TRR
JOURNAL OF THE TRANSPORTATION RESEARCH BOARD

Non-Stationary Time Series Model for Station-Based Subway Ridership During COVID-19 Pandemic: Case Study of New York City

Bahman Moghimi[1] (ID)**, Camille Kamga**[1]**, Abolfazl Safikhani**[2]**,
Sandeep Mudigonda**[3] (ID)**, and Patricio Vicuna**[1]

Abstract
The COVID-19 pandemic in 2020 has caused sudden shocks in transportation systems, specifically the subway ridership patterns in New York City (NYC), U.S. Understanding the temporal pattern of subway ridership through statistical models is crucial during such shocks. However, many existing statistical frameworks may not be a good fit to analyze the ridership data sets during the pandemic, since some of the modeling assumptions might be violated during this time. In this paper, utilizing change point detection procedures, a piecewise stationary time series model is proposed to capture the nonstationary structure of subway ridership. Specifically, the model consists of several independent station based autoregressive integrated moving average (ARIMA) models concatenated together at certain time points. Further, data-driven algorithms are utilized to detect the changes of ridership patterns as well as to estimate the model parameters before and during the COVID-19 pandemic. The data sets of focus are daily ridership of subway stations in NYC for randomly selected stations. Fitting the proposed model to these data sets enhances understanding of ridership changes during external shocks, both in relation to mean (average) changes and the temporal correlations.

During outbreaks caused by infectious diseases, providing an accurate and reliable subway station ridership model is crucial for transit operators and passengers. Subway ridership changes from time to time during a year, and a statistical model that can capture this variability would enable transit authorities to plan for appropriate resource allocation, obtain up-to-date train service frequencies for the expected ridership, and send required information to passengers during unexpected events. At the same time, with a good prediction of passenger ridership, commuters can adjust their departure times and choose other modes to reduce delays and improve their comfort.

Unforeseen and abrupt changes in urban mobility could happen for many reasons, such as crowding at the time of special events, severe weather conditions, natural disasters, infrastructure-related constructions, or global health crises like the COVID-19 pandemic. The pandemic generally reduced the overall efficiency of human mobility networks, which caused an exogenous shock to the economy and is functioning similar to a natural disaster (*1*).

The first case of the coronavirus in the State of New York in New York City (NYC), U.S., was confirmed on March 1, 2020, and from that date panic buying for food

[1]Department of Civil Engineering, City College of New York, NY
[2]Department of Statistics, University of Florida, Gainesville, FL
[3]Region-2 University Transportation Research Center, New York, NY

Corresponding Author:
Bahman Moghimi, smoghim000@citymail.cuny.edu

and household products was reported. On March 7, New York Governor Andrew Cuomo declared a state of emergency in the State of New York. A second case in the state was announced on March 10, which was the first known case in the state to be caused through community spread. On March 16, Governor Cuomo issued an executive order to close all public and private schools throughout the state, initially ordered to last until April 1, but later extended. On March 22, the New York State stay-at-home order took effect. The coronavirus cases in New York increased tremendously to the point that on March 22, the New York Times reported that NYC became the epicenter of the pandemic (*2*). As New Yorkers worked together to flatten the curve, the reopening strategies came into effect in June to let businesses to start reopening by phases. NYC entered Phase 1 of the reopening on June 8 and Phase 2 on June 22, 2020.

The coronavirus pandemic drastically altered the transportation choice behavior of all transportation users. Many travelers preferred or were obliged to stay home and work from home to comply with stay-at-home orders enacted by government officials (*3*). Those with essential tasks/jobs would have commuted to their work location under some orders enacted to slow the spread of the virus. To reduce the spread of the highly contagious virus, health professionals encouraged only essential travels. As a result, vehicular traffic volumes on roadways were very low. Thus, the vehicular speeds in highway and corridors increased significantly, and vehicle travel times and delays decreased when compared with pre-COVID times.

Public transit ridership was severely affected by the measures enacted to slow the spread of COVID-19. In NYC, transit ridership dropped on all transit modes operated by the Metropolitan Transportation Authority (MTA) which include the Subway, New York City Transit Bus, MTA Bus, Metro North Railroad, and Long Island Railroad. Specifically, for the subway ridership, the percentage change on March 12, March 16, March 23, April 2, April 17, May 15, and June 30, 2020, were reported at 19%, 60%, 87%, 92%, 93%, 90%, and 80% ridership declines, respectively, as compared with similar dates in 2019 (*4*).

To accurately forecast transportation data, and, more specifically, the subway ridership, such a drop in the observed time series should be considered. Many of the existing travel demand and transit ridership models are not trained to include anomalies in their models. Developing a rigorous model that can consider ridership variability and be resilient enough to take care of anomalies is the focus of this paper. In this research, an interpretable time series model to formulate subway ridership during the COVID-19 pandemic is developed. Then, statistically sound detection methods are utilized to detect

the anomalies of subway ridership time series data. It is a common practice to model data with temporal index using stationary time series models, and autoregressive integrated moving average (ARIMA) models are among the most well-known stationary time series models used in different scientific fields, including civil engineering and transportation. However, in the presence of structural breaks (shocks) in the temporal dynamical system under consideration, the stationarity assumption may be violated; thus, it is necessary to search for alternative modeling frameworks. Piecewise stationary models are interesting models which are easy to interpret since a time point at which the dynamical system receives an external shock can be called a "break point," and its location can be estimated using developed algorithms. Such time points are essentially the discontinuity points in the piecewise modeling framework.

An interesting fact about the proposed detection method used in this paper is that it can handle detection for second-order structure (autocorrelation). The difference between the fitted models before and after the change/break points shows that the dependence structure has changed significantly. Such changes need to be accounted when forecasting ridership. This area of research is relatively undeveloped with very little literature focused on the topic. In short, one of the main contributions of this research is to introduce a piecewise stationary time series model which can be utilized to model sequential data experiencing external shocks. The proposed modeling framework is applied to ridership data during the COVID-19 pandemic, which, to the best of the authors' knowledge, has been introduced for the first time to transportation datasets. To apply this modeling framework, a novel statistical algorithm is developed and utilized to detect break/change points, and, further, fitted models before and after each break point are summarized. Certain goodness of fit tests are applied to the residuals to confirm the satisfactory performance of the proposed methodology. In the next section, a brief review of the methods found in the literature is provided.

Literature Review

Subways play a crucial role in today's urban mobility. For many urban travelers, they are the convenient and first choice of mobility in large metropolitan areas. Developing a long-term transit ridership prediction is a product that may result from conventional four-step travel demand forecasting (*5*). One way to do this is to fit regression models based on many contributing factors such as transit attributes, geographic information, demographic, and economic factors (*5*). For example, using regression models, Singhal et al. analyzed the impact of weather on New York City Transit subway ridership

based on day of week and time of day combinations (6). They found that the impact of weather on transit ridership varies based on the time period and the location of subway stations.

Many machine learning algorithms were developed to model transit ridership. Liu and Chen developed a deep learning method to predict the ridership at four BRT stations in Xiamen, China (7). They used a three-stage hybrid deep network model centered on hourly basis (7). Also, Liu et al. proposed a multilayer deep-learning architecture to predict metro inbound/outbound passenger flow (8). Other machine learning models to predict subway station ridership were also developed, including gradient boosting decision trees, support vector machines, network Kriging method, and more (9–11). On the other hand, time series models have also been utilized in transportation-related problems as well as to model the subway ridership. For example, time series ARIMA models have been applied to many areas of transportation including traffic arrival demand modeling, seasonal variation of freeway traffic conditions, prediction of actuated signal cycle length, traffic speed modeling on a downstream link, and so forth (12–15). Also, time series models have been applied to predict transit ridership: for instance, Ding et al. presented an ARIMA-generalized autoregressive conditional heteroscedasticity (GARCH) time series model to predict short-term metro ridership, which is an ARIMA model that takes care of the deterministic part and the non-linear GARCH model for the stochastic part (16). Their proposed ARIMA-GARCH model was applied on three stations' ridership in Beijing and the result showed that it outperformed other proposed models (16).

Many of the existing time series models in the field of transportation have focused on stationary or smoothly varying models. However, these models cannot be used when there is a dramatic change in the system that could affect the time series model. For example, when Hurricane Sandy occurred in NYC in 2012, a sudden change of 11% decline happened in the number of trains in service on an average day (17). This major disruptive event could affect the future growth of transit ridership. Other disruptions include transit strikes, bridge closures, special events like the Olympic Games, and earthquakes—the impact of these events on transit ridership and mode choice were reviewed in the study by Zhu and Levinson (18). For example, considering transit strikes, a study of 13 cases between 1966 and 2000 in the U.S. and Europe showed that they caused short-term and long-term losses on transit ridership (19). Events like transit strikes provide a unique opportunity to study the change in transit ridership and travelers' behaviors, both of which are significant for drafting future transit policies. Stationary statistical models are not appropriate in the presence of anomalies caused by an unforeseen event. A global pandemic can be considered as an unforeseen event that significantly violates the stationary assumption of time series. Change point detection methods could be a remedy to unravel this issue in time series models.

