public health implications. Finally, in Section 6, we provide concluding remarks and give future work directions.

2 FRED SIMULATION SOFTWARE

FRED (Framework for Reconstructing Epidemiological Dynamics) is an agent-based, open-source software that is developed to simulate the temporal and spatial behaviors of epidemics. FRED is developed by the Public Health Dynamics Laboratory (PHDL) at the University of Pittsburgh School of Public Health. FRED was developed originally to study epidemic dynamics, however, it has shown broader potential for population studies that help guide public health policies and interventions. One of the key features of FRED is its synthetic population that is based on the US Census Bureau's public use Microdata files and Census aggregated data, which we discuss below (Guclu et al. 2016).

2.1 Synthetic Population

FRED represents every individual in a specified geographic area explicitly. FRED utilizes the US synthetic population database from RTI International (Wheaton 2012), where the synthetic population contains detailed geographically-allocated categories. The synthetic population has geographically-assigned, synthetic households. Each household has residents, and there are also residents for different facilities (e.g., dorms, prisons, etc.), schools, and mapping of students to schools. Each household, school, and workplace is assigned to a specific geographic region. These assignments reflect an implicit representation of the distance between agents and facilities. Each agent has its own assigned socioeconomic and demographic information (e.g., age, sex, race, employment, etc.) and assigned locations to their activities (e.g., school, workplace, household, etc.) (Grefenstette et al. 2013).

2.2 Discrete-Time Simulation

At every simulation step, the agent interacts with other agents likely to share similar daily activities. For example, agents in the same workplace interact with the same colleagues daily. If an infected agent interacts

with a susceptible agent, there is a chance for disease transmission from the infected agent to the susceptible one. Additionally, the agent can dynamically change their daily activities (e.g., stop going to work, change locations, etc.). Moreover, FRED allows detailed and effective evaluation of specific interventions as every disease transmission is recorded in the software logs.

2.3 Agent Model

Each agent has its own attributes (e.g., age, race, gender, employment, school, etc.) and activities' locations (e.g., household, school, workplace, etc.). Newborns are assigned to the same household as their mothers; if children reach school age, they are assigned to a school, and if an agent dies, it is erased from the synthetic population.

FRED allows agents to take actions especially in contexts related to public health. For example, agents can decide to stay home if sick, prevent their children from attending school if sick, or accept or reject vaccinations. At each step, the agent's actions involve interactions between the agent's internal actions and factors of their environment (e.g., the availability of vaccines given their decision to take it).

2.4 Disease Model

FRED supports the propagation of one or more infectious diseases in the population. Each disease is determined by specific parameters for transmission, contact, history, etc. For example, typically the agent will pass into three stages for infectious disease, Susceptible, Infected, Recovered, or S-I-R. The agent will move from each state to another based on the transmission rate (defined by the modeler) and also based on each agent's data (e.g., if the agent is registered in a school where the disease is spreading, then the agent will have an increased infection rate). Epidemics in FRED are propagated by assigning random agents in the population to have the infectious disease as an initialization. It is also possible to modify these assignments based on specified groups (e.g., related to race, geography, age, etc.).

2.5 Place Model

FRED supposes that all disease-specific interactions between agents occur in a "place," where each type of place specifies a unique environment for spreading infection. The default type of places in FRED include: schools, workplaces, households, etc. Agents are expected to do their daily activities in proximity to assigned locations. However, agents can move to another location during the day. Based on its specific information each agent has specific default daily schedule (e.g., times to go to the workplace, weekends, or annual holiday schedules). Places can also be location-free, allowing us to model social network interactions in FRED.

