Dana E. Goin¹ / Jennifer Ahern¹

Identification of Spikes in Time Series

¹ University of California, Division of Epidemiology & Biostatistics, School of Public Health, California Berkeley, USA, E-mail: degoin@berkeley.edu, jahern@berkeley.edu. <https://orcid.org/0000-0002-7557-7977>.

Abstract:

Researchers interested in the effects of exposure spikes on an outcome need tools to identify unexpectedly high values in a time series. However, the best method to identify spikes in time series is not known. This paper aims to fill this gap by testing the performance of several spike detection methods in a simulation setting. We created simulations parameterized by monthly violence rates in nine California cities that represented different series features, and randomly inserted spikes into the series. We then compared the ability to detect spikes of the following methods: ARIMA modeling, Kalman filtering and smoothing, wavelet modeling with soft thresholding, and an iterative outlier detection method. We varied the magnitude of spikes from 10 to 50 % of the mean rate over the study period and varied the number of spikes inserted from 1 to 10. We assessed performance of each method using sensitivity and specificity. The Kalman filtering and smoothing procedure had the best overall performance. We applied each method to the monthly violence rates in nine California cities and identified spikes in the rate over the 2005–2012 period.

Keywords: time series, spikes, violence, additive outlier

DOI: 10.1515/em-2018-0005

Received: March 6, 2018; **Revised:** March 12, 2019; **Accepted:** July 24, 2019

1 Introduction

Identification of unexpectedly high values in a time series is useful for epidemiologists, economists, and other social scientists interested in the effect of an exposure spike on an outcome variable. Exposure spikes may be of interest when they are considered to be caused by something exogenous to the general patterning of the series, as in the case of income shocks and infant mortality (Baird, Friedman, and Schady 2011), thus strengthening the inference that can be drawn from the estimated effect (Humphreys et al. 2016). Furthermore, spikes may be of interest when extreme increases in the exposure series are hypothesized to have disproportionate effects on the response compared to more usual disturbances from expected values or compared to corresponding decreases in the exposure. Previous studies have examined the effects of weather or economic spikes on outcomes as diverse as civil conflict, birth weight, and nutrition (Bhattacharya et al. 2003; Jacob, Lefgren, and Moretti 2007; Miguel, Satyanath, and Sergenti 2004; Margerison-Zilko et al. 2011; Wilde, Apouey, and Jung 2017).

Community violence is another exposure that exhibits spikes. Exposure to community violence has been linked to stress-related health outcomes, including depression, asthma, cardiovascular disease, and birth weight (Kane 2011; Brown, Hill, and Lambert 2005; Apter et al. 2010; Clark et al. 2008; Koppensteiner and Manacorda 2016; Miller et al. 1999; Martinez and Richters 1993; Masi et al. 2007; Curry, Latkin, and Davey-Rothwell 2008; Ahern et al. 2018). However, many studies of community violence and health have suffered from structural confounding, in which the strong correlation of community violence with factors such as segregation, poverty, and unemployment means the effects are challenging to disentangle (Ahern et al. 2013; Michael Oakes 2006). Examination of spikes in community violence offers advantages when between-community comparisons would suffer from structural confounding, because researchers can compare individuals within a community over time.

To study the impacts of spikes in an exposure on health outcomes, spikes must be well characterized. In this study, we consider a spike to be an acute increase in the series followed by an immediate return to the underlying level of the series. This type of spike is described in the time series literature as an additive outlier Cryer and Chan (2008). The best method for identification of spikes in time series is not known. Many previous studies have defined spikes using pre-specified critical values above or below yearly averages (Bhattacharya et al. 2003; Jacob, Lefgren, and Moretti 2007; Miguel, Satyanath, and Sergenti 2004; Wilde, Apouey, and Jung 2017). However, such methods do not effectively account for underlying trends or autocorrelation. Some studies have used time series methods to identify spikes (Margerison-Zilko et al. 2011), but a comparison of different methods has not been done.

Dana E. Goin is the corresponding author.

© 2019 Walter de Gruyter GmbH, Berlin/Boston.

Identification of places in a time series where unanticipated changes occur has been considered from numerous perspectives. These include methods of change detection, structural breaks, level shifts, variance changes, or other types of outliers, (e. g. Besculides et al. 2005; Dafni et al. 2004; Unkel et al. 2012; Balke 1993; Tsay 1988; Chen and Liu 1993). In this work, we focus on a subset of these methods that detect additive outliers, evaluate each method's performance in simulations, and illustrate an application to monthly violence rates in nine California cities. The code used for the simulations is provided in the supplemental materials for researchers interested in applying these methods to different exposure series.