To find a change point detection, it is possible to heuristically approximate locations of break points and estimate parameters within each (heuristically) estimated stationary segment. However, there are three main advantages in using a data-driven method to perform break point detection. First, in general, it is not clear whether an external shock would necessarily yield to a discontinuity point in the model. Therefore, it is possible to think that a time point should be considered as a break point while there are no break points in the data. Second, an important assumption in change point detection analysis is that the total number of break points is unknown and must be estimated from data. Therefore, an estimate on the number of change points in the data should first be obtained, and then an attempt should be made to approximate location of break point(s) using either heuristic methods or data-driven ones. However, estimating the number of change points in the data is a complicated task and there is no clear way how to perform it heuristically. Third, estimating the location of break points is a complicated task for heuristic methods, since there may be a delay from the time of external shock to real changes in the parameters in the data, or it may be the case that model parameters start to change before the external shock is known to humans. As a result, heuristic methods may estimate the location of break points with a certain error, which may yield to inaccurate estimation of model parameters. Therefore, data-driven methods can estimate the number of change points and their locations more accurately than heuristic methods.

Data-driven methods to perform change point detection have recently been applied to the transportation research domain. Tang and Gao developed a nonparametric model for traffic flow prediction utilizing anomaly detection, while Tsiligkaridis and Paschalidis applied anomaly detection to detect traffic jams (20, 21). Margarieter leveraged anomaly detection methods in her research to detect incidents (22). Riverio et al. presented an analytical framework to detect anomalous events to a large real road traffic dataset collected from various areas in Europe (23). An anomaly-detection-based analytical module was developed by Zhang et al. to visualize abnormal passenger traffic flow on an urban network (24). In this paper, the authors model the subway station ridership during the COVID-19 pandemic by considering specific time series models combined with possible change points. Next, this rich family of time series models is introduced.

Time Series Models

Time series models are very well-known statistical models designed for data sets indexed by time. Time series models have been employed in different disciplines, including finance, water resources, climate change, transportation, and so forth. (14, 25, 26). The main objective of time series models is to grasp the underlying behavior of the dataset over time, and then to detect the dependence among such data points to forecast the future. More specifically, in these models, the hope is to find the linear dependency structure among the data points over time, and then use that dependency to predict the new points in the future. This research focuses on applying autoregressive integrated moving average (ARIMA) models. The ARIMA model is a well-established yet simple statistical tool to deal with data indexed by time (27). In the next section, a short introduction of this family of models is provided.

ARIMA Model

The ARIMA model is a powerful statistical tool when univariate time series or one data set are being dealt with. Suppose there is a data set $X_1, X_2, \ldots, X_n$, which is observed through time. For example, X_1 is the observation at the first time point, X_2 is the observation at the second time point, and so on. The autoregressive moving average (ARMA) model assumes that the present value of a time series is a linear combination of its past observations together with a linear combination of noises in the past observations. Thus, the time series X_t is called $ARMA(p, q)$ as follows:

$$X_t - \phi_1 X_{t-1} - \ldots - \phi_p X_{t-p} = Z_t + \theta_1 Z_{t-1} + \ldots + \theta_q Z_{t-q} \tag{1}$$

where
$\phi_1, \phi_2, \ldots, \phi_p = AR$ constants,
$\theta_1, \theta_2, \ldots, \theta_q =$ moving average or MA constants, and
$Z_t =$ white noise (WN) with mean 0 and variance $\sigma^2 (WN(0, \sigma^2))$.

In the stated model, the current value of a data point stems from the past p observations through ϕ_i's and the past q observation noises through θ_i's. ARMA models are stationary models. In stationary models, the covariance between two observations X_t and X_{t+h} depends only on the lag h and not on the time t. To put it simply, in a stationary model, the dependence structure of points exits on their distances and not on their locations (27). ARIMA models are an extension of ARMA models to capture the non-stationary behavior in the model. ARMA models are generally denoted by $ARIMA(p, d, q)$, where p is the order (number of time lags) of the autoregressive part, d is the degree of differencing (the difference between the current observations and d time lags in the past is calculated as the new data set), and q is the order for the moving-average part. Note, $ARIMA(p, 0, q)$ is also an $ARMA(p, q)$ model; $ARIMA(p, d, q)$ will be non-stationary when $d \neq 0$.

To measure the degrees of dependency among data points at different times, an autocorrelation function (ACF) is used. ACF at lag h is calculated as

$$\rho(h) = \frac{\gamma(h)}{\gamma(0)} \tag{2}$$

where
$\gamma(h) =$ the autocovariance function (ACVF) at lag h and is defined as:

$$\gamma(h) = Cov\ (X_t,\ X_{t+h}), \tag{3}$$

where
Cov $=$ the covariance of two datasets.

The ACF at lag h is the normalization version of ACVF. The partial auto-correlation function (PACF) provides the partial correlation of a time series with its own lagged values, controlling for the values of the time series in between them. For the time series X_t, the *PACF* at lag h denoted by $\alpha(h)$ is the auto-correlation between X_t and X_{t+h} given the points $X_{t+1}, \ldots, X_{t+h-1}$. ACF and PACF are important factors in estimating the ARIMA model's parameters, p, d, q, which will be used to predict the future values.

The main assumption for parameters in ARIMA models is that the roots of the auto-regressive polynomial should not be on the unit circle, that is, should not have a complex norm of one. Indeed, all estimated ARIMA models satisfy this property. This ensures that each model is stationary/stable. Moreover, to ensure the model is causal (invertible), all roots of the auto-regressive polynomial (moving average polynomial) should be outside of the unit disc, that is, should have a complex norm greater than one (see more details in Brockwell and Davis) (27). This assumption for the case of p = 1 or q = 1, that is, Autoregressive (AR) (1) or moving average (MA) (1) model, is equivalent to the phi/theta parameter being less than one. However, for higher lags (p or q more than 1), it is possible to have phi or theta parameters which have magnitudes more than one.

A piecewise stationary ARIMA consists of several different and independent ARIMA models concatenated at certain time points called change points (or break points). In general, the number of change points and their locations is unknown and they need to be estimated using statistical techniques. In the next section, a brief description of one such procedure to detect change points is provided.

Change Point Detection

Change point detection is an active line of research in statistics, specifically in the field of time series analysis, with applications in many scientific fields such as economics and health sciences. The main objective of this line of research is to find time point(s) at which the parameters of the data generating process have changed. This change may be in the mean, variance, covariance structure, or spectral density of the time series. In this paper, the procedure developed in Safikhani and Shojaie is leveraged to detect the time when the change or break point occurs in the data sets under investigation (28).

The detection algorithm developed in Safikhani and Shojaie has three main steps—the first two are for detection purposes and the last step is for model parameter estimation within each stationary segment (28). The first step assumes every time point is a (potential) break point. This is done by expanding the model parameter space. Then, the fused lasso penalty combined with the least squares objective function is utilized to find a set of candidate change points by reformulating the detection problem as a variable selection problem in a high-dimensional regression model. This step ensures that no true break point remains isolated, that is, there will be at least one estimated break point close to any true break point with high probability (under certain regularity conditions). On the other hand, the total number of candidate change points selected in the first step may be larger than the true number of break points. To that end, a screening step is added (step 2) to search over all candidate change points and only keep the ones which would reduce the combined mean squared error (MSE) on the left- and right-hand side of the estimated break point significantly compared with MSE by ignoring the estimated break point. This screening step is shown to remove the redundant break points estimated in the first step with high probability, therefore the final estimated change points are consistent estimates for location of break points and the number of estimated change points is a consistent estimate for the true number of change points asymptotically. Finally, the third step estimates the model parameters in all stationary segments by applying penalized estimation techniques using lasso in high dimensions (or the simple least squares method in low dimensions).

The developed algorithm in Safikhani and Shojaie works for general multivariate vector auto-regressive models, and thus is a perfect fit to this modeling framework (28). This algorithm can handle detection for autocorrelation (the second-order structure). The basic idea is to first assume every time point could potentially be a break point and fit a smooth version of time-varying parameters to the data. This step is performed using the fused lasso technique by Rinaldo (29). Specifically, the

time series model can be written as a linear regression model with a very large design matrix, as:

$$Y = X\boldsymbol{B} + E, \tag{4}$$

where

Y = vector of currently observed data points,

X = a lower-triangular design matrix consisting of past values of time series (up to p lags), and

E = the usual measurement error, which is generally assumed to be independent and identically distributed (i.i.d) while this assumption can be relaxed (see e.g., Safikhani and Shojaie [28]).