3 GAUSSIAN PROCESS REGRESSION

A Gaussian Process (GP) is a type of stochastic process where the value of any random variable within the process follows Gaussian distribution. It can be seen as a generalization of the Gaussian probability distribution (hence the name). Each realization of a Gaussian process is a function $f(\cdot)$, and one way to look at GP is to think of its realizations as long vectors taking values $f(x)$ at specified inputs x_1, x_2, \dots . In this view, a good GP fit contains information about infinite queries and allows us to make inferences about any finite collection of points (Williams and Rasmussen 2006). Mathematically, a GP can be specified as follows:

$$f(x) \sim GP(\mu(x), k(x, x')) \quad (1)$$

where $\mu(x)$ is the expected value given input location x and $k(x, x')$ is the kernel function of the GP (also known as the covariance function) (Duvenaud 2014). Starting from prior parameters and given the observed data, Gaussian process regression (GPR) provides a Bayesian procedure to update GP kernel parameters based on their fit to the observed data. The Bayesian procedure consists of a sequence of iterative steps

where a loss function is calculated at every step, and based on the result the model's parameters are updated, until a predefined criterion is met and the procedure stops.

3.1 Kernels

A kernel (also known as covariance function, kernel function, or covariance kernel) is a function that defines the similarity between two points x, x' where x and x' are typically vectors in the Euclidean space, but can also be defined for images, graphs, or categorical inputs.

3.1.1 Kernel Types

There are different types of kernels due to the variety of applications for GP. We will mention kernels that are widely used in general including the one that we use in our model.

Radial Basis Function Kernel Radial basis function (RBF) kernel (also known as the squared-exponential kernel) is defined as

$$k(x, x') = \exp\left(-\frac{d(x, x')^2}{2l^2}\right),$$

where l is the length scale parameter and $d(\cdot, \cdot)$ is the Euclidean distance between the input points. As an advantage, RBF kernel is infinitely differentiable, implying that GP with the corresponding kernel has mean square derivatives and is very smooth. However, RBF has limitations in modeling some applications, especially physical processes where other kernels are recommended. Nevertheless, RBF kernel is the most widely used kernel in the field (Williams and Rasmussen 2006).

Matern Kernel The Matern class of kernels is a generalization of RBF kernels. The kernel is

$$k(x, x') = \frac{1}{\Gamma(\nu)2^{\nu-1}} \left(\frac{\sqrt{2\nu}}{l}d(x, x')\right)^\nu K_\nu\left(\frac{\sqrt{2\nu}}{l}d(x, x')\right),$$

where $d(\cdot, \cdot)$ is the Euclidean distance between the two points, $K_\nu(\cdot)$ is a modified Bessel function and $\Gamma(\cdot)$ is the gamma function. The parameter ν controls the smoothness of the kernel, where smoothness increases with the value of ν . The kernel becomes especially simple when ν is a half-integer: $\nu = p + 1/2$, where p is a non-negative integer. For example when $\nu = 3/2$ the kernel becomes

$$k_{\nu=3/2}(x, x') = \left(1 + \frac{\sqrt{3} * d(x, x')}{l}\right) \exp\left(-\frac{\sqrt{3} * d(x, x')}{l}\right).$$

Moreover, as $\nu \rightarrow \infty$, the Matern kernel converges to the RBF kernel.

3.1.2 Combining Kernels

In some applications, a single kernel may not be expressive enough, and custom kernel functions are developed. Usually developing this type of custom kernel is by combining already known kernels to get the desired properties, for example, by adding two kernels as follows:

$$k_a + k_b = k_a(x, x') + k_b(x, x'),$$

or by multiplying them together:

$$k_a * k_b = k_a(x, x') \cdot k_b(x, x').$$

It is worth mentioning that the above examples are in one dimension, but creating custom kernels can be done also over multi-dimensional inputs by adding or multiplying different kernels over different dimensions (Duvenaud 2014).

3.2 Spatial Autocorrelation

Spatial autocorrelation measures the degree of correlation between the observed values for spatial units. Moran's I is a common measure for evaluating the level of spatial autocorrelation (Moran 1950). Given a set of features, Moran's I statistic evaluates whether the observed pattern shows significant clustering or if it can be adequately described as random, with values close to 1 or -1 indicating that the location data is clustered positively (occurring more closely to each other than would be expected from a random spread) or negatively (spread out more than expected from a random distribution of locations). A Moran's I value around zero indicates that there is no significant evidence for clustering and the spatial spread of events' locations can be adequately described as random. Moran's I is given by

$$I = \frac{N \sum_{i=1}^N \sum_{j=1}^N w_{ij} (x_i - \bar{x})(x_j - \bar{x})}{W \sum_{i=1}^N (x_i - \bar{x})^2}, \quad (2)$$

where N is the number of spatial units indexed by i and j , x is the variable of the pattern (e.g., disease incidence or prevalence at different locations), $\bar{x} = (\sum_{i=1}^N x_i)/N$ is the mean of the variable x , w_{ij} represents the spatial weight between i and j , and $W = \sum_{i=1}^N \sum_{j=1}^N w_{ij}$ is the summation of all w_{ij} values that constitute the spatial weight matrix defined next.