2 Simulation study

Assessment of outliers and extreme values in time series differs from the approaches used in non-time ordered data due to the potential for autocorrelation, trends, and cyclical patterns. We compared the following methods for identifying spikes: ARIMA modeling (Cryer and Chan 2008; Metcalfe and Cowpertwait 2009) using the `auto.arima` function from the `forecast` package (Hyndman et al. 2019; Hyndman and Khandakar 2008), Kalman filtering and smoothing (De Jong and Penzer 2000; Durbin and Koopman 2012) using the `KFS` function from the `KFAS` package (Helske 2014, 2017), wavelet modeling with soft thresholding (Nason 2008) using the `wd` and `wr` functions from the `wavethresh` package (Nason 2010), and an iterative outlier detection method (Cryer and Chan 2008) using the `detectAO` function from the `TSA` package (Chan and Ripley 2012). We selected these classes of methods because they represent the most common time series modeling and outlier detection approaches. In preliminary simulations we compared these approaches with variations, and selected for the main analysis those which performed best. Other variations that were considered included different types of wavelet thresholding, and several other iterative outlier detection methods including those described by Chen and Liu (1993) and Tsay (1988). All methods were implemented using R packages available from CRAN (Core Team 2017). Additional details about each spike detection method are provided in the following sections.

2.1 Summary of simulation

To compare each method's ability to identify spikes, we devised a collection of simulation studies parameterized based on violence rates in several California cities. The cities we selected had monthly violence rates that differed in terms of mean, variance, and autocorrelation.

Each city's simulation study included (1) simulating a series parameterized to be similar to city-level violence data, (2) adding a pre-specified magnitude and number of spikes to the series, (3) applying each detection method and calculating sensitivity (the probability of detecting a spike, given that one has occurred) and specificity (the probability of not detecting a spike, given that one has not occurred) of each method, and (4) iterating the procedure 1,000 times.

We conducted simulations with a range of spike numbers and magnitudes. The magnitudes considered were 10, 20, 30, 40 and 50 percent increases over the average rate during the study time period. The number of spikes inserted into the series were integers from 1 to 10. These variations allowed us to assess the ability of each method to detect different magnitudes of spikes and to determine whether the number of spikes present in the series influenced the performance of each method.

2.2 Data description

We selected the following cities in California for our study: Berkeley, Fresno, Oakland, Los Angeles, Richmond, Sacramento, San Diego, San Francisco, and Stockton. We selected these cities because they range in population size and their violence rates have a range of characteristics.

Interpersonal violence totals by city were created by summing the total number of deaths and injuries attributable to assault or homicide from the emergency department records and patient discharge and inpatient hospitalization records from the Office of Statewide Health Planning and Development (OSHPD) and the death records from Vital Statistics. To estimate monthly rates, we divided the number of cases in each city by the estimated number of people living in each city in each month and multiplied by 100,000. The population denominators came from the intercensal and postcensal population estimates from the U.S. Census Bureau. These data capture all assaults severe enough to require an emergency department visit or hospital stay and all homicides of California residents during 2005–2012.

2.3 Simulation details

Each city had its own simulation study. First, we fit the city's actual monthly violence series from 2005 to 2012 (96 time points) with an ARIMA model whose parameters were selected by the Aikake Information Criteria (AIC). We then simulated from this model in order to capture general properties of the series (such as mean, variance, autocorrelation, and trend).

2.4 Fitting initial ARIMA models

An ARMA model predicts current values of the response based on past values (autoregressive (AR) parameters) and innovations or past error values (moving average (MA) parameters). If differencing is required, the ARMA model is integrated and described as an ARIMA model. The standard equation of an ARIMA model is

$$y_t = \mu + \phi_1 y_{(t-1)} + \cdots + \phi_p y_{(t-p)} - \theta_1 \epsilon_{(t-1)} - \cdots - \theta_q \epsilon_{(t-q)}$$

which is commonly expressed as

$$\phi(B)(1 - B^d)y_t = \mu - \theta(B)\epsilon_t$$

where B is the backshift operator (where $B^d y_t = y_{(t-d)}$), ϕ and θ are polynomials of order p and q , respectively, d is the amount of differencing, and ϵ is a noise process with assumed distribution $N(\theta, \sigma^2)$.

The order of the AR portion is usually referred to as p and the order of the MA portion is usually referred to as q . An ARIMA(p,d,q) model has AR degree p , is differenced d times, and has MA degree q . We used *auto.arima* function from the *forecast* package in R, which uses AIC to select the model order (Hyndman and Khandakar 2007). The *auto.arima* function uses a unit root test with a null hypothesis of no unit root to pick the amount of differencing required. Once the proper level of differencing is determined, a stepwise algorithm is used to select the model order, which is described in detail in the documentation for the *forecast* package (Hyndman and Khandakar 2007). The model coefficients are estimated using maximum likelihood. The model parameters selected for each city are listed in Table 1.

Table 1: ARIMA models and parameters by city.