The unknown parameter vector $\boldsymbol{B}$ can be estimated by minimizing a double-regularized least squares:

$$\hat{\mathbf{B}} = argmin_{\mathbf{B}} \ \frac{1}{n} \| \boldsymbol{Y} - \boldsymbol{Z}\mathbf{B} \|_2^2 + \lambda_{1,\,n}\|\mathbf{B}\|_1 + \lambda_{2,\,n} \sum_{k=1}^{n} \left\| \sum_{j=1}^{k} \beta_j \right\|_1 . \tag{5}$$

where

$\| \ \|_2$ = the Euclidean norm of a vector,

$\| \ \|_1$ = the sum of absolute values of elements in a vector,

n = the number of time points in the data,

β_j = the jump parameter at time point j,

$\sum_{j=1}^{k} \beta_j$ = the model parameter at time point k, and

$\lambda_{1,\,n}$ and $\lambda_{2,\,n}$ = tuning parameters, and are non-negative sequences of real numbers which are selected using data-driven techniques.

The first term in the objective function in Equation 5 is the MSE term, followed by two l_1 penalties controlling the number of break points and the sparsity of the auto-regressive parameters. The first one (i.e., $\lambda_{1,\,n}\|\mathbf{B}\|_1$) is the main component, since it fuses the differences between estimated parameters to ensure smoothly estimated parameters over time, and it is called fused lasso (30, 31, 32).The second one (i.e., $\lambda_{2,\,n} \sum_{k=1}^{n} \left\| \sum_{j=1}^{k} \beta_j \right\|_1$) is needed only in the high-dimensional regime where the number of time series components is much larger than the number of time points at which the data was observed. The two tuning parameters $\lambda_{1,\,n}$ and $\lambda_{2,\,n}$ can be approximated using rolling-window-type cross-validations (24). Specifically, since the focus is on univariate modeling in this paper, it is simple to put $\lambda_{2,\,n} = 0$ which does not violate the assumptions in Safikhani and Shojaie (28). The optimization problem (5) is convex and can be solved efficiently. In fact, by Proposition 1 of the Friedman et al. problem (5) can be solved by first finding a solution $\hat{\mathbf{B}}^{(0)}$ for $\lambda_{2,\,n} = 0$ and then applying element-wise soft-thresholding to $\hat{\mathbf{B}}^{(0)}$ to obtain the final estimate

$\hat{\mathbf{B}}$ for $\lambda_{2,n} \neq 0$ (*32*). Refer to Safikhani and Shojaie for more details on the optimization problem (5) and how to select the tuning parameters in finite sample (*28*).

An interesting fact about the first step (fused lasso) is that this step might lead to over-estimation of number of break/change points in the model. Thus, in the second step, the redundant candidate change points detected in the first step are removed through a careful screening step. The details are omitted in the second step. Refer to Safikhani and Shojaie, Section 4, Equations 9–12, for more details on the screening step (*28*). Under basic regularity conditions and enough jumps in the parameters, this procedure is proved to consistently estimate the number of change points and their locations (*29*). In the third and final step, model parameter estimates in each segment in which there are no more change points are estimated using the least squares ideas combined with penalization in the high-dimensional regime (similar to the estimation procedure) (5).

Data Preparation

The data used for the model is the subway station ridership from January 2019 to July 2020, retrieved from the New York City Transit open data system. New York City Subway system is operated by the MTA. It includes 472 stations in operation and approximately 4,600 turnstiles. The subway station turnstile ridership, either entries or exits, are provided on a 4 h increment by the MTA. In this research, the authors picked some subway stations in different boroughs in NYC. The selection of subway station was made randomly while making sure that there are no inconsistencies in the data during the study period. The busiest subway stations functioning as transport hubs were of limited use for the purpose of this research because they include many subway lines that might not work all the time, as they do not follow the data continuity of the time series.

To perform statistical analysis as well as anomaly detection, data preparation is done on some subway stations in NYC. The authors look first at the data for South Ferry Subway station located in Lower Manhattan. It is located close to the ferry station from the borough of Manhattan to the borough of Staten Island, and it is also near the City's financial district, home to Wall Street and skyscrapers. The data was aggregated to a daily basis from subway turnstile entries. During data processing and cleaning, some errors were noticed concerning turnstile records.

In the preprocessing part of calculating the entry ridership, a data warehouse was created using the MTA turnstile data. In that process, three categories were created: (i) subunit channel position, which represents a specific address of the used device; (ii) date of the associated data

point; and (iii) time of the associated datapoint. There is a metric that shows the cumulative entry of registered value of an associated device. These cumulative entries were subtracted from the ones of the previous days, to give the daily "entries." The MTA records subway ridership at every turnstile, which basically report both entry and exit. The number of ridership reported, either entries or exits, are turnstile counter values. When the counter reaches its maximum limit, the counter is reset. A drastic shift in the absolute (cumulative) number reported in the data is noticed during the time period when the reset happens. This issue is accounted for in the data cleaning process. To do this, when the cumulative sequence value changed significantly, the rolling average of the "entries" subway values was calculated.

To understand the trends of subway ridership, the ridership data over time at the South Ferry Subway station is used for illustration. Figure 1 shows the daily turnstile entries and moving average of the mean over a 7-day rolling window at the South Ferry Subway station, from January 1, 2020, to June 19, 2020. This ridership includes all turnstile entries of the South Ferry Subway station. The rolling average method reduces noise in time series data, giving the ability to look at obvious trends. The subway ridership decreased gradually, then plunged suddenly in March 2020, and continued to perform with a very low ridership during April 2020 as shelter-in-place orders were mandated by state officials.

In the following, more descriptive analyses of this subway ridership are provided. Figure 2 shows the histogram, ACF, and PACF of the raw ridership at the South Ferry Subway station. The histogram is skewed to the left, as a portion of the dataset includes a time period during the COVID-19 pandemic which imposed a low ridership on the subway. From the ACF and PACF plots, it is noticeable that there is a seasonal pattern in the data. This seasonality is mostly because of the weekly mobility of commuters utilizing subway for work purposes on weekdays.