Spatial Weight Matrix The spatial weight matrix, $(W_{ij})_{N \times N}$, is a representation of the spatial structure over the data. Binary and variable weighting are two main ways of specifying the spatial weights between locations. Examples of binary weights include: contiguity-based, fixed distance, K nearest neighbours, etc. For variable weighting inverse distance is one of the main techniques. In binary weighting, contiguity means that two spatial units share a common boundary (with von Neumann and Moore being the two main ways for determining boundaries). Moore defines neighbours by the existing edge between the two spatial units, which usually defines four neighbours while von Neumann defines neighbours by the existing edge or vertices between the spatial units, which usually ends up with more neighbours compared to the Moore method. In this paper, the calculation of the weights is based on the distance between ZIP code centers.

4 METHODOLOGY

4.1 Simulated Models

In this section we will discuss the two simulation models used in our study, both models were built using FRED. The first model is a little simpler than the second one, however, it will help illustrate our points.

4.1.1 Influenza Model

Influenza (INF) model is one of the well-studied models provided with the installation of FRED simulation. INF model has five defined states as shown in Figure 1: Susceptible, Exposed, Infected with Symptoms, Infected Asymptomatic, and Recovery. The transition of the agent from state to state is based on explicitly defined probabilities, moreover, the agent's residing time in each state is explicitly defined for each state. We simulated the INF model over Allegheny County, Pennsylvania. The INF model can be found on FRED library on the following link: <https://github.com/PublicHealthDynamicsLab/FRED/tree/FRED-v5.0.0/library/Influenza>.

4.1.2 Opioid Use Disorder Model

The Opioid Use Disorder (OUD) model is a model developed to understand the OUD epidemic across the U.S. The growth of substance use disorder and drug overdose deaths is a growing public health concern in the U.S. Currently, opioids are the leading cause of drug overdose deaths in the U.S. (CDC 2023) (including prescription opioids, heroin, and synthetic opioids). Jalal et al. (2018) studies the epidemic dynamics over the last 40 years and concluded that the current opioid overdose deaths wave is part of a long-trend that




Figure 1: State transition diagram for the INF model.

is sustained over several decades stressing the importance of studying the epidemic dynamics. The OUD model that we use in this paper was developed in the Public Health Dynamics Laboratory at the University of Pittsburgh, based on data provided by Centers for Disease Control and Prevention (CDC) as a part of their sponsored research project. Similar to the INF model, OUD model has explicitly defined transition probabilities between different states and dwell times for each agent at different states. (See Figure 2.) Nevertheless, OUD model differs from INF in that its simulation update steps occur every month. Similar to the INF model and to have standard comparisons, the OUD model was simulated over the synthetic population of Allegheny County, PA. Both simulations were conducted over the same time frame between January 1, 2016 and December 31, 2017. The reason for this specific time frame is that OUD model state transitions were calibrated to the real data over the same period.




Figure 2: State transition diagram for the OUD model.

4.2 FRED Simulation Results

The INF and OUD models resulted in detailed records for each specified state over the simulated time window. We picked a single state from each simulation model to compare their behaviors. In the INF model, we picked the Infected with Symptoms (IS) state numbers over the simulated time window, while for the OUD model, we picked the Overdose Deaths (ODD) state over the same time window. Therefore, the behaviors we were comparing are the number of Infected with Symptoms numbers monthly from INF model against the number of ODD monthly from the OUD model. We aimed to use an inferential technique to measure statistical differences between the spatial spread of the epidemic dynamics in each case (i.e., spatial spread of the number of IS in the INF model and the number of ODD in the OUD model). In this study, we used GP as an inference model for our target parameters, where we select a single month (simulation step) to study the spatial spread. We avoided the starting and ending months in the two-year simulation period. For choosing a kernel, we tested different kernels to better infer the target parameters. To

decide which kernel to use we tested selected kernels on a separate set (i.e., test-set) and selected the kernel that has the lowest means square error (MSE) for each model, Table 1 shows MSE for different kernels RBF, Matern, and a combined kernel for the INF and OUD models. The training was done using Python language (<https://github.com/abdulrahmanfci/gpr-abm>) and the GPFlow library Matthews et al. (2017) for GP regression.