City	ARIMA model	Parameterization	Mean	SD	$\frac{\text{Mean}}{\text{SD}}$
Oakland	ARIMA(4,1,2)	$(1 - B)y_t = 0.8481y_{t-1} - 0.1436y_{t-2} + 0.3572y_{t-3} - 0.6178y_{t-4} + \epsilon_t - 1.6616\epsilon_{t-1} + 0.7814\epsilon_{t-2}$	79.28	11.79	6.72
San Diego	ARIMA(2,0,0)	$y_t = 0.3605y_{t-1} + 0.1875y_{t-1} + \epsilon_t$	30.72	3.20	9.60
Los Angeles	ARIMA(1,0,0)	$y_t = 0.436y_{t-1} + \epsilon_t$	35.53	3.40	10.45
Sacramento	ARIMA(2,0,1)	$y_t = -0.465y_{t-1} + 0.5180\epsilon_{t-2} + \epsilon_t + 0.9315\epsilon_{t-1}$	50.99	6.85	7.44
Fresno	ARIMA(1,1,1)	$(1 - B)y_t = 0.5306y_{t-1} + \epsilon_t - 0.9504\epsilon_{t-1}$	49.20	6.38	7.71
San Francisco	ARIMA(1,0,0)	$y_t = 0.3151y_{t-1} + \epsilon_t$	46.65	5.01	9.31
Stockton	ARIMA(1,0,0)	$y_t = 0.3440y_{t-1} + \epsilon_t$	54.58	7.90	6.91
Richmond	ARIMA(0,1,1)	$(1 - B)y_t = \epsilon_t - 0.7940\epsilon_{t-1}$	73.59	12.49	5.89
Berkeley	ARIMA(1,0,1)	$y_t = 0.8762y_{t-1} + \epsilon_t - 0.6531\epsilon_{t-1}$	28.41	6.84	4.15

2.5 Inserting spikes

We simulated from each city's fitted ARIMA model and inserted spikes randomly into the series, with all 96 time points equally likely to receive a spike. The number of spikes were varied from integers 1 to 10 and the magnitude of spikes considered were 10 %, 20 %, 30 %, 40 % and 50 % of the mean violence rate. Each combination of spike magnitudes and numbers were run as separate simulation studies and replicated 1,000 times.

2.6 Applying spike detection methods

A spike was considered correctly identified when a method identified a spike in violence at the time point where one was inserted by our algorithm. Any spike identified at a time point where we had not inserted a spike was considered to be incorrectly identified. We summarized the performance of each method using sensitivity and specificity. Descriptions of each method follow.

2.7 ARIMA

As described previously, an ARIMA model predicts the current value of the series using a linear combination of past values and innovations. The innovations are assumed to be $N(0, \hat{\sigma}^2)$. When using the ARIMA model to identify violence spikes, we fit the simulated series with an ARIMA model selected via AIC using the *auto.arima* package (Hyndman and Khandakar 2007). We then calculated the residuals and their standard deviation. Any time point with a residual value greater than two times the standard deviation of the residuals was identified as a spike. To assess whether the choice of critical value influenced the results, we replicated the analysis identifying spikes as values greater than 2.5 times the standard deviation of the residuals.

2.8 Kalman filter and smoother

The Kalman filter is a recursive data processing algorithm most famously used in engineering problems to predict trajectories based on position and velocity (Faragher 2012; Durbin and Koopman 2012). It uses the observed values of a series to update its predictions, resulting in a weighted average of the predicted and observed values where the weights are determined by the uncertainty around each. The Kalman filter uses a state space approach to time series modeling, and attempts to model a latent state that is unobserved but for which there are recorded measurements related to the state at discrete points in time. In this way, the state space approach is similar to a hidden Markov model in which the state space of the latent variable is continuous and the latent and observed variables are assumed to be Gaussian distributed.

The state space approach requires two equations, called the state equation and the observation equation. For details of the equations and derivations in this section, see Durbin and Koopman (2012).

The equation for the state of the system is

$$\alpha_t = T_t \alpha_{t-1} + R_t \eta_t$$

where α_t is the state of the system at time t , T_t is the transition matrix, which applies characteristics of the system at time $t-1$ to generate a prediction of the state at time t , and η_t is a vector of error terms with assumed distribution $N(0, Q_t)$. In our example, the matrix R_t is the identity matrix.

There is also an observation equation for the system, defined as

$$y_t = Z_t \alpha_t + \epsilon_t$$

where y_t is the measured value at time t , α_t is the true state at time t , Z_t is the matrix that maps the state to the measured value, and ϵ_t is a vector of measurement error terms assumed to be distributed $N(0, H_t)$.

The following matrices must be specified or estimated:

- T_t , the state-transition matrix, which maps the state at time $t-1$ to the state at time t
- Z_t , the observation matrix which maps the true state to the observed state at time t
- Q_t , the covariance of the state process at time t
- H_t , the covariance of the observation measurement at time t

The Kalman filter itself consists of the following recursion equations:

$$v_t = y_t - Z_t a_t$$

where v_t is the residual of the observed minus predicted value of the (latent) state,

$$a_t = E(\alpha_t | y_1, \dots, y_t - 1)$$

where a_t is the expected value of the state given past observations,

$$a_{t|t} = a_t + P_t Z_t F_t^{-1} v_t$$

where $a_{t|t}$ is the expected value of the state given the past and present observations,

$$P_t = \text{var}(\alpha_t | y_1, \dots, y_{t-1})$$

where P_t is the variance of the state given past observations,

$$F_t = \text{var}(v_t | y_1, \dots, y_{t-1}) = Z_t P_t Z_t + H_t$$

where F_t is the variance of the innovations, given past observations

$$a_{t+1} = T_t a_{t|t} = T_t a_t + K_t v_t$$

where a_{t+1} is the prediction for the value of the state at the next time point given past observations,

$$K_t = T_t P_t Z_t F_t^{-1}$$

where K_t , referred to as the Kalman gain, represents the change in the estimate of the state at time t after incorporating the information from the measurement,

$$P_{t|t} = P_t P_t Z_t F_t^{-1} Z_t P_t$$

where $P_{t|t}$ is the variance of the state given past and present observations, and

$$P_{t+1} = T_t P_t (T_t K_t Z_t) + R_t Q_t R_t$$

where P_{t+1} is the variance of the prediction for the state at the next time point, given past observations.