To better understand the patterns between different data points in time series, the time series' trend and seasonality are removed from the dataset. Then, the focus will just be on the time series of the residuals to decipher any temporal dependencies among them. First, a linear regression was applied to incorporate the imbalance between weekdays and weekends ridership. To do so, Saturday and Sunday (weekend) are labeled as 1, other days (weekdays) are considered as 0. In the linear regression model, the dummy categorial weekday/weekend indicator, being 0 as weekday and 1 as weekend, is defined as the independent variable. After applying the linear regression, the residual from the regression model is left. Second, to remove seasonality in the dataset, a seasonal autoregressive moving average model (SARIMA) was
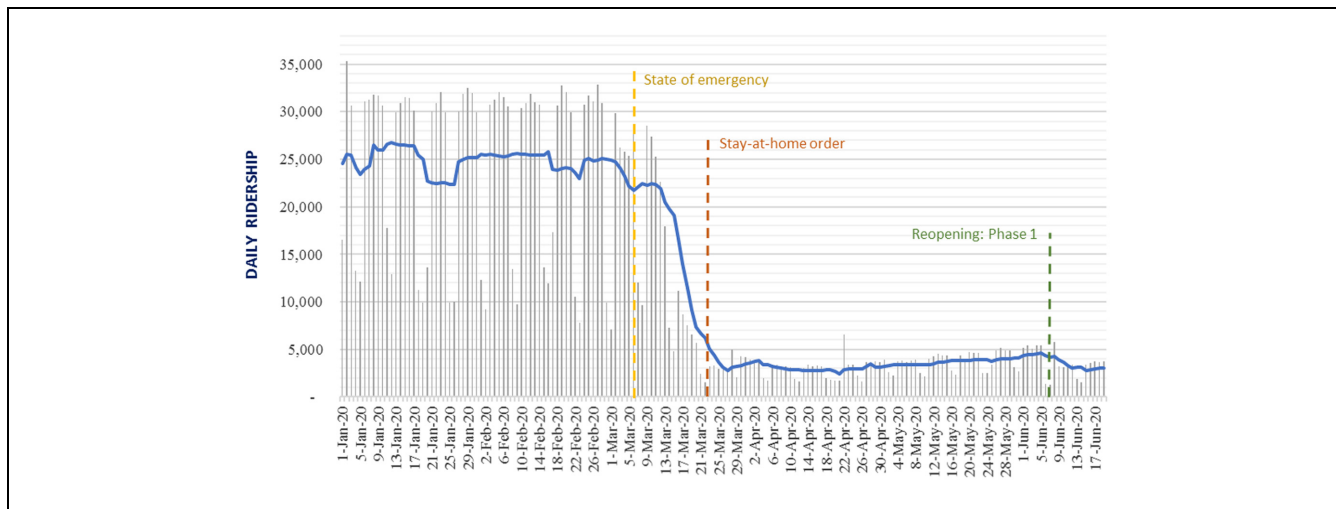
Figure 1. Daily entries and 7-day rolling average of the mean of South Ferry Subway station.
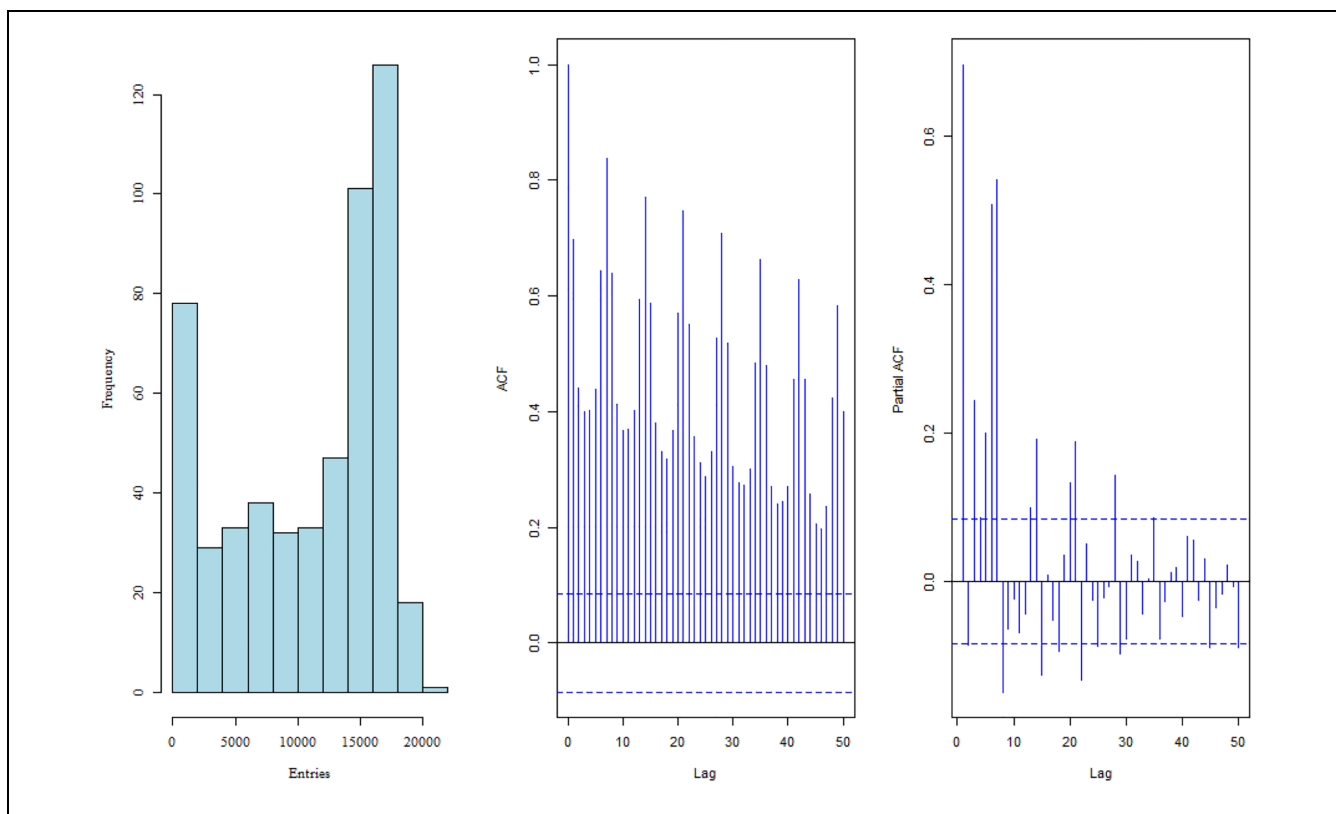
Figure 2. Histogram, autocorrelation function (ACF), and partial autocorrelation function (PACF) of the turnstile entries raw data at South Ferry Subway station, from Jan 2019 to April 17, 2020.

applied with periodic operator of 7. It treats the dataset as if there is a weekly seasonal trend. Finally, the residual from the SARIMA model, which is the absolute error of

the ridership, is the output to be analyzed for the rest of this paper. It is worth noting that no transformation—that is, log-transform—was applied to the raw dataset.
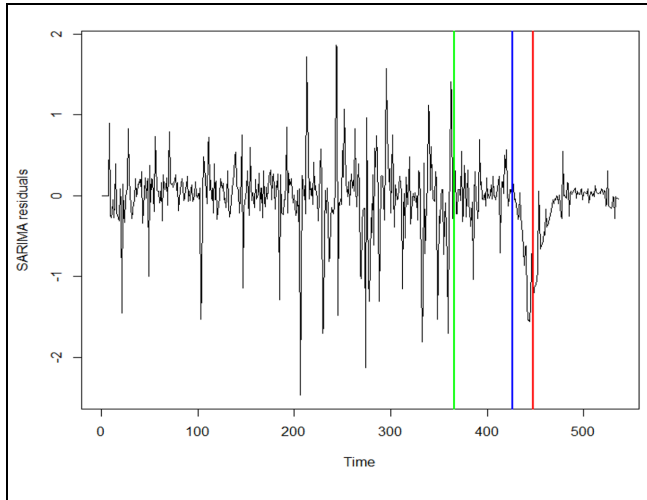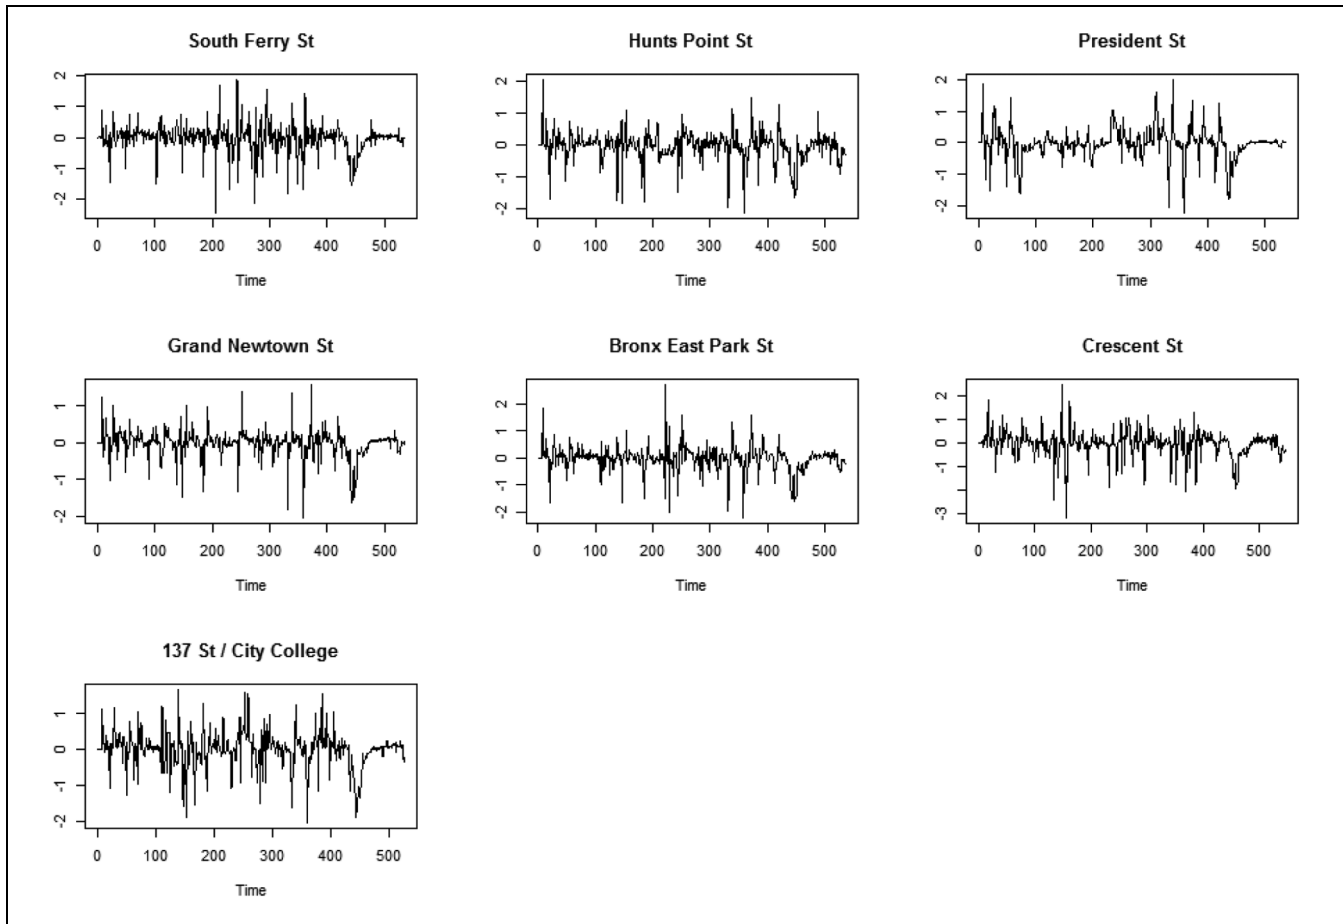
Figure 3. Residuals of the seasonal autoregressive moving average model (SARIMA) model (South Ferry Subway station).