Model	RBF	Matern	RBF*Matern
INF	48.2	51.06	51.9
OUD	0.77	0.59	0.592

Table 1: Showing the MSE for different kernels tested for the two models (INF and OUD).

5 RESULTS AND DISCUSSION

The FRED-synthesized data are the list of coordinates where the IS and ODD happened at each simulation step. To better visualize the data we used the map of Allegheny county, PA ZIP codes to assign each incident location to its ZIP code. For example, if the first month of the INF model had two IS incidents happening at coordinates (x_1, y_1) and (x_2, y_2) within the 15261 ZIP code, then we assign a count of two to that ZIP code area. Afterward, we fitted a GP to the processed location data from the FRED output, taking the input, x , as Allegheny County's, PA ZIP code areas and the response y as the number of events in each ZIP code area. Figure 3 shows the simulated data for the INF model over the Allegheny County, PA ZIP codes map with the corresponding GP fitted values in part (b). The values shown in (b) are the mean of the posterior distribution. Because the INF model has relatively abundant numbers of IS over the region, a per-made RBF kernel was capable of modeling the INF disease's behavior. Figure 4 (a) shows a sample from the posterior distribution plotted over the longitude and latitude coordinates.

Similarly for the OUD model the FRED-synthesized data are the list of coordinates representing death events at each simulation step. We applied the same post-processing to map OD death locations to their ZIP code areas. In the case of the OUD model, we tested different GP kernels and combinations to improve the GP fit. The Matern kernel provided the best fit to the simulated data. In contrast to the INF model, the OUD model incident data are sparser and more spread out over the region. Figure 5 shows the simulated data for the OUD model for Allegheny County, PA and the corresponding GP fitted values in part (b), where the values shown in (b) are the mean of the GP posterior distribution. We can see from Figure 5 that the fitted values from GPR captured the center of the spread, however, they failed to pin down the location of the surge. This could be due to the sparsity of the ODD incident locations which makes tuning the parameters of the GP for the OUD model harder than for the INF model. Similarly, Figure 4 (b) shows a sample from the posterior distribution plotted over the longitude and latitude coordinates.

Spatial Autocorrelation Analysis To demonstrate the utility of GPR for inferring properties of the underlying epidemic dynamics, we conducted a spatial autocorrelation analysis using our GPR results. We used Moran's I statistics (equation (2)) to compare our two fitted GPs. The mean values for Moran's I statistics over 200 random draws for INF was 0.4 and for OUD was 0.68. We conducted a difference in means test where we assumed Gaussian distribution for population means. This resulted in a 95% confidence interval of $\mu_{INF} - \mu_{OUD} \in [-0.30, -0.27]$ for the difference in means. Figure 6 shows the distribution of Moran's I statistics for each model. The results indicate that incident locations in the INF model are less clustered than the OUD model. This could be verified by inspecting Figures 3 and 5 where occurrences of ZIP codes with high incident rates are fewer in the former case. Moreover, the results for OUD indicate more spatial correlation between adjacent ZIP codes compared to the INF model which does not show a strong relationship between incident counts in neighboring ZIP codes. The significant difference between spatial clustering may be attributed to strong socioeconomic correlates that derive the OUD spread, whereas INF is primarily driven by the spatial spread of the viral disease agents.