The smoothing algorithm combines the results from the Kalman filtering and does backward recursion to create new smoothed estimates for the state. This is done for each time t by combining both the estimate of the state at time t from the filtering process and the residuals at times $> t$. The state smoothing recursion equations are:

$$\begin{aligned}\hat{\alpha}_t &= a_t + P_t r_{t-1} \\ r_{t-1} &= Z_t F_t^{-1} v_t + L_t r_t \\ N_{t-1} &= Z_t F_t^{-1} Z_t + L_t N_t L_t \\ V_t &= P_t - P_t N_{t-1} P_t\end{aligned}$$

where r_t is a weighted combination of the innovations that occur after time $t-1$ and $L_t = T_t K_t = T_t - T_t P_t Z_t F_t^{-1}$.

We used the *KFAS* package in R to apply Kalman filtering and smoothing to the data (Helske 2014), which uses the algorithms from Durbin and Koopman (2012). For clarity, we have tried to stay as consistent as possible with the notation from these sources.

For our example, the latent state α_t is the violence level at time t and the observed data is the measured violence rate y_t at time t . To apply Kalman filtering and smoothing, we first fit an ARIMA model to the series, which was selected using the AIC, and transformed it into a state-space model formulation. The recursive filtering and smoothing algorithms described above were applied to the data, creating smoothed predictions for the violence level at each time point. Finally, we calculated residuals and identified spikes as any value greater than two times the standard deviation of the residuals. To assess whether the choice of critical value influenced the results, we replicated the analysis identifying spikes as values greater than 2.5 times the standard deviation of the residuals.

2.9 Wavelets

Wavelets are functions that oscillate around zero and satisfy $\int_{-\infty}^{\infty} \psi(x) dx = 0$ (Nason 2008). The wavelet transform is similar to a Fourier transform, in that for a given function $f(x)$, we can represent the function as a sum of orthonormal functions (Nason and Silverman 1994). However, instead of expressing the function as a sum of sine and cosine or complex exponential terms like in a Fourier series, wavelet analysis uses functions that are dilations and translations of a function called a mother wavelet ψ (Nason 2008).

For example, we can define a set of wavelets

$$\psi_{j,k}(x) = 2^{j/2} \psi(2^j x - k)$$

where j and k are integers, which have the effect of dilating ψ by a factor of 2^j and translating by $2^{-j}k$, and $2^{j/2}$ is a normalizing constant. These wavelets form an orthonormal set, which implies that a function $f(x)$ can be decomposed into a generalized Fourier series

$$f(x) = \sum_{j=-\infty}^{\infty} \sum_{k=-\infty}^{\infty} b_{j,k} \psi_{j,k}(x)$$

where j and k are integers that index the dilation number (j) and the translation number (k) (Nason 2008). The wavelet coefficients $b_{j,k}$ are the inner product of $f(x)$ and $\psi_{j,k}(x)$ such that

$$b_{j,k} = \int_{-\infty}^{\infty} f(x) \psi_{j,k}(x) dx$$

The wavelet transform is the process of representing the function in terms of the wavelets and their coefficients. The goal of using a wavelet transform is to calculate the coefficients that allow us to approximate the series. Thresholding the coefficients allows for different levels of smoothing. Soft thresholding shrinks the coefficients around the threshold. For a more complete description of wavelets, their properties and usage, (see Nason 2008; Nason and Silverman 1994).

We used the R package *wavethresh* to apply a wavelet transform to the data (Nason 2008, 2010). First a discrete wavelet transform was performed according to Mallat's pyramidal algorithm, and the coefficients were thresholded using a soft threshold. The wavelets were reconstructed by applying the inverse discrete wavelet transform. We calculated residuals and identified spikes as values greater than two times the standard deviation of the residuals. To assess whether the choice of critical value influenced the results, we replicated the analysis identifying spikes as values greater than 2.5 times the standard deviation of the residuals.

2.10 Outlier detection

We also tested the performance of an iterative model fitting and outlier detection procedure based on (Chang, Tiao, and Chen 1988; Tsay 1988).

The outlier detection process goes as follows:

1. Fit an ARIMA model to the observed time series.
2. Calculate the residuals and the variance of the residuals.
3. Compute the likelihood ratio test statistic λ_t for an additive outlier at time t

$$\lambda_t = \frac{\hat{\omega}_t}{\hat{\rho}_t \hat{\sigma}}$$

where

$$\hat{\omega}_t = \hat{\rho}_t^2 \left(1 - \sum_{i=1}^{n-t} \hat{\pi}_i \epsilon_t\right)$$

and

$$\hat{\rho}_t^2 = \left(1 - \sum_{i=1}^{n-t} \hat{\pi}_i^2\right)^{-1}$$

The π -weights are functions of the estimated coefficients of the ARIMA model, and can be expressed as $\pi(B) = \frac{\phi(B)}{\theta(B)}$. These weights can also be derived by representing the model as a recursive AR model.