Figure 3 shows the time series of residual after removing the trend and seasonality. The vertical green, blue, and

red lines in the graph correspond to January 1, March 22 (the day the state-at-home order was announced), and June 8 (the day for reopening of Phase 1), 2020, respectively. Also, Figure 4 displays time series trends of all seven subway stations after removing the trend and seasonality.

Figure 5 illustrates the ACFs of the residuals at every season from the beginning of 2019 to June 2020. It shows how the autocorrelation changes from one season to another. All four ACF plots in 2019 have the same shape, mostly significant at lag 1 and lag 7 with a different magnitude. However, the ACF for the winter season in 2020 shows a different pattern when it is compared with other seasons, specifically to the corresponding season in 2019. Specifically, there is a strong autocorrelation among data points in the winter season 2020 compared with the same time in 2019. Such extreme temporal correlations resemble the existence of change points in the data which appear as long-range dependence (*30*). A similar pattern has been observed for other subway stations considered here. This observation justifies the use of piecewise stationary time series modeling framework for the data sets focused on in this project.

Figure 4. Residuals of the seasonal autoregressive moving average model (SARIMA) model for all seven subway stations.
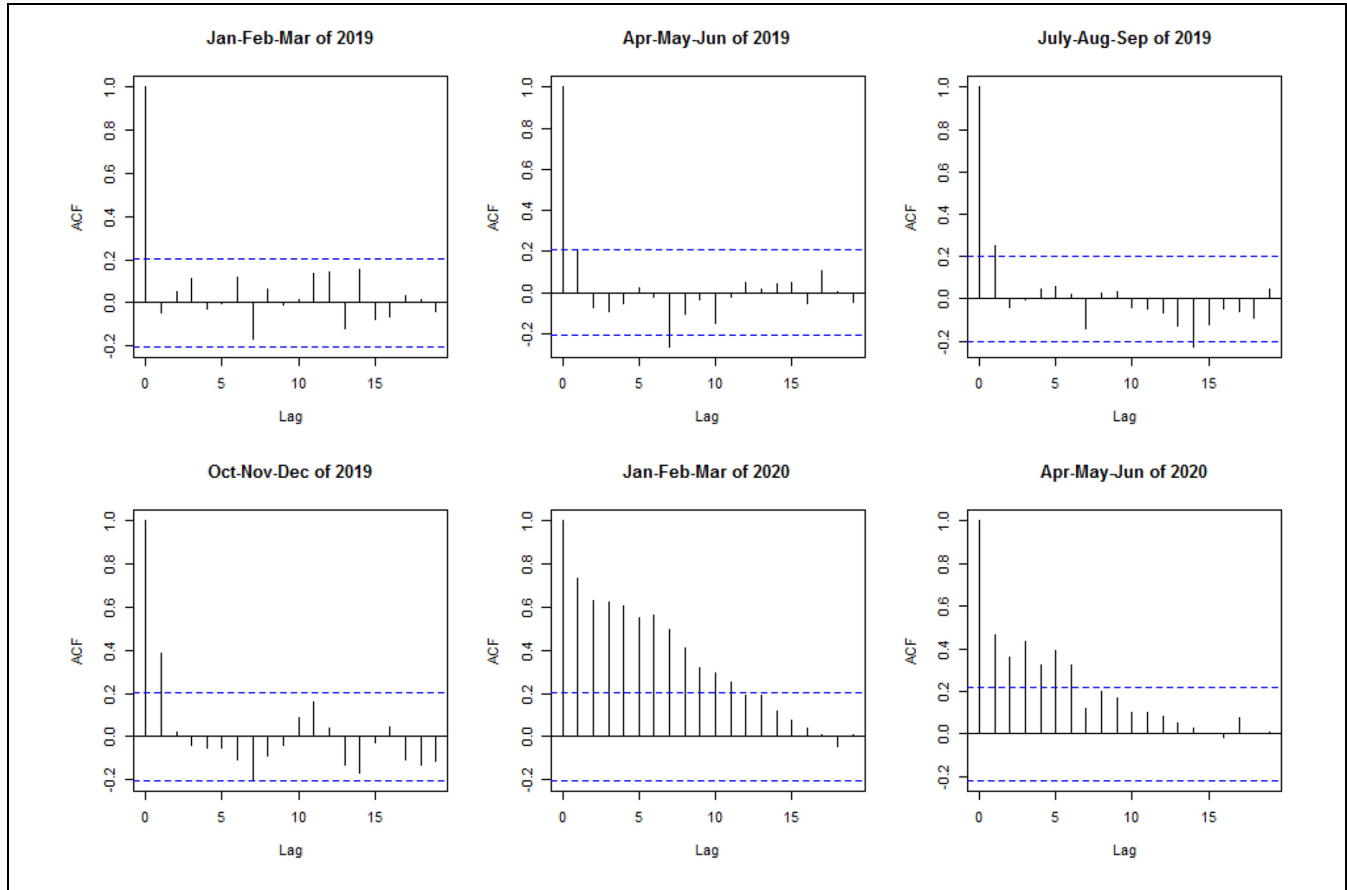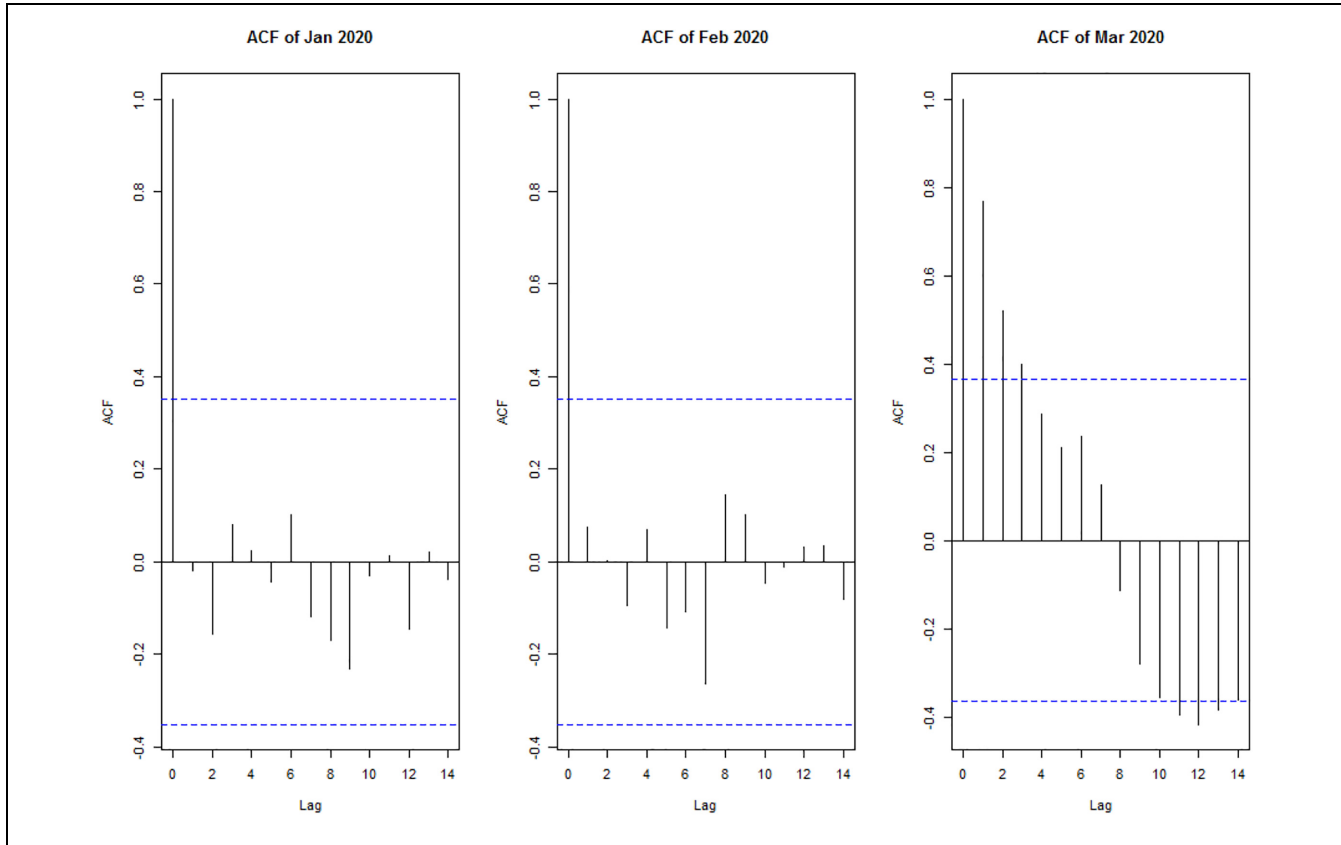
Figure 5. Autocorrelation functions (ACFs) of each season from beginning of 2019 until June 2020 (South Ferry Subway station).