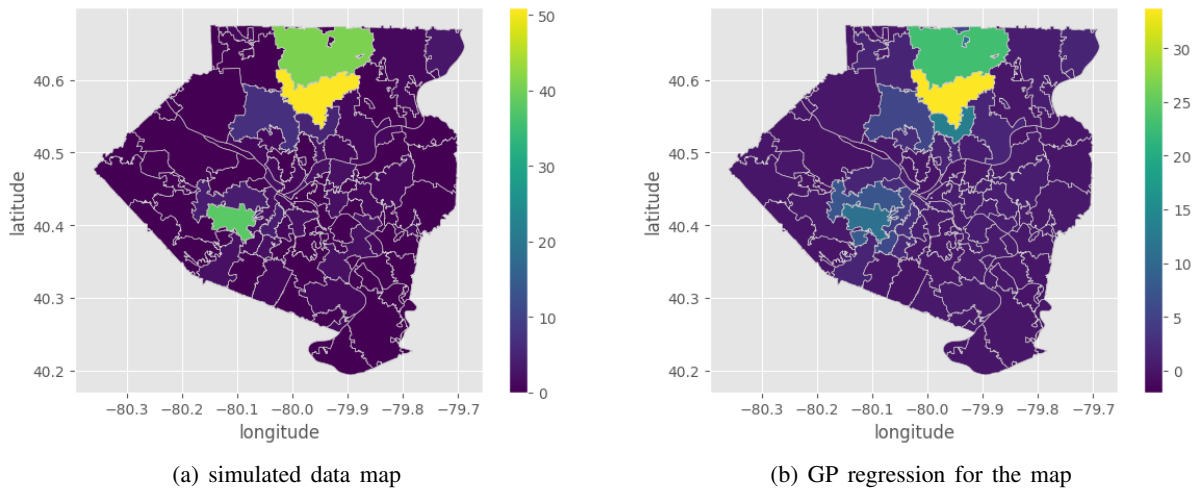


Figure 3: INF model simulated data versus GP regression for Allegheny county, PA.

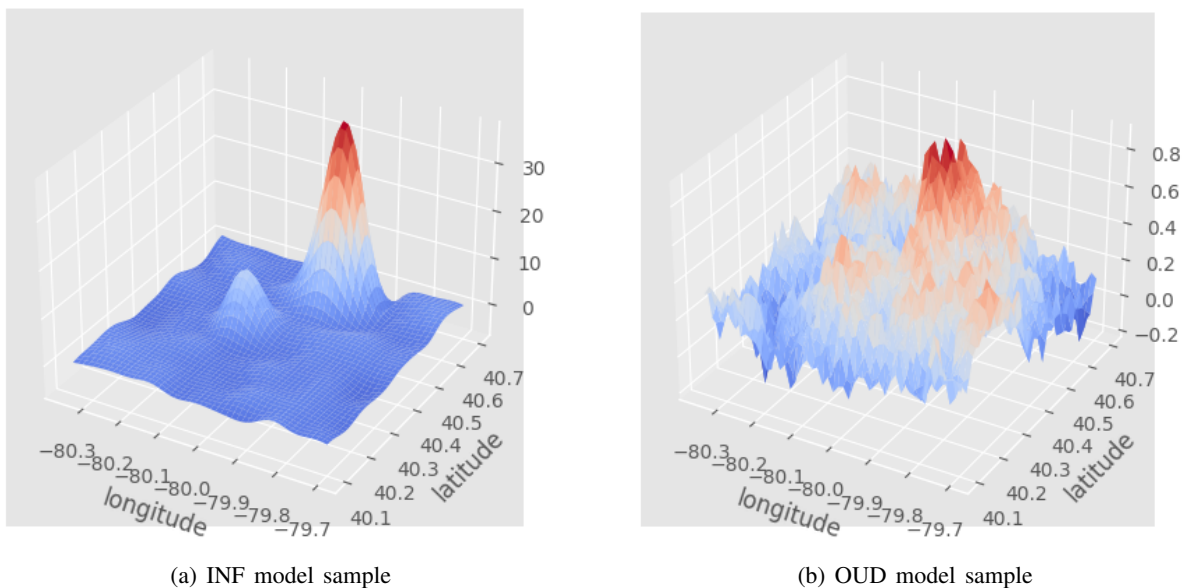


Figure 4: Posterior samples from GPR for different values of x and y in INF and OUD models, respectively.

