Bayesian Structural Time Series for Mobile Health and Sensor Data: A Flexible Modeling Framework for Evaluating Interventions

Jason Liu[1*], Daniel J. Spakowicz[2,3*], Rebecca Hoyd[2], Garrett I. Ash[4,5], Shaoke Lou[1], Donghoon Lee[1], Jing Zhang[1], Carolyn Presley[2], Ann Greene[6], Andrew V Papachristos[7], Mark Gerstein[8+]

[*] Co-first authors
[+] To whom correspondence should be addressed.

[1]Program in Computational Biology and Bioinformatics, Yale University, Bass 437, 266 Whitney Ave, New Haven, CT 06520, USA; Department of Molecular Biophysics and Biochemistry, Yale University, Bass 437, 266 Whitney Ave, New Haven, CT 06520, USA.
[2]Division of Medical Oncology, The Ohio State University Comprehensive Cancer Center, A455A Starling-Loving Hall, 320 W. 10th Ave, Columbus, OH 43210, USA.
[3]Department of Biomedical Informatics, The Ohio State University College of Medicine, BRT 482, 460 W 12th Ave, Columbus, OH 43210, USA.
[4]The Pain, Research, Informatics, Medical comorbidities, and Education (PRIME) Center, Veterans Affairs Connecticut Healthcare System, Building 35A, Rm 235, 950 Campbell Avenue, West Haven, CT 06516, USA
[5]Center for Medical Informatics, Yale University, Suite 501, 300 George St, New Haven, CT 06511, USA
[6]Yale School of Medicine, 333 Cedar St, New Haven, CT 06510, USA
[7]Department of Sociology, Northwestern University, 1810 Chicago Ave, Room 223, Chicago, IL 60208, USA
[8]Program in Computational Biology and Bioinformatics, Yale University, Bass 432, 266 Whitney Ave, New Haven, CT 06520, USA; Department of Molecular Biophysics and Biochemistry, Yale University, Bass 432, 266 Whitney Ave, New Haven, CT 06520, USA; Department of Computer Science, Yale University, Bass 432, 266 Whitney Ave, New Haven, CT 06520, USA. Electronic address: pi@gersteinlab.org.

ABSTRACT

The development of mobile health technology has the potential to contribute greatly to personalized medicine. Wearable sensors can assist with determining the proper treatment plans for individuals, provide quantitative information to physicians, or give individuals an objective measurement of their health. However, though treatments and interventions have become more targeted and specific, measuring the causal impact of these actions require more careful considerations of complex covariate structures as well as temporal and spatial properties of the data. Thus, emerging data from sensors and wearables in the near future will make use of and require complex models. Here, we describe a general statistical framework for sensor and wearable data that applies a Bayesian structural time series model to analyze and understand various behavior and health data collected in different environments. We show the wide applicability of this modelling framework, and how it corrects for covariates and biases to provide accurate assessments of intervention. Furthermore, it allows for a time dependent confidence interval of impact through its use of Bayesian estimation. We give three main examples, physical sensor data, environmental air sensors and longitudinal behavioral data to show the effect of various interventions through parameter estimation and comparison in pre- and post- intervention periods. The Bayesian structural time series model shows robust performance in a wide variety of tasks, further supporting its applicability to current and future mobile health and sensor data types.

1. INTRODUCTION

1.1 Background

The wearable sensor market was valued at $9 billion USD in 2018 and is expected to triple in the next decade (1). Investment is bolstered by the great potential for advancing personalized medicine in the near future (2-4). Whereas medicine has previously focused on determining the right interventions, it is now more focused on the individual and temporal aspects of treatment (5). Identifying the right time for treatments has still been relatively understudied, but this trend is expected to change soon as large streams of sensor data are released (6). As sensor technology develops, data-rich features such as physical, chemical, and behavioral variables will be measurable. In addition to time series data, spatial information is becoming more popular as well (7), all of which can be used for a more detailed understanding of sensors. A survey of distinct types of data are given in Figure 1.

Though there is an increasing growth of sensor data being released, there still is a lack of analytical methods in the biomedical field that can accommodate the complex covariate structures as well as the temporal and local trend considerations necessary to analyze these longitudinal data. Existing methods applied to these situations are limited. First, they only accommodate data that employs course sampling, known as low intensity data (8-10). Second, they address longitudinally observed patterns that could inform the timing of an intervention but do not evaluate the impact and efficacy of such intervention.

While the nature of data from sensors and wearables will vary depending on the context, most of them share numerous properties. A variety of such data can be seen in Figure 1, all of which come from various contexts. It is important to discern these properties and to establish a flexible model for emerging sensor and wearable data, as it will have broad implications in fields such as personalized medicine in the near future (2). Therefore, in this study we aim to contextualize a model and establish a flexible statistical framework to model various sensor and mobile health data. The model we use here is the Bayesian structural time series model and relies also on the Causal Impact package from Google (21). While these models have been extensively used in fields such as finance (11), ad campaigns (12), and marketing (13), they have rarely been used in the context of wearable and sensor data analysis. Therefore, we hope to bring this widely accepted model to a context more closely related to personalized medicine.

1.2 Modelling Framework

To establish an intuitive modeling framework for data taken from sensors and other mobile health sources, it is important to lay out the structure of how the framework is formed as well as any assumptions that may be essential for the model. First, we must recognize that the dataset is a time series data with a known response variable that evolves over time. For example, this could be a person's weight. This response variable could be a direct measurement taken from a sensor, or it could be a derived value – some scalar calculated from various underlying variables measured by a device – such as calories burned (14) or a Nike Fuel Score (15). The response variable's evolution in time is important, as it should be dynamic, and change based on covariates and interventions. In the example of weight, the measured weight fluctuates over time based on things such as seasonality (16), temperature (17), or diet (18).

Because we are interested in the impact that an intervention has on such a variable, an important assumption is that we know when the intervention occurs as well as the duration. There exist methods that detect interventions (19, 20) though to jointly assess the impact would result in a significant statistical disadvantage due to the multiple hypothesis corrections since the predictions from these models only account for one time point rather than a set of time points. Therefore, by having prior knowledge on the occurrence and duration of the intervention, we are able to more accurately assess the impact on an individualized level.

Given a specific intervention and its duration, we can split our time series into two segments – the pre- and post- intervention periods. These periods are used to describe all variables, known and unknown, as well as any parameters associated with them. For the example of weight and diet, the pre-intervention would be the period before the diet starts and the post-intervention is the period after the diet starts.