4. Find the maximum in absolute value of the test statistics, and compare to a pre-specified critical value.

5. If the maximum exceeds the critical value, an outlier has been detected. Remove its effect by defining a new residual for the relevant time points and recalculate the residual variance. The adjusted residuals should have the form

$$\tilde{\epsilon}_t = \hat{\epsilon}_t - \hat{\omega}\hat{\pi}(B)I(t = T) \text{ for } t \geq T$$

where $I()$ is the indicator function.

6. Calculate new test statistics using the modified residuals and residual variance.
7. Continue until no more outliers are identified.
8. Assume the outliers identified above are known outlier times. Estimate the model parameters and the ω at each outlier time.
9. Use the parameter estimates from the models with assumed known outlier times and begin the outlier detection process again. Continue until no outliers are found.
10. If additional outliers are detected, re-estimate the parameters, incorporating the new outliers into the known outlier times.
11. Once no more outliers are identified, the procedure is complete. The locations of outliers (if any) in the time series have been identified and model parameters that exclude the effects of outliers have been estimated.

We used the *detectAO* function in the *TSA* R package to implement this outlier detection procedure (Chan and Ripley 2012).

2.11 Hypothesized performance across methods

The methods we selected have different strengths and are thus likely to perform best under different scenarios. We expect ARIMA modeling to perform best when there are few spikes and the data have simple parameterizations. When there are many spikes, the coefficients fit by the ARIMA model may be biased due to the spikes, and spike detection may suffer. While the Kalman filter and smoother assumes linear equations and Gaussian errors like the ARIMA models, the Kalman filter is more adaptive than ARIMA because of the extra uncertainty it incorporates in the measurement equation and the updating step. Therefore, we expect the Kalman filter to out-perform the ARIMA model in situations of high variance. However, because of the updating step, we expect the Kalman filter to have worse specificity as it may incorrectly adjust toward lower spike levels, obscuring their effects. In contrast to the ARIMA and Kalman methods, wavelets are useful when data have localized patterns, non-linearities, and discontinuities. Therefore, we expect these methods to perform well when the series are not well characterized by ARIMA models or when there are many spikes in a series. Outlier detection methods may perform similarly to the ARIMA and Kalman methods in situations with large shocks, but may have better performance with small shocks. The iterative procedure may capture small spikes better than other methods because they remove large spikes before searching for smaller magnitude spikes.

3 Results

The performance of all methods varied substantially by city and by series characteristics, although the patterns of performance were similar across magnitudes of spikes (Table 2, Table 3). Performance improved across all methods with increasing spike magnitude, as expected (Table 4, Table 5, and Figure 1). In general, the places in which methods had higher overall performance also tended to have a higher ratio of mean to standard deviation (Table 1).

Table 2: Average sensitivity of spike identification methods for spikes of magnitudes ranging from 10 to 50 % increase over series mean.

City	Sensitivity			
	ARIMA	Kalman	Wavelet	Outlier detection
Oakland	72.21	75.92	63.44	22.92

San Diego	72.45	74.73	64.72	46.86
Los Angeles	69.60	71.43	59.91	42.01
Sacramento	68.42	70.85	59.30	43.79
Fresno	63.60	67.38	57.07	30.08
San Francisco	66.14	67.14	55.85	38.01
Stockton	52.70	54.51	45.43	26.20
Richmond	46.06	48.59	42.94	7.66
Berkeley	38.25	40.07	35.48	13.13

Table 3: Average specificity of spike identification methods for spikes of magnitudes ranging from 10 to 50 % increase over series mean.

City	Specificity			
	ARIMA	Kalman	Wavelet	Outlier detection
Los Angeles	98.77	98.97	97.41	99.84
Oakland	98.68	98.84	97.03	99.54
San Diego	98.59	98.76	97.20	99.60
Fresno	98.28	98.73	97.55	99.65
San Francisco	98.56	98.68	97.23	99.76
Sacramento	98.37	98.56	97.22	99.51
Stockton	98.12	98.31	97.10	99.69
Richmond	97.89	98.14	96.85	99.79
Berkeley	97.26	97.38	96.33	99.23

Table 4: Average sensitivity of spike identification methods for spikes with magnitude 50 % of series mean.

City	Sensitivity			
	ARIMA	Kalman	Wavelet	Outlier detection
San Diego	95.26	96.08	89.50	74.53
Oakland	94.10	95.95	89.18	42.76
Los Angeles	94.16	95.01	87.20	70.90
Sacramento	93.75	94.82	86.66	71.42
San Francisco	93.01	93.97	85.20	68.68
Fresno	89.67	91.57	83.77	50.63
Stockton	85.07	86.18	76.12	52.19
Richmond	78.49	80.10	74.78	17.00
Berkeley	75.03	76.13	68.05	28.72

Table 5: Average specificity of spike identification methods for spikes with magnitude 50 % of series mean.