This result shows the potential of GPR to be used in inferring epidemic dynamics. Such capabilities are especially useful when simulating many interventions across large populations where GPR can be used to interpolate across population characteristics and treatment conditions, thus circumventing the need to simulate each condition across the entire population (Ahmed et al. 2023). In this part, we used functions from PySAL library (Rey and Anselin 2007) — PySAL is a library specialized in providing geospatial tools and applications for spatial analysis in Python.

6 CONCLUSION

Because the INF and OUD models were tested on the same synthetic population over the same time window and geographic location, the observed differences between the fitted parameters of GP, and in particular, the

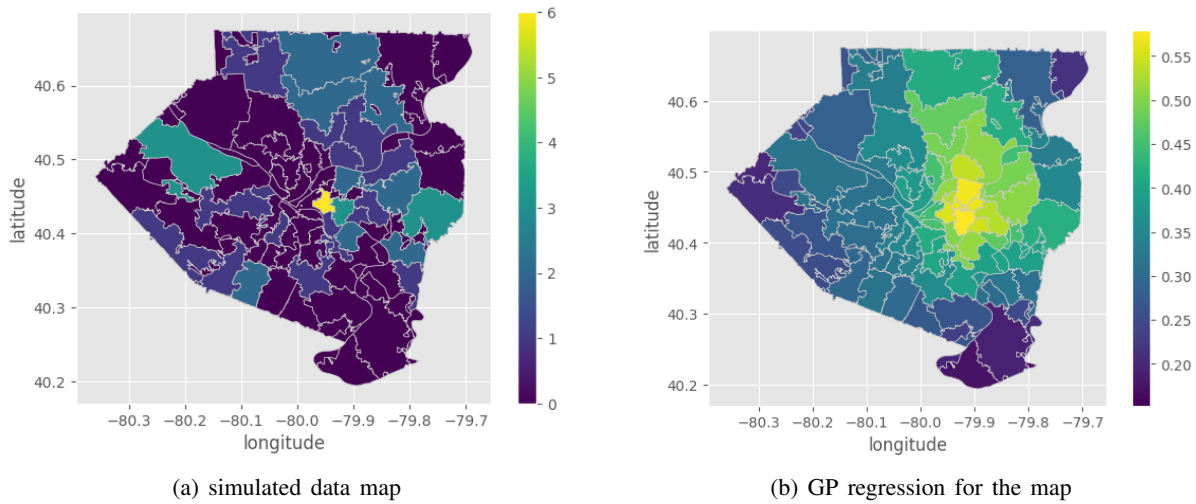


Figure 5: OUD model simulated data and GP regression for Allegheny county, PA.




Figure 6: Moran's I test values for INF and OUD.

different spatial autocorrelation measures can be attributed completely to the differing epidemic dynamics — all the relevant conditions in the synthetic population (whom each agent lives with within their household, whom they meet in their workplaces, etc.) are otherwise fixed. Such sharp inferences are typically impossible with real data due to the variety of factors that confound our comparisons. In this paper, we showed that the same area of Allegheny County, PA, with the same agent's information will end up having different inferences for each disease due to their differing dynamics. These results demonstrate the potential of FRED simulations for comparing treatment conditions on identical synthetic populations, and thus controlling for every confounding factor that may otherwise compromise our contrasts. We also illustrated the utility of Gaussian process regression as an inference tool for comparing epidemic dynamics on FRED-synthesized data. Such surrogate modeling techniques help us circumvent the computational burden of generating extensive simulation samples from large population dynamics.

This initial analysis opens several directions for future work, including a comparison of more complex spatial dependencies, e.g., spillovers across adjacent counties that can directly inform public health policy and intervention design Holtz et al. (2020), and extending the methodology for spatiotemporal regression. Our current analysis focused only on spatial relations and investigating inferences with spatiotemporal GPR is an important future direction.

DATA AND CODE

In this study, we are not able to share detailed data about the OUD model for contractual reasons. The repository link for the paper's code can be found at <https://github.com/abdurahmanfci/gpr-abm>.

ACKNOWLEDGMENTS

This research was funded by contract 75D30121C12574 from the Centers for Disease Control and Prevention. The findings and conclusions in this work are those of the authors and do not necessarily represent the official position of the Centers for Disease Control and Prevention.