Our goal is to model the pre-intervention period using covariates to most accurately assess any changes in behavior in the response variable before any intervention has occurred. We also assume that any covariates used in our model should not be affected by the intervention, to provide a similar control in the post-intervention period. Though the post-intervention covariates themselves may change in value, the assumption is that they are derived from the same distribution as in the pre-intervention. Finally, using the model derived from the pre-intervention period and the covariates in the post-intervention, we can calculate a counterfactual in the post-intervention. The counterfactual predicts the response variable without intervention. It serves as a baseline to compare the actual observation in the post-intervention and ultimately is what we use to calculate the impact of an intervention. Compared to using linear models, this framework allows for an evolving measure of impact, due to the dynamic confidence interval for the difference between counterfactual and observation that is inherent to using Bayesian structural time series. This temporal consideration, in addition to the more common advantage of using hyper-parameters and priors, is an important consideration for the Bayesian framework, and sets it apart from other models.

This modeling framework can make use of complex covariate structures, which is useful and necessary to get an unbiased measure of impact. Two main types of covariates can be used: those that have known effects on the response and those that can account for hidden effects. The first type refers to covariates that could be correlated to the response, or cause changes to the response unrelated to the intervention. In relation to our example about weight, covariates of this type may include temperature, weather, or season. Since it is very unlikely to know all covariates of the first type that may perturb the response variable, we also introduce the notion of a second type of covariate, termed a "paired covariate". The paired covariate is an independent stream of sensor data of the same type as our original data in question. They could be collected in different locations, but the ideal case is they receive little to no treatment or intervention. If we imagine a scenario where the roommate of our subject is subject to many of the same external factors, but the roommate does not participate in the dieting, then this roommate's weight could be considered the paired covariate. Since both the original data and paired covariate share underlying biases, by using the paired covariates we can determine the true impact of the intervention.

As a whole, this modeling framework addresses many of the limitations of other methods and serves to advance the field by providing a structured method for analyzing data from sensors and mobile health sources. We give more details in the Statistical Formalism section and provide a schematic of the framework in Figure 2.

2. STATISTICAL FORMALISMS

2.1 Bayesian Structural Time Series

In order to understand the causal impact of interventions on longitudinal datasets that may be affected by a wide range of factors it is important to have a statistical framework that considers carefully the prior and posterior measurements as well as covariates that may affect the response variable. We detail aspects of this modeling framework below.

We assume for a given time point, t, there exists an observation y_t, linked to a variety of other parameters, a_t: $\{\mu_t, \delta_t, \boldsymbol{x}_t, \boldsymbol{\eta_t} \sim N(0, \boldsymbol{\sigma_t^2})\}$. As time progresses from t to $t+1$, the other parameters also progress due to their time dependency.

There also exists a set of hyper-parameters π: $\{\varepsilon_t, \eta_t \sim N(0, \sigma_t^2), \dots\}$ that define the initial conditions of the parameters above. The hyper-parameter $\pi$is a prior that reflects previous knowledge about the system and is used to estimate the best set of parameters a_t. More specifically, we use the following equations to represent our time series data.

$$y_t = \mu_t + \tau_t + \beta^T \boldsymbol{x}_t + \eta_t^y \tag{1-1}$$

$$\mu_{t+1} = \mu_t + \delta_t + \eta_t^\mu \tag{1-2}$$

$$\delta_{t+1} = \delta_t + \eta_t^\delta \tag{1-3}$$

Here we see that the equation involving the observation or response variable, y_t, depends on μ_t, and additionally a variable $\boldsymbol{x}_t$ to represent the covariates. Furthermore, we additionally include another layer of dependency, δ_t, which μ_t depends on. Because our observation, y_t, can be biased, it is important to take into account as many covariates in the form of $\boldsymbol{x}_t$. Each equation has an error term represented by η_t. Since each of the processes are different, the error term is also different. For example, the error associated with μ_t is η_t^{μ}. These η_t are drawn from some distribution satisfying $N(0, \sigma_t^2)$. Similar to other regression-based models, the covariates here are scaled by a coefficient vector β^T. The model also employs a spike and slab method to penalize and select for important covariates when the covariate matrix is large. The variable μ_t, the underlying "state" of the response, contributes to our response variable, and its representation is given as a time dependent equation. If we ignore the δ_t variable, our equation becomes

$$\mu_{t+1} = \mu_t + \eta_t^{\mu} \tag{2}$$

and this is a representation for the random walk. That is to say, in this simplified form of μ_t, our observation is similar to a random walk, where at each time point, there is some progression to the next time point in random fashion, based on the parameter $\eta_{t,\mu}$. However, one optimization that can be done for this parameter is the inclusion of δ_t, which results in the pair of equations

$$\mu_{t+1} = \mu_t + \delta_t + \eta_t^{\mu} \tag{3-1}$$
$$\delta_{t+1} = \delta_t + \eta_t^{\delta} \tag{3-2}$$

Here δ_t follows a random walk, and the μ_t equation is dependent on δ_t. That is to say, δ_t serves as a trajectory or slope parameter that helps to guide the behavior of μ_t. We can still imagine μ_t as a random walk, but now with some additional trend vector. If μ_t is steadily increasing, it is likely that at μ_{t+1} there is also an increase, due to the inclusion of δ_t. In other words, the δ_t allows for more stability between each time point and serves a purpose similar to a slope parameter. We provide more examples in Supplemental Figures S1-5.

2.2 Evaluating Intervention Impact

While the above model describes longitudinal sensor data well, it is important to consider how the pre- and post-intervention periods differ. To determine if an intervention was effective in bringing about a change in our dependent measurement, y_t, it is important to have a framework in which comparisons can be quantified. To do so, we define the pre-intervention period as time points $t = 1, \dots, n$, while the post-intervention period is defined as $t = n + 1, \dots, N$. Furthermore, we define observations as y and predictions of the model as y'. Therefore, the set of observations in the pre-intervention period, $y_{1,\dots,n}$, serve as the training data, based on all covariates. At each time point $t = 1, \dots, n$, we update our parameter set of the model to better fit $y_{1,\dots,n}$ using a Markov Chain Monte Carlo (MCMC) method of sampling from the distribution. More specifically the model parameters are estimated to minimize the difference between $y_{1,\dots,n}$ and $y'_{1,\dots,n}$. After estimation of parameters in the pre-intervention period, we predict the counterfactual in the post-intervention, $y'_{n+1,\dots,N}$. This is done by using the parameter set defined in the pre-intervention period and the covariates from the post-intervention period. Because $y'_{n+1,\dots,N}$ has an associated prediction interval, we can calculate a significance p-value associated with the difference between $y'_{n+1,\dots,N}$ and $y_{n+1,\dots,N}$. While we can get a p-value at every time point in the post intervention, it is more useful to consider the impact that an intervention had on the whole post-intervention period. The p-value associated with the full post-intervention time period ($t = n + 1, \dots, N$) is known as the *cumulative impact*. It should be noted that the prediction interval associated with $y'_{n+1,\dots,N}$ generally will increase as time progresses. This is one advantage of the Bayesian model and allows for an evolving cumulative impact. This is due to the fact that at every time point after the intervention, the variance associated with the distribution that the prediction is drawn from is compounded at every time point. It is useful to factor in this temporal aspect since the confidence that an intervention resulted in a causal impact may be dependent on time. For our analysis, we use the package CausalImpact in R (21).