Since the ACF of the winter season in 2020 is different, this season has been divided into three segments resembling the three months of winter. Figure 6 displays the ACFs of the first three months of 2020 (January, February, and March) which illustrates how the autocorrelation changes from one to another. It is possible to observe that the ACF in March has a different behavior. March 2020 corresponds to the time when the positive coronavirus cases emerged, and increased to the point when New York state officials announced the state of emergency, then the New York State's stay-at-home order took effect (3).

One of the main goals of this paper is to understand the subway's demand variability during unexpected events, in the face of the COVID-19 pandemic. As explained before, the approach is to consider the existence of anomalies at the points at which the subway demand suddenly changed—either surged or plunged. The existence of such anomalies in subway ridership would violate the assumption of stationarity in the time series models which is an underlying assumption in many non-stationary times series models in the literature.

In the first attempt at modeling, the existence of change points is ignored and ARIMA models are simply applied by applying the function "auto.arima" in the R programming language to the processed data. Table 1 illustrates the model summary of applying ARIMA models to the processed data. The table includes the model's name and its estimated parameters and the model's performance errors by Akaike information criterion (AIC), corrected AIC (AICC), and Bayesian information criterion (BIC), as well as the p-values from the white noise Li-McLeod's test (27). The first row in the table describes the summary of the fitted ARIMA model to the whole data. The fitted model for the whole data is ARIMA (0,1,2), which gives the best ARIMA model based on AICC. However, it turns out that this model is not a good fit for the data, since the residuals turn out not to be white noise.

The white noise test was performed using the Li-McLeod test method (see Brockwell and Davis for further information about the white noise test) (27). The p-value of the white noise test for the fitted mode was equal to 0.012, which indicates that there are still some

Figure 6. Autocorrelation functions (ACFs) of three months: January, February, and March, 2020 (South Ferry station).

Table 1. Time Series Models for the Whole Data of the South Ferry Subway Station Ridership

Data	Fitted model	Estimated parameters	Performance errors	Li-McLeod's test *p*-values
Whole data: January 2019 to June 2020	ARIMA (0,1,2)	$\theta_1 = -0.6368$ $\theta_2 = -0.2590$	AIC = 653.81 AICC = 653.85 BIC = 666.65	0.012 (less than 0.05)

Note: AIC = Akaike information criterion; AICC = corrected AIC; ARIMA = autoregressive integrated moving average; BIC = Bayesian information criterion.

dependencies existing among lagged data points, so further modeling and modification are required.

The same simple analysis was performed on the subway ridership for other randomly selected subway stations in NYC. Specifically, seven subway stations that have been chosen randomly across different boroughs are: South Ferry Subway station in lower Manhattan, Crescent Street in Brooklyn, Bronx Park East station in Bronx, 137 Street (City College) station in Upper Manhattan, Hunts Point Avenue station in Bronx, President Street station in Brooklyn, and Grand Avenue-Newtown station in Queens. The subway stations were selected so that each time series of ridership does not have any missing data. Table 2 shows the fitted ARIMA

models at all seven selected subway stations, with their coefficient values.

The existence of long-range dependence in the data together with failing to model the processed data sets as a stationary ARIMA model motivated the use of ARIMA modeling with change points. The first step to develop such a model is to find change points in each data set.

After running the change point detection algorithm developed in Safikhani and Shojaie, the results of each subway station data are reported in Table 3 (*28*). The change point for the 137 Street station (City College) in Upper Manhattan is March 13, 2020, while the change point for Bronx East Park station is March 12, 2020. Further, the change point for Crescent Street station in

Table 2. Time Series Model for all Seven Subway Stations

Subway station	Fitted model	Estimated parameters	Performance errors	Li-McLeod's test p-values
South Ferry station	ARIMA (0,1,2)	$\theta_1 = -0.6368$ $\theta_2 = -0.2590$	AIC = 653.81 AICC = 653.85 BIC = 666.65	<0.05
137 Street, City College	ARIMA (1,0,2)	$\phi_1 = 0.886$ $\theta_1 = -0.342$ $\theta_2 = -0.299$	AIC = 610.36 AICC = 610.44 BIC = 627.42	<0.05
Bronx East Park station	ARIMA (1,1,1)	$\phi_1 = 0.219$ $\theta_1 = -0.871$	AIC = 587.61 AICC = 587.65 BIC = 600.45	<0.05
Crescent Street, Brooklyn	ARIMA (1,0,1)	$\phi_1 = 0.435$ $\theta_1 = 0.148$	AIC = 754.63 AICC = 754.58 BIC = 767.54	<0.05
Hunts Point station	ARIMA (1, 0, 2)	$\phi_1 = 0.819$ $\theta_1 = -0.397$ $\theta_2 = -0.145$	AIC = 527.04 AICC = 527.12 BIC = 544.18	<0.05
President Station	ARIMA (1, 0, 0)	$\phi_1 = 0.682$	AIC = 390.7 AICC = 390.72 BIC = 399.27	<0.05
Grand Newton station	ARIMA (2, 1, 4)	$\phi_1 = -0.4081$ $\phi_2 = -0.7528$ $\theta_1 = -0.1235$ $\theta_2 = 0.3702$ $\theta_3 = -0.6168$ $\theta_4 = -0.2732$	AIC = 311.76 AICC = 311.97 BIC = 341.74	<0.05

Note: AIC = Akaike information criterion; AICC = corrected AIC; ARIMA = autoregressive integrated moving average; BIC = Bayesian information criterion.

Table 3. Change Point of Subway Station Ridership

Subway station	Change point (time)
137 Street, City College	March 13, 2020
Bronx East Park station	March 12, 2020
Crescent Street, Brooklyn	March 13, 2020
South Ferry station	March 13, 2020
Hunts Point station	June 28, 2019, and March 12, 2020
President station	March 13, 2020
Grand Newton station	March 13, 2020

Brooklyn is March 13, 2020, and the change point for South Ferry Subway station in Lower Manhattan is March 13, 2020. Two change points detected for Hunts Point station are June 28, 2019, and March 12, 2020. The change points for President Station and Grand Newton station are detected to be March 13, 2020.

Understandably, all change points detected for the subway station ridership considered are within almost a week from the state of emergency reported in the state of New York. Detecting another change point for the Hunts Point station around June 28, 2019, is related to maintenance and partial closure of this subway station which caused a significant drop in its ridership. Such a consistent performance of the detection algorithm on all

subway stations considered here is an indication that the piecewise stationary time series model focused on this project is a reasonable choice. Next, ARIMA models are applied to each stationary segment estimated by this method to further investigate the goodness of fit for the proposed model.

With change points values being found, instead of decomposing the data through an eyeballing procedure, the subway ridership data at each station has been separated into before and after the change point(s). Table 4 shows the fitted ARIMA model (again, applying the function "auto.arima" in the R programming language) for each subway station before and after the change point(s). It should be noted again that, before running the change points algorithm, the seasonality and trend of each dataset were already considered.