This research was supported in part by the University of Pittsburgh Center for Research Computing, RRID:SCR_022735, through the resources provided. Specifically, this work used the HTC and VIZ clusters, which are supported by NIH award number S10OD028483.

REFERENCES

- Ahmed, A. A., M. A. Rahimian, and M. S. Roberts. 2023. "Estimating Treatment Effects Using Costly Simulation Samples from a Population-Scale Model of Opioid Use Disorder". In *Proceedings of IEEE International Conference on Biomedical and Health Informatics*. In Press.
- Ball, F., and T. House. 2017. "Heterogeneous Network Epidemics: Real-Time Growth, Variance and Extinction of Infection". *Journal of Mathematical Biology* 75(3):577–619.
- Buckingham-Jeffery, E., V. Isham, and T. House. 2018. "Gaussian Process Approximations for Fast Inference from Infectious Disease Data". *Mathematical Biosciences* 301:111–120.
- Cajka, J. C., P. C. Cooley, and W. D. Wheaton. 2010. "Attribute Assignment to a Synthetic Population in Support of Agent-Based Disease Modeling". *Methods Report (RTI Press)* 19(1009):1.
- CDC. 2023. "Understanding Drug Overdose and Deaths". <https://www.cdc.gov/drugoverdose/epidemic/index.html>, accessed March 20, 2023.
- Chao, D. L., M. E. Halloran, V. J. Obenchain, and I. M. Longini Jr. 2010. "FluTE, a Publicly Available Stochastic Influenza Epidemic Simulation Model". *PLoS Computational Biology* 6(1):e1000656.
- Duvenaud, D. 2014. *Automatic Model Construction with Gaussian Processes*. Ph. D. thesis, University of Cambridge.
- Fearnhead, P., V. Giagos, and C. Sherlock. 2014. "Inference for Reaction Networks Using the Linear Noise Approximation". *Biometrics* 70(2):457–466.
- Ferguson, N. M., D. A. Cummings, C. Fraser, J. C. Cajka, P. C. Cooley, and D. S. Burke. 2006. "Strategies for Mitigating an Influenza Pandemic". *Nature* 442(7101):448–452.
- Grefenstette, J. J., S. T. Brown, R. Rosenfeld, J. DePasse, N. T. Stone, P. C. Cooley, W. D. Wheaton, A. Fyshe, D. D. Galloway, A. Sriram et al. 2013. "FRED (A Framework for Reconstructing Epidemic Dynamics): An Open-Source Software System for Modeling Infectious Diseases and Control Strategies Using Census-Based Populations". *BMC Public Health* 13(1):1–14.
- Guclu, H., S. Kumar, D. Galloway, M. Krauland, R. Sood, A. Bocour, T. B. Hershey, E. van Nostrand, and M. Potter. 2016. "An Agent-Based Model for Addressing the Impact of a Disaster on Access to Primary Care Services". *Disaster Medicine and Public Health Preparedness* 10(3):386–393.
- Holtz, D., M. Zhao, S. G. Benzell, C. Y. Cao, M. A. Rahimian, J. Yang, J. Allen, A. Collis, A. Moehring, T. Sowrirajan et al. 2020. "Interdependence and the Cost of Uncoordinated Responses to COVID-19". *Proceedings of the National Academy of Sciences* 117(33):19837–19843.
- Jalal, H., J. M. Buchanich, M. S. Roberts, L. C. Balmert, K. Zhang, and D. S. Burke. 2018. "Changing Dynamics of the Drug Overdose Epidemic in the United States from 1979 through 2016". *Science* 361(6408):eaau1184.
- Krauland, M. G., R. J. Frankeny, J. Lewis, L. Brink, E. G. Hulsey, M. S. Roberts, and K. A. Hacker. 2020. "Development of a Synthetic Population Model for Assessing Excess Risk for Cardiovascular Disease Death". *JAMA Network Open* 3(9):e2015047–e2015047.
- Liu, F., W. T. Enanoria, J. Zipprich, S. Blumberg, K. Harriman, S. F. Ackley, W. D. Wheaton, J. L. Allpress, and T. C. Porco. 2015. "The Role of Vaccination Coverage, Individual Behaviors, and the Public Health Response in the Control of Measles Epidemics: An Agent-Based Simulation for California". *BMC Public Health* 15(1):1–16.
- Lukens, S., J. DePasse, R. Rosenfeld, E. Ghedin, E. Mochan, S. T. Brown, J. Grefenstette, D. S. Burke, D. Swigon, and G. Clermont. 2014. "A Large-Scale Immuno-Epidemiological Simulation of Influenza A Epidemics". *BMC Public Health* 14:1–15.
- Matthews, A. G., M. van der Wilk, T. Nickson, K. Fujii, A. Boukouvalas, P. León-Villagrà, Z. Ghahramani, and J. Hensman. 2017. "GPflow: A Gaussian Process Library Using TensorFlow". *Journal of Machine Learning Research* 18(40):1–6.