3. RESULTS

3.1 Physical Sensor Data

Using Google Science Journal app on Android, we collected various mobile sensor data for a simple experiment involving a person spinning in different scenarios. The goal of this analysis was to demonstrate the flexibility of our model in determining significant causal impact of human physical actions before and after an intervention. A schematic of the results is given in Figure 3A, B, and C. First, a single user's longitudinal data was used to estimate the effect of an intervention (Figure 3A). An example of this could be sleep data and the effect of changing one's pillow, and where no other data are available to the user. Here we show an individual spinning with a sensor measuring acceleration when held at arm's length; the intervention is to bring the sensor closer to the body while spinning at the same rate, thereby increasing the acceleration. The model successfully detected the state change associated with the intervention, as shown by the cumulative change relative to the predicted values.

Next, we demonstrate how the previous scenario is susceptible to noise, i.e. something that affects the signal but is not related to the intervention. In this case the noise is represented by a hop, which occurred after the intervention began and was therefore not predicted by the model (Figure 3B). In this case the model does not detect an effect of the intervention – the hop affected the signal to such a degree that the confidence in the prediction decreased, i.e. the credible intervals widened greatly. To continue the sleep study analogy mentioned above, the noise could be loud neighbors moving in next door (adding noise in a literal sense), which leads to worse sleep in a way that is unrelated to the pillow intervention.

Finally, we show how including a paired covariate -- another data stream that affects the state being measured but is not related to the intervention – can effectively correct for noise (Figure 3C). In this case two sensors are held while spinning, but only one is subject to the intervention (brought close to the chest). Both, however, experience the noise (hop), which the model is then able to control and correctly identifies the intervention like the case without noise (Figure 3A). In the sleep analogy, an effective paired covariate could be another device measuring the sleep of one's partner sharing the same bed, but who did not change their pillow. This is distinct from a non-paired covariate which could be similarly effective, e.g. a device measuring the decibel level in the room; it is relevant to the state (sleep quality), is not affected by the intervention (pillow change) and identifies the noise that affects the state (loud neighbors). However, a paired covariate has the characteristic that it is able to control for unknown confounders, akin to a control arm of a randomized controlled trial.

3.2 Environmental Sensor Data

We performed a similar analysis using data collected from an AWAIR monitor, with measurements of CO_2, dust, humidity, and temperature (Figure 1E). The data are collected over a one-month period, with measurements taken every 15 minutes. The data are aggregated at each time point across all 30 days to give a smoother signal and limit the analysis to one intervention with a clearly delineated a pre- and post- intervention period, namely exposure of the room to people. It should be noted that though intervention often implies some treatment put in place, it can be widely adapted in the realm of Bayesian structural time series to allow for any disruption or change in status such as the effect that people have on environmental CO_2 levels. After correcting for covariates such as dust, humidity, and temperature (22), here is a significant increase in CO_2in the hours when workers are in in the room with the sensor (p-value < 0.001). Figure 3D shows the CO_2measurements across the aggregated time points in a 24-hour time frame. We can see that as CO_2levels increase drastically after the 9am time point, the cumulative causal impact in the post-intervention period increases. This cumulative impact increases until around 5pm in the evening, where the cumulative impact tapers off, signaling a state is reached in the post-intervention that is similar to that of the pre-intervention period. Although there is one p-value associated with the whole post-intervention period, this modeling framework can show confidence intervals for each time point in the post-intervention period as a function of both the observation and covariates. This helps us to define period of time where the intervention is most effective.

3.3 Human Behavioral Data

In addition to the longitudinal data, many mobile health devices also collect location information. For example, many watches collect GPS tracks of a run as well as steps and heart rate (e.g., Apple Watch, Garmin Forerunner). The Bayesian structural time series framework can flexibly accommodate location data through the inclusion of spatial correlation matrices. In addition, the spatial data can be used to segment the data where the intervention does and does not have an effect. By this method one can create "synthetic" paired-covariates, in that data are from different spatial segments from the same source. However, spatial data are often protected for privacy concerns and are therefore less accessible to researchers. We therefore demonstrate the utility of our modeling framework on more accessible behavioral data that share many characteristics: the effect of a social monitoring application called SeeClickFix on negative behavioral patterns (crime).

Crime patterns show similar characteristics to spatial mobile health data. They are affected by covariates such as temperature and precipitation (23) and can be linked to many other features such as census data of household incomes or education (24) (analogous to covariates of movement such as age, physical fitness, etc. (25)). SeeClickFix is a smartphone and web application developed to allow users to report issues in their communities including non-violent crimes. Posts can be voted on and supported by other users' comments and local government agencies acknowledge issues and post when they have been addressed. It has been hypothesized that SeeClickFix and similar tools may reduce crime through establishing social cohesion, promoting collective efficacy (26).

Following the intervention (creation of SeeClickFix) there was no detectable decrease in crime across the entire area (Figure 4A), nor in particular neighborhoods (Figure 4B). However, one would expect many other factors besides the introduction of SeeClickFix to affect crime in this time (e.g. increases in the police force, changes to local employment, city-wide initiatives). The modeling framework can leverage the spatial information in these behavioral data to search for a paired-covariate, in this case referring to locations not affected by the intervention (no SeeClickFix use) in order to control for other, unobserved effects on the outcome variable (increases in the city police force). We aggregated crimes and SeeClickFix posts by neighborhood (Figure 4C) and then modeled used a neighborhood that was not affected by the intervention as a control (i.e. a neighborhood without SeeClickFix posts), but that would experience any city-wide initiatives that might affect crime (Figure 4D). Neighborhoods with heavy SeeClickFix use showed an effect of the intervention on crime when controlling for unobserved factors with synthetic paired covariates (Figure 4E).