City	Specificity			
	ARIMA	Kalman	Wavelet	Outlier detection
Los Angeles	99.72	99.76	98.27	99.95
San Diego	99.64	99.69	98.18	99.86
San Francisco	99.60	99.63	98.12	99.92
Oakland	99.53	99.61	98.20	99.64
Fresno	99.43	99.60	98.40	99.86
Sacramento	99.44	99.52	98.21	99.83
Stockton	99.24	99.32	97.98	99.89
Richmond	98.99	99.15	97.84	99.85
Berkeley	98.26	98.37	97.30	99.59

Figure 1: ROC curve for spike detection across all simulations, by method of detection and magnitude of spikes.

Overall, the Kalman method had the highest sensitivity and the outlier detection method had the highest specificity (Figure 1). However, the outlier detection method was by far the worst performer in terms of sensitivity. The Kalman method had both high sensitivity and specificity, and was consistently the best performer with an average sensitivity of 89.98 % and an average specificity of 99.41 % across cities in simulations with 1 to 10 spikes inserted and spikes with magnitude 50 % increase over the series mean (Table 4). Sensitivity was lower for spikes of lower magnitude, which makes sense given these are much more likely to blend into the background variation (Figure 1). Averaging over all simulation scenarios, in which we included spikes of all magnitudes (10 %, 20 %, 30 %, 40 %, and 50 % increase over the series mean) and numbers of spikes from 1–10, the Kalman filter and smoother had an average sensitivity of 63.40 % and specificity of 98.49 % across cities (Table 2, Figure 1). The ARIMA method was a close second best performer, with an average sensitivity of 61.05 % and specificity of 98.28 % across all simulations.

All methods performed worst in the simulation parameterized to be similar to Berkeley, which is likely due to the high relative variance (Table 1), which made distinguishing spikes from background variation difficult. However, the Kalman filter method was still the best performer, with an average sensitivity of 40.07 % and specificity of 97.38 % for all magnitude spikes and across simulations with 1 to 10 spikes (Table 2, Table 3). Among the highest magnitude spikes (50 % increase over Berkeley series mean), the Kalman method had average sensitivity of 76.13 % and specificity 98.37 % (Table 4, Table 5).

In order to assess how performance changed based on the critical values specified as part of each detection method, we also ran simulations increasing the threshold values. For the Kalman filter and smoother, ARIMA, and wavelet methods, this meant using a critical value of 2.5 standard deviations of the residuals rather than 2. There was a small decrease in sensitivity and increase in specificity, but the results were not substantively different.

We applied each of the methods to the true violence series for each of the nine cities to illustrate their applications (Figure 2–Figure 10). In general, the Kalman filter and smoother was able to follow the true series much more closely than the other methods, which likely improved its performance in the simulations. The ARIMA method sometimes detected the same spikes as the Kalman method (Figure 4, Figure 6), but not always. The wavelet method created a much smoother fit of the series, which resulted in identification of spikes that were not identified by the Kalman or ARIMA, especially near the beginning or end of the series (Figure 3, Figure 4, Figure 5, Figure 7, Figure 8, and Figure 9). The outlier detection method did not produce fitted values for the series because it relied on calculation of test statistics to identify outliers, and did not identify spikes in violence in any of the cities. Because the Kalman filter and smoother had the best performance in the simulations, we listed the months with spikes detected by the Kalman method in Table 6.

Figure 2: Spikes detected by method in monthly violence rate Berkeley, CA from 2005 to 2012.

Figure 3: Spikes detected by method in monthly violence rate Fresno, CA from 2005 to 2012.

Figure 4: Spikes detected by method in monthly violence rate Los Angeles, CA from 2005 to 2012.

Figure 5: Spikes detected by method in monthly violence rate Oakland, CA from 2005 to 2012.

Figure 6: Spikes detected by method in monthly violence rate Richmond, CA from 2005 to 2012.

Figure 7: Spikes detected by method in monthly violence rate Sacramento, CA from 2005 to 2012.

Figure 8: Spikes detected by method in monthly violence rate San Diego, CA from 2005 to 2012.

Figure 9: Spikes detected by method in monthly violence rate San Francisco, CA from 2005 to 2012.

Figure 10: Spikes detected by method in monthly violence rate Stockton, CA from 2005 to 2012.

Table 6: Cities and months with violence spikes detected by the Kalman filter and smoother detection method.

City	Months with Spikes Detected
Los Angeles	2005–07, 2006–07
Oakland	2007–03, 2009–09
San Diego	2005–07, 2007–05, 2007–10
Berkeley	2006–07, 2008–05, 2009–05, 2012–08
Fresno	2008–08, 2012–05
Sacramento	2009–08, 2010–08
San Francisco	2005–09, 2006–05, 2006–07, 2011–08, 2012–10
Stockton	2005–05, 2008–08
Richmond	2006–07, 2007–07, 2010–03, 2012–07

4 Discussion

This work contrasted four methods that can be used to identify spikes in time series. We found that applying a Kalman filter and smoother and identifying values whose residuals were greater than two standard deviations from the predicted value was the best performing method in simulations parameterized based on violence time series from nine cities. The features of the simulations varied with respect to the mean, variance, autocorrelation, and trend, suggesting that across these features Kalman consistently does best. However, series of other data or violence series substantially different on these features might lead to different relative performance. The code provided in the supplemental material facilitates extension of this work to other series of interest that may have different features.