The results indicate that with the analyzed data, there are two models (before and after the change point) that could be very well developed to represent the variability in the data set. The two models (or three with multiple change points) have performed better compared with using only one model. To perform a goodness of fit, all residuals are tested using the same procedure as before, that is, the Li-McLeod test method (27). Most of the computed p-values are greater than 5% which indicates that the new residuals don not resemble any additional

Table 4. Time Series Model for all Four Subway Stations before and after Change Points

Subway station	Fitted model	Estimated parameters	Performance errors	Li-McLeod's test p-values
South Ferry station (before change point)	ARIMA (0,0,1)	$\theta_2 = 0.2677$	AIC = 563.07 AICC = 563.1 BIC = 571.23	>=0.05
South Ferry station (after change point)	ARIMA (3,1,3)	$\phi_1 = 1.3867$ $\phi_2 = -1.0280$ $\phi_3 = 0.0795$ $\theta_1 = -1.9053$ $\theta_2 = 1.5240$ $\theta_3 = -0.370$	AIC = −34.28 AICC = −33.03 BIC = −16.26	>=0.05
137 Street, City College (before change point)	ARIMA (1,0,1)	$\phi_1 = 0.310$ $\theta_1 = 0.492$	AIC = 486.4 AICC = 486.5 BIC = 498.7	<0.05
137 Street, City College (after change point)	ARIMA (5,0,1)	$\phi_1 = 0.0536$ $\phi_2 = -0.8238$ $\phi_3 = 0.4975$ $\phi_4 = -0.1208$ $\phi_5 = 0.2307$ $\theta_1 = 0.4355$	AIC = 531.2 AICC = 531.53 BIC = 563.84	>=0.05
Bronx East Park station (before change point)	ARIMA (0,0,2)	$\theta_1 = 0.1761$ $\theta_2 = 0.1947$	AIC = 491.35 AICC = 491.4 BIC = 503.58	>=0.05
Bronx East Park station (after change point)	ARIMA (2,1,1)	$\phi_1 = 0.7612$ $\phi_2 = -0.1742$ $\theta_1 = -0.8191$	AIC = −15.55 AICC = −15.12 BIC = −5.25	>=0.05
Crescent Street, Brooklyn (before change point)	ARIMA (0,0,2)	$\theta_1 = 0.5124$ $\theta_2 = 0.1316$	AIC = 659.27 AICC = 659.32 BIC = 671.58	>=0.05
Crescent Street, Brooklyn (after change point)	ARIMA (0,1,0)		AIC = 4.83 AICC = 4.87 BIC = 7.42	>=0.05
Hunts Point (before change point)	ARIMA (0,0,1)	$\theta_1 = 0.2199$	AIC = 178.2 AICC = 178.27 BIC = 184.57	>=0.05
Hunts Point (after change point, 2 models)	ARIMA (3,0,0) ARIMA (2,1,1)	$\phi_1 = 0.4740$ $\phi_2 = -0.0671$ $\phi_3 = 0.0522$ $\phi_1 = 0.3942$ $\phi_2 = 0.2132$ $\theta_1 = -0.9014$	AIC = 262.91 AICC = 263.07 BIC = 277.14 AIC = 64.67 AICC = 65.1 BIC = 74.96	>=0.05 >=0.05
President station (before change point)	ARIMA (1,0,0)	$\theta_1 = 0.6636$	AIC = 394.15 AICC = 394.18 BIC = 402.31	>=0.05
President station (after change point)	ARIMA (6,1,0)	$\phi_1 = 0.1121$ $\phi_2 = -0.6223$ $\phi_3 = -0.3153$ $\phi_4 = -0.1176$ $\phi_5 = -0.5350$ $\phi_6 = 0.3523$	AIC = −140.1 AICC = −138.48 BIC = −119.42	>=0.05
Grand Newton station (before change point)	ARIMA (1,0,0)	$\phi_1 = 0.3244$	AIC = 270.06 AICC = 270.09 BIC = 278.23	<0.05
Grand Newton station (after change point)	ARIMA (2,1,2)	$\phi_1 = 1.3181$ $\phi_2 = -0.9618$ $\theta_1 = -1.4192$ $\theta_2 = 0.8813$	AIC = −52.96 AICC = −52.3 BIC = −40.09	<0.05

Note: AIC = Akaike information criterion; AICC = corrected AIC; ARIMA = autoregressive integrated moving average; BIC = Bayesian information criterion.

temporal dependence, therefore the piecewise ARIMA model fitted to the data is a reasonable fit. It should be noted that, in ARIMA models, for higher lags (p or q more than 1), it is possible to have phi or theta parameters with magnitudes more than one, which is the case for models in the stations at South Ferry and at Grand Newton.

Another interesting finding is that the changes which occurred in the ridership data are not only on means (average of ridership), but the second-order statistics are changed as well. Reviewing the fitted models in Tables 2 and 4 shows that the order of ARIMA models (selected p, d, q) and the estimated auto-regressive and moving average parameters in different segments for each subway station are different. For example, the selected model for South Ferry data before the break point is a simple MA(1) while after the break point is an ARIMA(3,1,3). The former is a very simple model with minor correlations among data points with one time-lag apart from each other, while the latter is a more complicated model with six parameters, and the temporal dependence is much stronger than the first model. Therefore, the correlations among time-lagged ridership have changed during the pandemic. Such second-order changes are not easy to detect by eyeballing or just plotting the data as a time series. This shows the necessity of developing statistical modeling frameworks which would include anomalies in the data to analyze and understand better the behavior of subway ridership over time during such extreme times (e.g., pandemic).

Finally, looking at the results in the Grand Newton station, the white noise test p-values are relatively low which means the current model may not be a good fit for this data. A non-linear time series model be a better choice for this specific station.

Conclusion

In this paper, a linear and non-stationary ARIMA-based time series model is developed to understand the temporal pattern of station-based subway ridership during the COVID-19 pandemic. The main objective of the paper is to propose an interpretable modeling framework to characterize the temporal pattern of subway ridership during the COVID-19 pandemic.

Interpretability lies in the modeling framework utilized in this paper. It is a common practice to model data with temporal index using stationary time series models, and ARIMA models are among the most well-known stationary time series models used in different scientific fields including civil engineering and transportation. However, in the presence of structural breaks (shocks) in the temporal dynamical system under consideration, the

stationarity assumption may be violated; thus, it is necessary to search for alternative modeling frameworks. Nonparametric statistical models using B-splines and wavelet basis functions are models which may be able to handle the changes in the dynamical system. However, they are not easy to interpret. This is one of the drawbacks of nonparametric methods, while they are useful in practice. Piecewise stationary models, on the other hand, are interesting models which are easier to interpret, since a time point at which the dynamical system receives an external shock can be called a "break point," and its location can be estimated using developed algorithms. Such time points are essentially the discontinuity points in the piecewise modeling framework; this is why the modeling framework is called an interpretable one. For example, the developed algorithm detects a break point in the ridership of South Ferry Station on March 13, 2020. This date could be linked to an external shock in the dynamical system of ridership at this station, and is indeed close to the date of the stay-at-home order at the state level in NYC. Note that it is almost impossible to make such connections to real events using nonparametric models.

To that end, a simple yet powerful family of time series models are utilized (ARIMA models) while the inclusion of the sudden changes (break points) in the model accounts for discontinuity/anomalous behavior of subway ridership. A unique feature of the developed model is that it balances between model complexity and model interpretability. In other words, the statistical model designed for the time series data under investigation is deliberately selected to be simple and interpretable while it is a valid model (i.e., fits very well to the data). It is possible to note not only that the detected break points found by the algorithm are near the date of the "stay-at-home" order in NYC, which validates the proposed modeling framework, but also that the model shows the "covariance structure" of the time series have changed before and after the pandemic.

In summary, the main contributions of the paper include (a) developing a simple and interpretable statistical modeling framework to analyze the subway ridership in NYC during the COVD-19 pandemic; (b) utilizing novel machine learning algorithms from the statistical literature to detect the location of change/break points while the number of such break points are assumed to unknown; and (c) enhancing the scientific community's understanding of changes in the subway ridership during such pandemics by shedding some light on the second-order changes of subway ridership, a topic worth of further investigation.

It is worth noting that there are two different approaches to change point detection in the statistical

and engineering literature: (1) online detection and (2) off-line detection (*33*). In the latter, the whole time series data is given to the modeler and the objective is to find the set of break/change points, while in the former, a streaming data is observed (new data points are observed continuously) and the objective is to raise an alarm as soon as an anomalous pattern is observed in the data. The change point model developed in this paper is off-line, since the whole observed data was used and then novel algorithms were used to locate all break points. Developing online detection algorithms needs different statistical treatments (mainly, likelihood ratio tests), different modeling frameworks (distributional assumption on the time series), and different algorithms. Further, certain assumptions must be made to define identifiable break/change points, which may not hold for the subway ridership data in NYC during the COVID-19 pandemic. For these reasons, the authors focused on off-line change point detection and leave the online detection as a fruitful future research direction. Moreover, analyzing and comparing the effect of COVID-19 on transit ridership of other metropolitan cities would be in the continuation of the author's research to better justify the need for elaborated time series models. Also, how the transit ridership reacted during this pandemic with reopening strategies or how the subway ridership returned in the aftermath of the COVID-19 crisis could add other anomalies to this time series modeling. These questions are a few important research directions that the authors plan to pursue in the near future.