4. DISCUSSION

In this paper, we collected data involving three various kinds of longitudinal data that closely reflect the properties seen in mobile health sensor and wearable data. While some studies demonstrating the success that linear models have on understanding such data do exist (8-10, 19, 20), there is a lack of emphasis on considerations for the temporal aspect of interventions and how effective interventions are at different time points in the post-intervention period. Linear modeling frameworks lack the flexibility to evaluate the intervention's effect strength at all points in the post-intervention period. In this study, we illustrate the benefit of using a Bayesian structural time series model in modeling the behavior of various longitudinal data collected via apps and sensors. These data demonstrate properties that are commonly found across other wearable sensor data, which are increasingly gaining in popularity (3). In order to effectively integrate such data with personalized medicine and health in the future, it is important to understand these data and what kinds of models can be used to analyze them (2-4).

It is also important to consider the contexts in which this modeling framework does not yield significant advantage over a linear framework. A stable and consistent intervention effect over the entire post-intervention period should be evaluated equally well by linear models. Given the computational complexity of the Bayesian method due to MCMC sampling for each time point, such considerations could be very important for large datasets.

We show that the Bayesian modeling framework carefully considers complex covariates and can begin to see the richness of data in our behavioral sensor dataset, as the covariate structure between different locations is

complex and results in biases that require careful consideration when analyzing. Furthermore, the modeling framework also considers the temporal properties, ensuring that predictions of intervention impact are conditional on duration of the intervention. This is especially important for future data from sensors and mobile health sources as one of the major concerns of personalized medicine is ensuring the right time and duration of a treatment or intervention is considered (5). With time varying confidence intervals, we can begin to understand not only the effectiveness of an intervention but determine the most effective intervention plan necessary to reach a desired result.

5. METHODS AND MATERIALS

5.1 Data Collection

There exist many datasets similar to those of mobile sensors and we should not confine ourselves to just the traditional accelerometer and gyroscope data that one would traditionally think of when envisioning sensor data. In studying the causal impact of an intervention, we show that most data sharing longitudinal qualities allow for development of algorithms and exploring how such algorithms can be useful in analyzing various sensor data. The data we analyze in this paper consists of environmental sensor data, physical activity sensor data, and human behavioral sensor data.

Due to the longitudinal nature of sensor and wearable data, these data not only demonstrate interesting patterns in a response variable, but also are closely tied to the temporal property of a phenomena. By introducing the aspect of time, it becomes important to find models that leverage temporal considerations and use them to make accurate assessments about the data. Also, some of the data showcase complex covariates (paired) and can be used to better correct for unknown and hidden biases. More details about the data collection and data types are given below.

5.1.1 Google science journal data collection

Data were collected using the linear accelerometer measurement function in the Google Science Journal application on a Samsung Galaxy S8 smartphone. Each experiment lasted 20 seconds. One or two instruments were held at arm's length while spinning at a constant rate for 10 seconds. Next, one instrument was brought in close to the body for 10 seconds while maintaining the spinning rate. In the case of the noise simulation, the experimenter hopped once after 15 seconds (halfway through the intervention period). Data were exported from Science Journal and analyzed using the CausalImpact package in R.

5.1.2 AWAIR Collection

Data was collected from a one-month period from the AWAIR device. The device was placed in an office lab setting where individuals frequented on a daily basis. CO_2 levels were measured in units of ppm. AWAIR also measures dust, temperature, humidity, which were used as covariates.

5.1.3 SeeClickFix Data collection

The application SeeClickFix is a smartphone and web application developed in New Haven, Connecticut, where users report issues in their communities including non-violent crimes. SeeClickFix posts can be supported and commented on by other users, and local government agencies acknowledge and address issues. The SeeClickFix data are publicly available, providing a rich longitudinal and spatial dataset for monitoring behavior and interactions with other users and city representatives. Posts were aggregated by month for the New Haven metropolitan area and by neighborhood (n = 19) from 2007-2015.

Aggregated crime data were shared through a memorandum of understanding with the New Haven Police Department for 2000-2013. Rates were calculated using the 2014 ACS 5-year population estimates (crimes / 10K population per unit area).

FUNDING

This work was funded by the National Institutes of Health (NIH). GA was supported by a fellowship from the Office of Academic Affiliations at the United States Veterans Health Administration.