The results are in accordance with our hypotheses that the Kalman filter and smoother would outperform the ARIMA when there are high magnitude spikes. However, while we thought the Kalman method would

not perform as well with low magnitude spikes, it was the best performer for every magnitude and number of spikes. This method is easily implemented via the *KFAS* package in R (Helske 2014). Using this method, we found at least one spike in the monthly violence rate for each city. While the Kalman filter method was very successful in detecting large spikes, it inconsistently identified smaller magnitude spikes.

The causes of the spikes in violence identified in this analysis are outside the scope of this study. However, it is interesting to note that all but two of the violence spikes detected by the Kalman method occurred between the months of May and October. This suggests explanations for the spikes in violence could include exacerbations of gang activity, muggings, robberies, or civil unrest that are generally higher in the summer months.

There are some limitations to our approach. It is possible that simulating from ARIMA models to characterize the violence series does not completely capture the true underlying violence distributions. This may bias our results toward methods that rely on series well-described by ARIMA models. However, we would be more concerned about this if the ARIMA method were performing best in each scenario. While the Kalman filter also assumes linear functions and Gaussian errors, it incorporates a model for measurement error, which improves its ability to differentiate between normal variance and spikes in the series.

It is also likely there are spikes in the violence series used to parameterize the simulations. This may affect the assumed data generating distribution for each city. However, we expect any bias in the parameter estimates will not affect the overall behavior of the simulated series. Repeating each simulation 1,000 times should incorporate sufficient variability around the true data generating mechanism so that small bias in the parameters should not affect the results.

This simulation study compared the performance of several methods for identifying spikes in time series. Among the methods considered, our results suggest that applying a Kalman filter and smoother and identifying spikes as values above two standard deviations of the residuals is the best choice to identify additive outliers or spikes in time series. This finding was determined based on its performance with respect both sensitivity and specificity. We also found that the performance was best in cities that had low relative variance. Kalman filtering and smoothing is straightforward to apply using standard statistical software and should be considered by other researchers interested in the effect of exposure spikes on a response.

Funding

This work was supported by National Institutes of Health (Funder Id: <http://doi.org/10.13039/1000000002>, Grant Number: DP2 HD080350).

References

- Ahern, J., Cerdá, M., Lippman, S. A., Tardiff, K. J., Vlahov, D., and Galea, S. (February 2013). Navigating non-positivity in neighbourhood studies: An analysis of collective efficacy and violence. *Journal of Epidemiology and Community Health*, 67(2):159–165.
- Ahern, J., Matthay, E. C., Goin, D. E., Farkas, K., and Rudolph, K. E. (2018). Acute changes in community violence and increases in hospital visits and deaths from stress-responsive diseases. *Epidemiology*, 29(5):684–691.
- Apter, A. J., Garcia, L. A., Boyd, R. C., Wang, X., Bogen, D. K., and Ten Have, T. (September 2010). Exposure to community violence is associated with asthma hospitalizations and emergency department visits. *The Journal of Allergy and Clinical Immunology*, 126(3):552–557.
- Baird, S., Friedman, J., and Schady, N. (2011). Aggregate income shocks and infant mortality in the developing world. *Review of Economics and Statistics*, 93(3):847–856.
- Balke, N. S. (1993). Detecting level shifts in time series. *Journal of Business & Economic Statistics*, 11(1):81–92.
- Besculides, M., Heffernan, R., Mostashari, F., and Weiss, D. (2005). Evaluation of school absenteeism data for early outbreak detection, New York city. *BMC Public Health*, 5(1):105.
- Bhattacharya, J., DeLeire, T., Haider, S., and Currie, J. (2003). Heat or eat? Cold-weather shocks and nutrition in poor American families. *American Journal of Public Health*, 93(7):1149–1154.
- Brown, J. R., Hill, H. M., and Lambert S. F. (2005). Traumatic stress symptoms in women exposed to community and partner violence. *Journal of Interpersonal Violence*, 20(11):1478–1494.
- Chan, K. S., and Ripley, B. (2012). *Tsa: Time series analysis*. R Package Version 1.01. URL: <http://CRAN.R-project.org/package=TSA>.
- Chang, I., Tiao, G. C., and Chen, C. (1988). Estimation of time series parameters in the presence of outliers. *Technometrics*, 30(2):193–204.
- Chen, C., and Liu, L.-M. (1993). Joint estimation of model parameters and outlier effects in time series. *Journal of the American Statistical Association*, 88(421):284–297.
- Clark, C., Ryan, L., Kawachi, I., Canner, M. J., Berkman, L., and Wright, R. J. (January 2008). Witnessing community violence in residential neighborhoods: A mental health hazard for urban women. *Journal of Urban Health*, 85(1):22–38.
- Cryer, J. D., and Chan, K.-S. (2008). *Time Series Analysis*. Springer Texts in Statistics. New York, NY: Springer New York.
- Curry, A., Latkin, C., and Davey-Rothwell, M. (July 2008). Pathways to depression: The impact of neighborhood violent crime on inner-city residents in Baltimore, Maryland, USA. *Social Science & Medicine*, 67(1):23–30.