Author Contributions

The authors confirm contribution to the paper as follows: study conception and design: B. Moghimi, C. Kamga, A. Safikhani, S. Mudigonda, P. Vicuna; data collection: P. Vicuna, B. Moghimi; analysis and interpretation of results: A. Safikhani, B. Moghimi; draft manuscript preparation: A. Safikhani, B. Moghimi, C. Kamga., S. Mudigonda. All authors reviewed the results and approved the final version of the manuscript.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

ORCID iDs

Bahman Moghimi https://orcid.org/0000-0003-4129-1388
Sandeep Mudigonda https://orcid.org/0000-0003-1734-673X

References

1. Bonaccorsi, G., F. Pierri, M. Cinelli, A. Flori, A. Galeazzi, F. Porcelli, A. L. Schmidt, C. M. Valensise, A. Scala, W. Quattrociocchi, and F. Pammolli. Economic and Social Consequences of Human Mobility Restrictions Under COVID-19. *Proceedings of the National Academy of Sciences*, Vol. 117, No. 27, 2020, pp. 15530–15535.
2. https://www.nytimes.com/2020/03/22/nyregion/Coronavirus-new-York-epicenter.html.
3. https://www.nytimes.com/interactive/2020/us/coronavirus-stay-at-home-order.html.
4. Kamga, C., P. Vicuna, S. Mudigonda, and R. Tchamna. *Mobility Trends in New York City During COVID-19 Pandemic: Analyses of Transportation Modes Throughout June 2020*. University Transportation Research Center, New York, NY, 2020.
5. McNally, M. G. The Four-Step Model. In *Handbook of Transport Modelling* (Hensher, D. A., and K. J. Button, eds.), Pergamon Press, Oxford, Vol. 1, 2000, pp. 35–41.
6. Singhal, A., C. Kamga, and A. Yazici. Impact of Weather on Urban Transit Ridership. *Transportation Research Part A: Policy and Practice*, Vol. 69, 2014, pp. 379–391.
7. Liu, L., and R.-C. Chen. A Novel Passenger Flow Prediction Model Using Deep Learning Methods. *Transportation Research Part C: Emerging Technologies*, Vol. 84, 2017, pp. 74–91.
8. Liu, Y., Z. Liu, and R. Jia. DeepPF: A Deep Learning Based Architecture for Metro Passenger Flow Prediction. *Transportation Research Part C: Emerging Technologies*, Vol. 101, 2019, pp. 18–34.
9. Ding, C., D. Wang, X. Ma, and H. Li. Predicting Short-Term Subway Ridership and Prioritizing Its Influential Factors Using Gradient Boosting Decision Trees. *Sustainability*, Vol. 8, No. 11, 2016, p. 1100.
10. Liu, S., and E. Yao. Holiday Passenger Flow Forecasting Based on the Modified Least-Square Support Vector Machine for the Metro System. *Journal of Transportation Engineering, Part A: Systems*, Vol. 143, No. 2, 2017, p. 04016005.
11. Zhang, D., and X. C. Wang. Transit Ridership Estimation With Network Kriging: A Case Study of Second Avenue Subway, NYC. *Journal of Transport Geography*, Vol. 41, 2014, pp. 107–115.
12. Barua, S., A. Das, and K. C. Roy. Estimation of Traffic Arrival Pattern at Signalized Intersection Using ARIMA Model. *International Journal of Computer Applications*, Vol. 128, No. 1, 2015, pp. 1–6.
13. Willams, B. M., and L. A. Hoel. Modeling and Forecasting Vehicular Traffic Flow as a Seasonal ARIMA Process: Theoretical Basis and Empirical Results. *Journal of Transportation Engineering*, Vol. 129, No. 6, 2003, pp. 664–672.
14. Moghimi, B., A. Safikhani, C. Kamga, and W. Hao. Cycle-Length Prediction in Actuated Traffic-Signal Control Using ARIMA Model. *Journal of Computing in Civil Engineering*, Vol. 32, No. 2, 2018, p. 04017083.
15. Duan, P., G. Mao, C. Zhang, and S. Wang. STARIMA-based Traffic Prediction With Time-varying Lags. *Proc., 19th International Conference on Intelligent Transportation*

System (ITSC), Rio de Janeiro, Brazil, IEEE, Piscataway, NJ, 2016.

16. Ding, C., J. Duan, Y. Zhang, X. Wu, and G. Yu. Using an ARIMA-GARCH Modeling Approach to Improve Subway Short-Term Ridership Forecasting Accounting for Dynamic Volatility. *IEEE Transactions on Intelligent Transportation Systems*, Vol. 19, No. 4, 2017, pp. 1054–1064.

17. Paul, J., and M. J. Smart. The Hangover: Assessing Impact of Major Service Interruptions on Urban Rail Transit Ridership. *Transportation Research Record: Journal of the Transportation Research Board*, 2017. 2648: pp. 79–85.

18. Zhu, S., and D. M. Levinson. Disruptions to Transportation Networks: A Review. In *Network Reliability in Practice* (Levinson, D., H. Liu, and M. Bell, eds.), Springer, New York, NY, 2012, pp. 5–20.

19. van Exel, N., A. Job, and P. Rietveld. Public Transport Strikes and Traveller Behaviour. *Transport Policy*, Vol. 8, No. 4, 2001, pp. 237–246.

20. Tang, S., and H. Gao. Traffic-Incident Detection-Algorithm Based on Nonparametric Regression. *IEEE Transactions on Intelligent Transportation Systems*, Vol. 6, No. 1, 2005, pp. 38–42.

21. Tsiligkaridis, A., and I. C. Paschalidis. Anomaly Detection in Transportation Networks Using Machine Learning Techniques. *Proc., MIT Undergraduate Research Technology Conference (URTC)*, Cambridge, MA, IEEE, Piscataway, NJ, 2017, pp. 1–4.

22. Margreiter, M. Automatic Incident Detection Based on Bluetooth Detection in Northern Bavaria. *Transportation Research Procedia*, Vol. 15, 2016, pp. 525–536.

23. Riveiro, M., M. Lebram, and M. Elmer. Anomaly Detection for Road Traffic: A Visual Analytics Framework. *IEEE Transactions on Intelligent Transportation Systems*, Vol. 18, No. 8, 2017, pp. 2260–2270.

24. Zhang, Y., H. Shi, F. Zhou, Y. Hu, and B. Yin. Visual Analysis Method for Abnormal Passenger Flow on Urban Metro Network. *Journal of Visualization*, Vol. 23, No. 6, 2020, pp. 1–18.

25. Moghimi, B., A. Safikhani, C. Kamga, W. Hao, and J. Ma. Short-term Prediction of Signal Cycle on an Arterial with Actuated-Uncoordinated Control Using Sparse Time Series Models. *IEEE Transactions on Intelligent Transportation Systems*, Vol. 20, No. 8, 2018, pp. 2976–2985.

26. Safikhani, A., C. Kamga, S. Mudigonda, S. S. Faghih, and B. Moghimi. Spatio-temporal Modeling of Yellow Taxi Demands in New York City Using Generalized STAR Models. *International Journal of Forecasting*, Vol. 36, No. 3, 2020, pp. 1138–1148.

27. Brockwell, P. J., and R. A. Davis. *Introduction to Time Series and Forecasting*. Springer, New York, NY, 2006.

28. Safikhani, A., and A. Shojaie. Joint Structural Break Detection and Parameter Estimation in High-Dimensional Non-Stationary Var Models. *Journal of the American Statistical Association*, Vol. 117, 2020, pp. 1–26.

29. Rinaldo, A. Properties and Refinements of the Fused Lasso. *The Annals of Statistics*, Vol. 37, No. 5B, 2009, pp. 2922–2952.

30. Aue, A., and L. Horváth. Structural Breaks in Time Series. *Journal of Time Series Analysis*, Vol. 34, No. 1, 2013, pp. 1–16.

31. Tibshirani, R., M. Saunders, S. Rosset, J. Zhu, and K. Knight. Sparsity and Smoothness via the Fused Lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, Vol. 67, No. 1, 2005, pp. 91–108.

32. Friedman, J., T. Hastie, H. Höfling, and R. Tibshirani. Pathwise Coordinate Optimization. *The Annals of Applied Statistics*, Vol. 1, No. 2, 2007, pp. 302–332.

33. Csörgö, M., and L. Horváth. *Limit Theorems in Change-Point Analysis*, Vol. 18. John Wiley & Sons Inc., New York, NY, 1997.