REFERENCES

(1) Wearable Sensors Market: Growth, Trends and Forecast (2019 - 2024) [Internet]. Available at: https://www.mordorintelligence.com/industry-reports/global-wearable-sensors-market. Accessed October 30, 2019.
(2) Hekler EB, Klasnja P, Chevance G, Golaszewski NM, Lewis D, Sim I. Why we need a small data paradigm. BMC Med 2019 Jul 17;17(1):133-019-1366-x.
(3) Sim I. Mobile devices and health. N Engl J Med 2019:956-968.
(4) Whitney RL, Ward DH, Marois MT, Schmid CH, Sim I, Kravitz RL. Patient Perceptions of Their Own Data in mHealth Technology-Enabled N-of-1 Trials for Chronic Pain: Qualitative Study. JMIR Mhealth Uhealth 2018 Oct 11;6(10):e10291.
(5) Corwin E, Redeker NS, Richmond TS, Docherty SL, Pickler RH. Ways of knowing in precision health. Nurs Outlook 2019 Jul - Aug;67(4):293-301.
(6) FitBit. Fitbit's 100+ billion hours of resting heart rate user data reveals resting heart rate decreases after age 40. 2014; Available at: https://investor.fitbit.com/press/press-releases/press-release-details/2018/Fitbits-100-Billion-Hours-of-Resting-Heart-RateUser-DataReveals-Resting-Heart-Rate-Decreases-After-Age-40/default.aspx. Accessed October 28, 2019.
(7) Seshadri DR, Li RT, Voos JE, Rowbottom JR, Alfes CM, Zorman CA, et al. Wearable sensors for monitoring the internal and external workload of the athlete. NPJ Digit Med 2019 Jul 29;2:71-019-0149-2. eCollection 2019.
(8) Suchting R, Hebert ET, Ma P, Kendzor DE, Businelle MS. Using Elastic Net Penalized Cox Proportional Hazards Regression to Identify Predictors of Imminent Smoking Lapse. Nicotine Tob Res 2019 Jan 4;21(2):173-179.
(9) Koslovsky MD, Hebert ET, Swartz MD, Chan W, Leon-Novelo L, Wilkinson AV, et al. The Time-Varying Relations Between Risk Factors and Smoking Before and After a Quit Attempt. Nicotine Tob Res 2018 Sep 4;20(10):1231-1236.
(10) Koslovsky MD, Swartz MD, Chan W, Leon-Novelo L, Wilkinson AV, Kendzor DE, et al. Bayesian variable selection for multistate Markov models with interval-censored data in an ecological momentary assessment study of smoking cessation. Biometrics 2018 Jun;74(2):636-644.
(11) Jacquier E, Polson N. Bayseian methods in finance. 2010:1-92.
(12) Jin Y, Wang Y, Yunting S, Chan D, Koehler J. Bayesian Methods for Media Mix Modeling with Carryover and Shape Effects. 2017:1-34.
(13) Rossi P, Allenby G. Bayesian Statistics and Marketing . Marketing Science 2003;22:304-328.
(14) Freedson PS, Melanson E, Sirard J. Calibration of the Computer Science and Applications, Inc. accelerometer. Med Sci Sports Exerc 1998 May;30(5):777-781.
(15) Nike. What is nikefuel? 2014; Available at: https://news.nike.com/news/what-is-nikefuel. Accessed October 28, 2019.
(16) Schoeller DA. The effect of holiday weight gain on body weight. Physiol Behav 2014 Jul;134:66-69.
(17) Naya DE, Naya H, White CR. On the Interplay among Ambient Temperature, Basal Metabolic Rate, and Body Mass. Am Nat 2018 Oct;192(4):518-524.
(18) Geddes R, Harvey JD, Wills PR. The molecular size and shape of liver glycogen. Biochem J 1977 May 1;163(2):201-209.
(19) Hebert ET, Stevens EM, Frank SG, Kendzor DE, Wetter DW, Zvolensky MJ, et al. An ecological momentary intervention for smoking cessation: The associations of just-in-time, tailored messages with lapse risk factors. Addict Behav 2018 Mar;78:30-35.
(20) Klasnja P, Smith S, Seewald NJ, Lee A, Hall K, Luers B, et al. Efficacy of Contextually Tailored Suggestions for Physical Activity: A Micro-randomized Optimization Trial of HeartSteps. Ann Behav Med 2019 May 3;53(6):573-582.
(21) Brodersen K, Gallusser F, Koehler J, Remy N, Scott S. Inferring Causal Impact Using Bayesian Structural Time-Series Models. ANN APPL STAT 2015;9:247-274.

(22) Lipsett MJ, Shusterman DJ, Beard RR. In patty's industrial hygeine and toxicology, 4th ed. ASHRAE Winter Conf Pap 1994;2:4552-7.
(23) Mares D. Climate change and crime: monthly temperature and precipitation anomalies and crime rates in St. Louis, MO 1990–2009. Crime Law Soc Change 2013;59:185-208.
(24) Reitzel JD. Socioeconomic status and offending. 2019:1-2.
(25) Aleksovska K, Puggina A, Giraldi L, Buck C, Burns C, Cardon G, et al. Biological determinants of physical activity across the life course: a "Determinants of Diet and Physical Activity" (DEDIPAC) umbrella systematic literature review. Sports Med Open 2019 Jan 8;5(1):2-018-0173-9.
(26) Mergel I. Distributed Democracy: SeeClickFix.com for Crowdsourced issue reporting. 2012 1/27/2012:1-19.