- Moran, P. A. 1950. "A Test for the Serial Independence of Residuals". *Biometrika* 37(1/2):178–181.
- Parker, J., and J. M. Epstein. 2011. "A Distributed Platform for Global-Scale Agent-Based Models of Disease Transmission". *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 22(1):1–25.
- Potter, M. A., S. T. Brown, P. C. Cooley, P. M. Sweeney, T. B. Hershey, S. M. Gleason, B. Y. Lee, C. R. Keane, J. Grefenstette, and D. S. Burke. 2012. "School Closure as an Influenza Mitigation Strategy: How Variations in Legal Authority and Plan Criteria can Alter the Impact". *BMC Public Health* 12:1–11.
- Rey, S. J., and L. Anselin. 2007. "PySAL: A Python Library of Spatial Analytical Methods". *The Review of Regional Studies* 37(1):5–27.
- Senanayake, R., S. O’Callaghan, and F. Ramos. 2016. "Predicting Spatio-temporal Propagation of Seasonal Influenza Using Variational Gaussian Process Regression". In *Proceedings of the AAAI Conference on Artificial Intelligence*, Volume 30.
- Tuite, A. R., A. L. Greer, M. Whelan, A.-L. Winter, B. Lee, P. Yan, J. Wu, S. Moghadas, D. Buckeridge, B. Pourbohloul et al. 2010. "Estimated Epidemiologic Parameters and Morbidity Associated with Pandemic H1N1 Influenza". *CMAJ* 182(2):131–136.
- W. D. Wheaton 2012. "US Synthetic Population Database 2005–2009: Quick Start Guide. RTI International".
- Williams, C. K., and C. E. Rasmussen. 2006. *Gaussian Processes for Machine Learning*, Volume 2. MIT press Cambridge, MA.
- Zimmer, C., and R. Yaesoubi. 2020. "Influenza Forecasting Framework Based on Gaussian Processes". In *International Conference on Machine Learning*, 11671–11679. PMLR.
- Zimmer, C., R. Yaesoubi, and T. Cohen. 2017. "A Likelihood Approach for Real-time Calibration of Stochastic Compartmental Epidemic Models". *PLoS Computational Biology* 13(1):e1005257.

AUTHOR BIOGRAPHIES

ABDULRAHMAN A. AHMED is a PhD student in the Department of Industrial Engineering at University of Pittsburgh. He obtained his MSc and BSc in Operations Research and Computer Science from Cairo University. His current research is on using surrogate models to develop inference procedures on complex sociotechnical systems with a focus on public health. His email address is aba173@pitt.edu.

M. AMIN RAHIMIAN is an Assistant Professor in the Department of Industrial Engineering at University of Pittsburgh. His current research focus is on challenges of inference and intervention design in complex, large-scale sociotechnical systems with applications ranging from social networks and e-commerce, to public health. His email address is rahimian@pitt.edu, and his website is <https://aminrahimian.github.io>.

MARK S. ROBERTS is a distinguished professor in the Department of Health Policy and Management at University of Pittsburgh, and holds secondary appointments in Medicine, Industrial Engineering, Business Administration, and Clinical and Translational Science. His recent research has concentrated in the use of mathematical methods from operations research and management science, including Markov Decision Processes, Discrete Event, and Agent-based Simulation. His email address is mroberts@pitt.edu and his page is <https://www.sph.pitt.edu/directory/mark-roberts>.