- Dafni, U. G., Tsiodras, S., Panagiotakos, D., Gkolfinopoulou, K., Kouvatsas, G., Tsourti, Z., and Saroglou, G. (2004). Algorithm for statistical detection of peaks—syndromic surveillance system for the athens 2004 olympic games. *MMWR Morb Mortal Wkly Rep*, 53(Suppl):86–94.
- De Jong, P., and Penzer, J. (August 2000). Diagnosing shocks in time series. Technical Report, London School of Economics.
- Durbin, J., and Koopman, S. J. (2012). *Time Series Analysis by State Space Methods*, Volume 38. Oxford: OUP.
- Faragher, R. (2012). Understanding the basis of the kalman filter via a simple and intuitive derivation [lecture notes]. *IEEE Signal Processing Magazine*, 29(5):128–132.
- Helske, J. (2014). Kfas: Kalman filter and smoothers for exponential family state space models. R Package Version, 1:4–1.
- Helske, J. (2017). KFAS: Exponential family state space models in R. *Journal of Statistical Software*, 78(10):1–39.
- Humphreys, D. K., Panter, J., Sahlqvist, S., Goodman, A., and Ogilvie, D. (2016). Changing the environment to improve population health: A framework for considering exposure in natural experimental studies. *Journal of Epidemiology and Community Health*, 70(9):941–946.
- Hyndman, R., Athanasopoulos, G., Bergmeir, C., Caceres, G., Chhay, L., O’Hara-Wild, M., Petropoulos, F., Razbash, S., Wang, E., and Yasmeen, F. (2019). *Forecast: Forecasting functions for time series and linear models*, 2019. R Package Version 8.5.
- Hyndman, R. J., and Khandakar, Y. (2008). Automatic time series forecasting: The forecast package for R. *Journal of Statistical Software*, 26(3):1–22.
- Hyndman, R. J., Khandakar, Y. (2007). Automatic time series for forecasting: The forecast package for R. Technical Report, Monash University, Department of Econometrics and Business Statistics.
- Jacob, B., Lefgren, L., and Moretti, E. (2007). The dynamics of criminal behavior evidence from weather shocks. *Journal of Human Resources*, 42(3):489–527.
- Kane, R. J. (2011). The ecology of unhealthy places: Violence, birthweight, and the importance of territoriality in structurally disadvantaged communities. *Social Science & Medicine*, 73(11):1585–1592.
- Koppensteiner, M. F., and Manacorda, M. (2016). Violence and birth outcomes: Evidence from homicides in brazil. *Journal of Development Economics*, 119:16–33.
- Margerison-Zilko, C. E., Catalano, R., Hubbard, A., and Ahern, J. (2011). Maternal exposure to unexpected economic contraction and birth weight for gestational age. *Epidemiology (Cambridge, Mass.)*, 22(6):855.
- Martinez, P., and Richters, J. E. (February 1993). The NIMH community violence project: II. Children’s distress symptoms associated with violence exposure. *Psychiatry*, 56(1):22–35.
- Masi, C. M., Hawkey, L. C., Harry Piotrowski, Z., and Pickett, K. E. (December 2007). Neighborhood economic disadvantage, violent crime, group density, and pregnancy outcomes in a diverse, urban population. *Social Science & Medicine*, 65(12):2440–2457.
- Metcalfe, A. V., and Cowpertwait, P. S. P. (2009). *Introductory Time Series with R*. New York, NY: Springer New York.
- Michael Oakes, J. (June 2006). Commentary: Advancing neighbourhood-effects research—selection, inferential support, and structural confounding. *International Journal of Epidemiology*, 35(3):643–647.
- Miguel, E., Satyanath, S., and Sergenti, E. (2004). Economic shocks and civil conflict: An instrumental variables approach. *Journal of Political Economy*, 112(4):725–753.
- Miller, L. S., Wasserman, G. A., Neugebauer, R., Gorman-Smith, D., and Kamboukos, D. (March 1999). Witnessed community violence and antisocial behavior in high-risk, urban boys. *Journal of Clinical Child Psychology*, 28(1):2–11.
- Nason, G. P., editor. (2008). *Wavelet Methods in Statistics with R*. New York, NY: Springer New York.
- Nason, G. P. (2010). Wavethresh: Wavelets statistics and transforms. R Package Version, 4:163.
- Nason, G. P., and Silverman, B. W. (1994). The discrete wavelet transform in s. *Journal of Computational and Graphical Statistics*, 3(2):163–191.
- R Core Team. (2017). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Tsay, R. S. (January 1988). Outliers, level shifts, and variance changes in time series. *Journal of Forecasting*, 7(1):1–20.
- Unkel, S., Paddy Farrington, C., Garthwaite, P. H., Robertson, C., and Andrews, N. (2012). Statistical methods for the prospective detection of infectious disease outbreaks: A review. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 175(1):49–82.
- Wilde, J., Apouey, B. H., and Jung, T. (2017). The effect of ambient temperature shocks during conception and early pregnancy on later life outcomes. *European Economic Review*.

Supplementary Material: The online version of this article offers supplementary material (DOI:<https://doi.org/10.1515/em-2018-0005>).