FIGURES

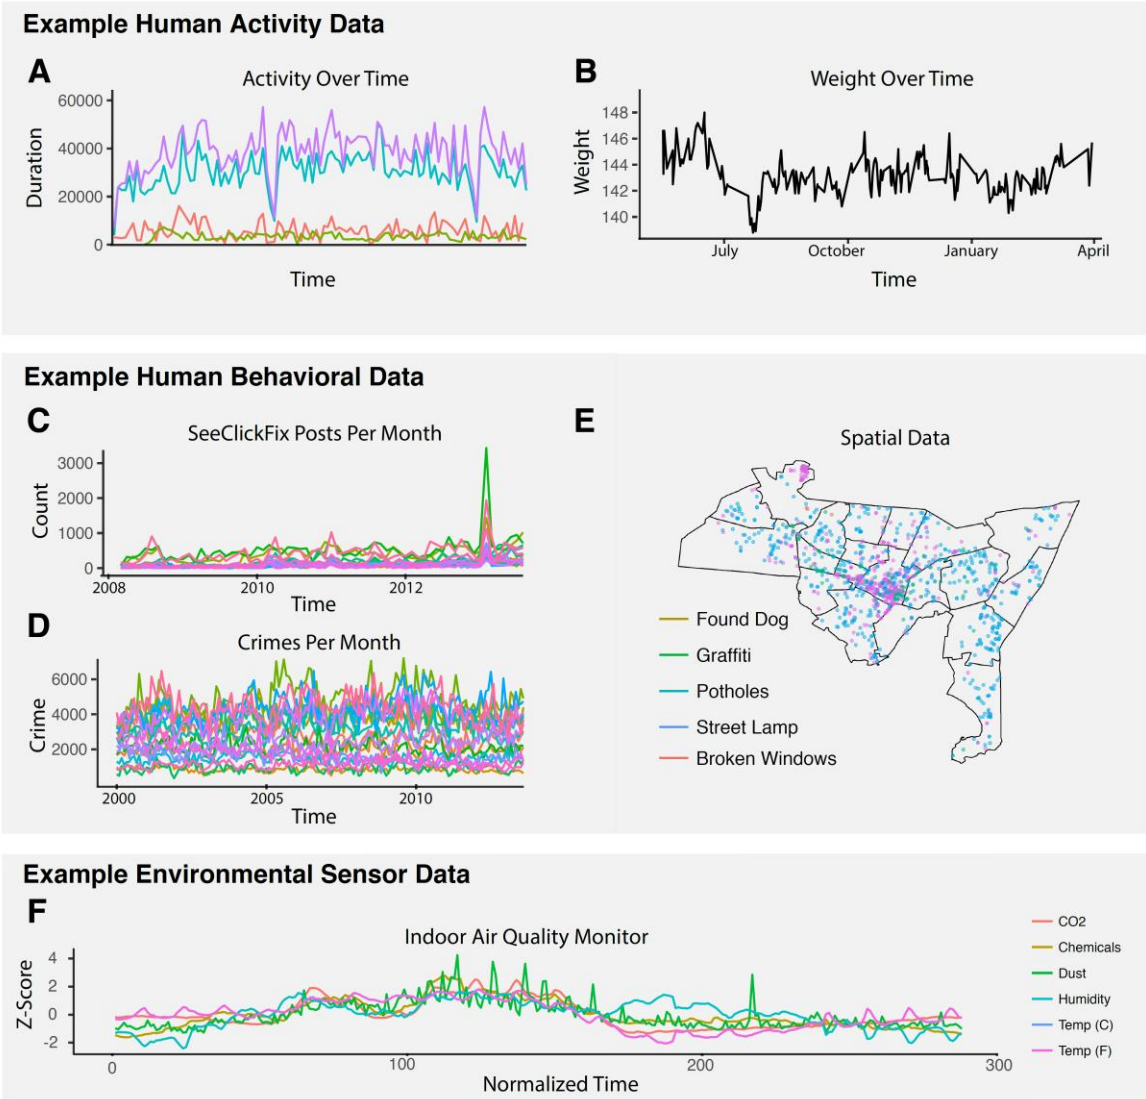

Figure 1. A) MOVES data set showing different exercise activities over time, with each color representing a different activity. B) Weight over time from the Withings data set. C) An example of human behavioral data, the use of SeeClickFix, over time. D) Various neighborhoods in New Haven and their crime data over time. E) Spatial map of New Haven illustrating the use of SeeClickFix and the associated geographical data. F) Various measurements taken from an indoor air quality sensor averaged over a day across a month.

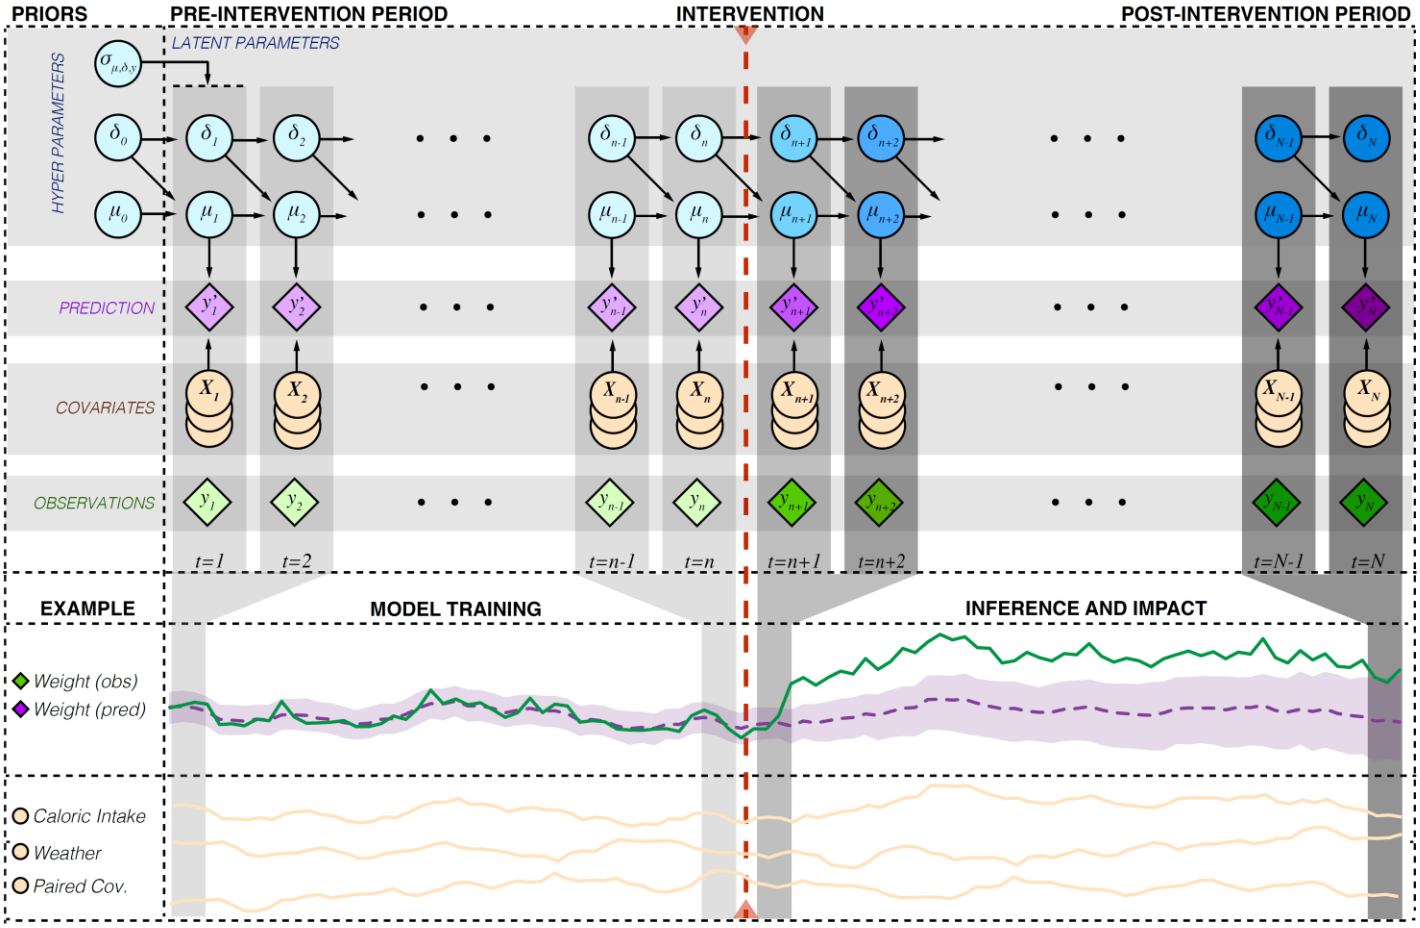

Figure 2. Schematic and illustration of the Bayesian structural time series. Latent parameters and hyperparameters are shown in blue, observations are shown in green, covariates in yellow, and predictions in purple. Predictions have an associated prediction interval shown in light purple. An illustrative example shows weight over time, with various covariates, being modeled. Assuming an intervention of increased diet, the model detects a strong impact on the post-intervention weight.

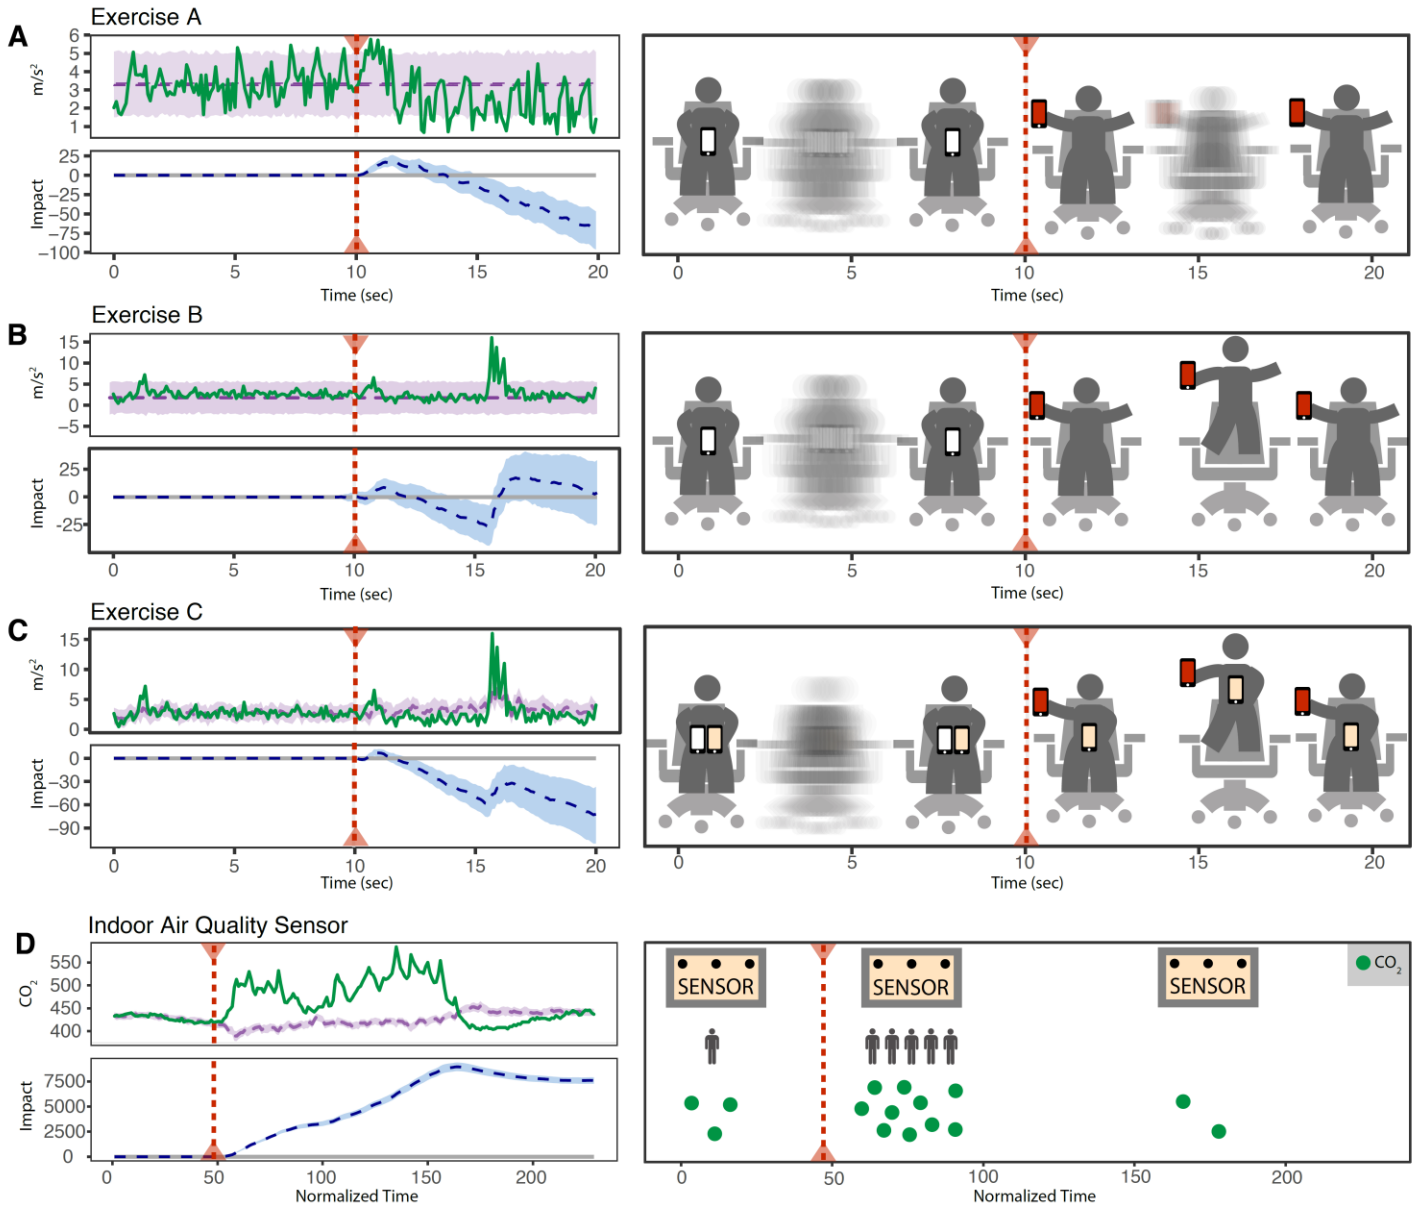

Figure 3. Performance of the Bayesian structural time series model in model experiments with known interventions. Accelerometer measurements of (A) a person spinning in a chair holding sensor near body and then extended to arm's length at the intervention; (B) with simulated "noise" produced by a hop during the intervention period and (C) with a paired covariate (second sensor) that is not affected by the intervention but experiences the "noise" hop. D) Using an indoor air quality sensor, CO_2 is measured with a variety of other covariates. The intervention (arrival of office members) causes an increase in CO_2 and is determined to be impactful using the model.

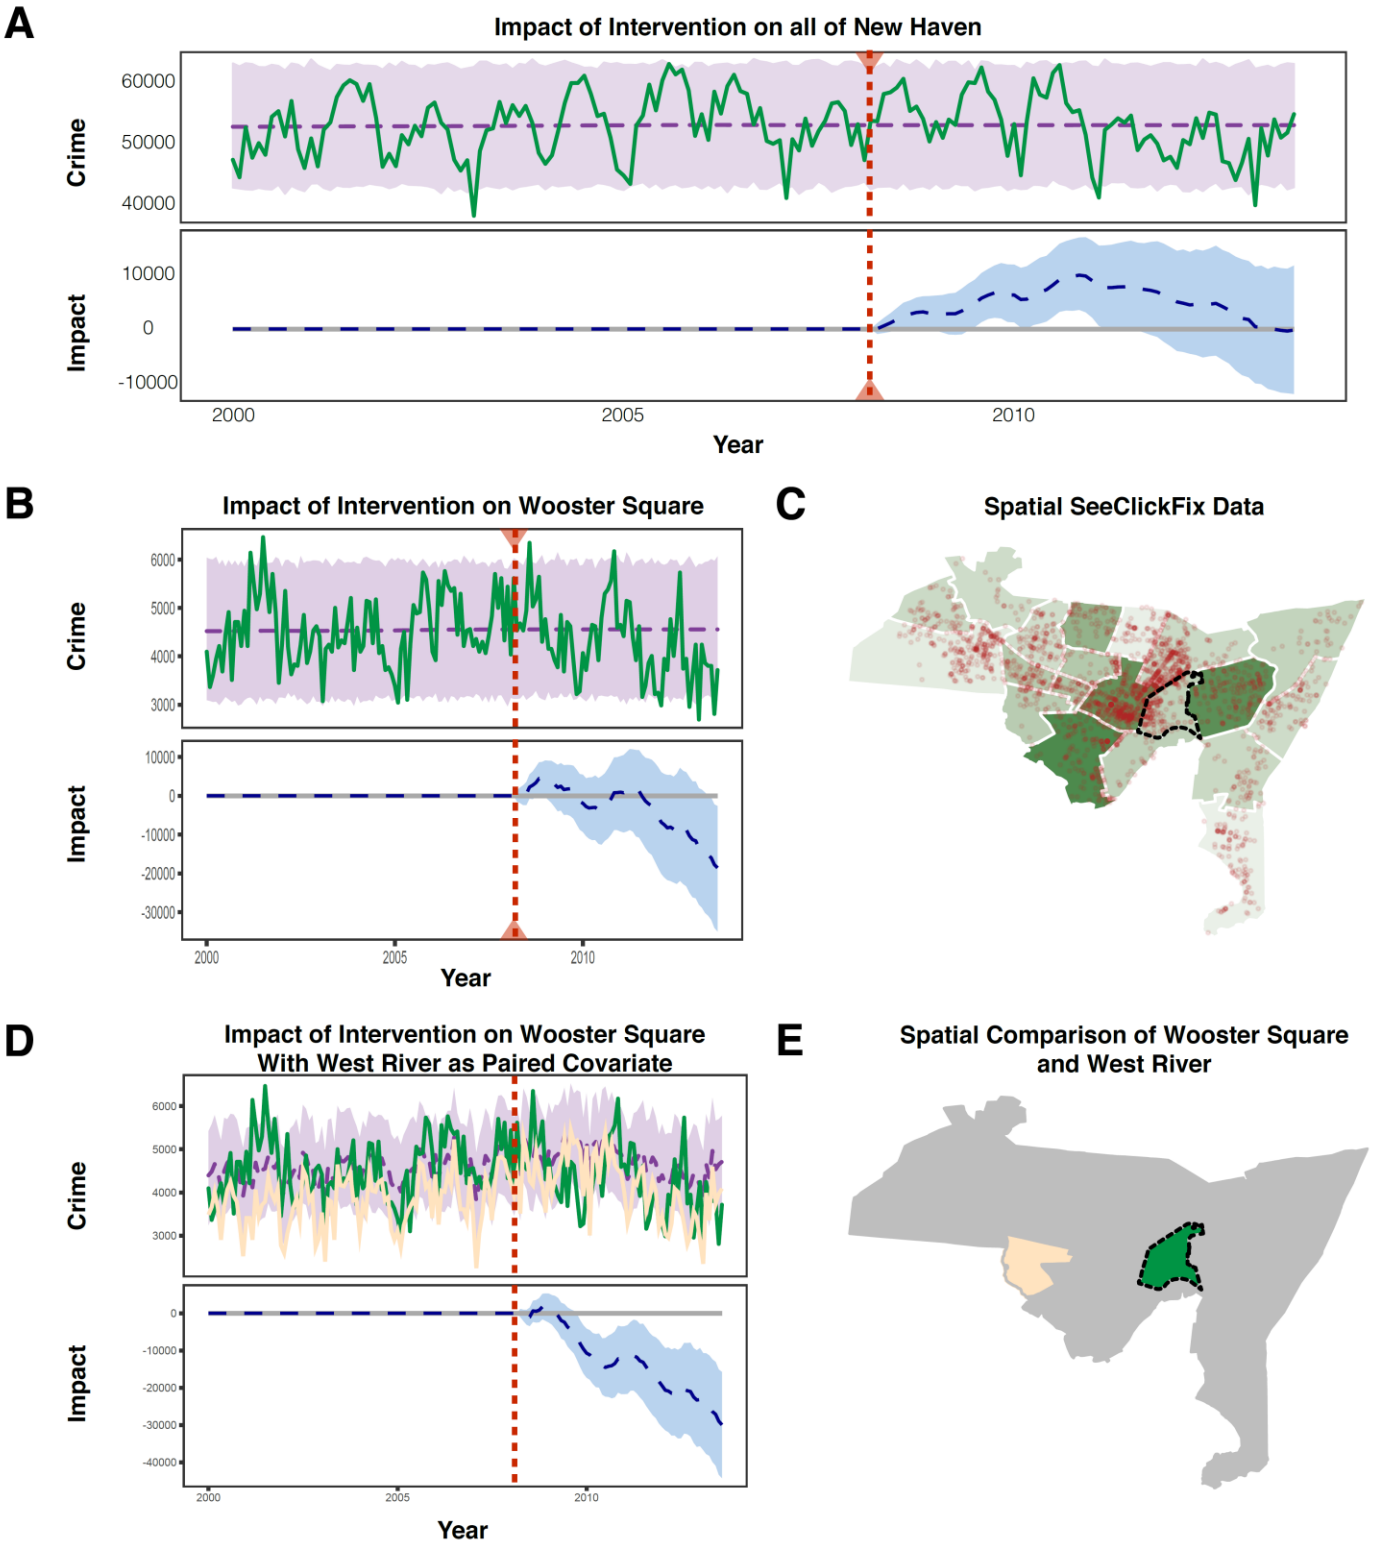

Figure 4. A) Impact of intervention (use of SeeClickFix) on all of New Haven. B) Impact of intervention on only one neighborhood, Wooster Square. C) Spatial data of New Haven showing variety of SeeClickFix usage throughout various neighborhoods. D) Impact of intervention on Wooster Square crime using West River as a paired covariate. E) Illustration showing the observed data, Wooster Square, versus the paired covariate, West River.