

Original Paper

Forecasting COVID-19 Hospital Census: A Multivariate Time-Series Model Based on Local Infection Incidence

Hieu M Nguyen^{*}, MSc; Philip J Turk^{*}, MSc, PhD; Andrew D McWilliams, MPH, MD

Center for Outcomes Research and Evaluation, Atrium Health, Charlotte, NC, United States

^{*}these authors contributed equally

Corresponding Author:

Hieu M Nguyen, MSc

Center for Outcomes Research and Evaluation

Atrium Health

1300 Scott Ave

Charlotte, NC, 28204

United States

Phone: 1 9706914892

Email: hieu.nguyen@atriumhealth.org

Abstract

Background: COVID-19 has been one of the most serious global health crises in world history. During the pandemic, health care systems require accurate forecasts for key resources to guide preparation for patient surges. Forecasting the COVID-19 hospital census is among the most important planning decisions to ensure adequate staffing, number of beds, intensive care units, and vital equipment.

Objective: The goal of this study was to explore the potential utility of local COVID-19 infection incidence data in developing a forecasting model for the COVID-19 hospital census.

Methods: The study data comprised aggregated daily COVID-19 hospital census data across 11 Atrium Health hospitals plus a virtual hospital in the greater Charlotte metropolitan area of North Carolina, as well as the total daily infection incidence across the same region during the May 15 to December 5, 2020, period. Cross-correlations between hospital census and local infection incidence lagging up to 21 days were computed. A multivariate time-series framework, called the vector error correction model (VECM), was used to simultaneously incorporate both time series and account for their possible long-run relationship. Hypothesis tests and model diagnostics were performed to test for the long-run relationship and examine model goodness of fit. The 7-days-ahead forecast performance was measured by mean absolute percentage error (MAPE), with time-series cross-validation. The forecast performance was also compared with an autoregressive integrated moving average (ARIMA) model in the same cross-validation time frame. Based on different scenarios of the pandemic, the fitted model was leveraged to produce 60-days-ahead forecasts.

Results: The cross-correlations were uniformly high, falling between 0.7 and 0.8. There was sufficient evidence that the two time series have a stable long-run relationship at the .01 significance level. The model had very good fit to the data. The out-of-sample MAPE had a median of 5.9% and a 95th percentile of 13.4%. In comparison, the MAPE of the ARIMA had a median of 6.6% and a 95th percentile of 14.3%. Scenario-based 60-days-ahead forecasts exhibited concave trajectories with peaks lagging 2 to 3 weeks later than the peak infection incidence. In the worst-case scenario, the COVID-19 hospital census can reach a peak over 3 times greater than the peak observed during the second wave.

Conclusions: When used in the VECM framework, the local COVID-19 infection incidence can be an effective leading indicator to predict the COVID-19 hospital census. The VECM model had a very good 7-days-ahead forecast performance and outperformed the traditional ARIMA model. Leveraging the relationship between the two time series, the model can produce realistic 60-days-ahead scenario-based projections, which can inform health care systems about the peak timing and volume of the hospital census for long-term planning purposes.

(*JMIR Public Health Surveill* 2021;7(8):e28195) doi: [10.2196/28195](https://doi.org/10.2196/28195)

KEYWORDS

COVID-19; forecasting; time-series model; vector error correction model; hospital census; hospital resource utilization; infection incidence

Introduction

SARS-CoV-2 is a novel member of the coronavirus family, and infections in humans can result in the disease COVID-19. The virus is transmitted primarily through droplets from coughing and sneezing and is highly infectious. Its basic reproduction rate is estimated to be in the low to mid 2s based on different models [1], compared to 2 for severe acute respiratory syndrome (SARS) and 1.3 for the 2009 swine flu [2]. Moderate to severe disease typically manifests with acute hypoxemia, and can progress to acute respiratory distress syndrome, multiorgan dysfunction, and death. Furthermore, an estimated 25%-30% of patients admitted to hospitals require intensive care admission [2]. In December 2019, the first cases were recorded in Wuhan, China, with subsequent spread across the world. In early 2020, the World Health Organization declared COVID-19 to be a global health emergency [3]. At the end of December 2020, SARS-CoV-2 had resulted in over 82 million documented cases and nearly 2 million deaths [4].

Our work is motivated by the need of hospital leaders to have timely and accurate forecasts to guide planning for surges in hospital demands due to the pandemic. Adequate preparation can help prevent or mitigate strains on hospital resources that result when hospitals exceed their historical capacity. On the contrary, being caught off-guard under a pandemic can devastate the population and health care systems. For example, previous models in India suggested falsely that it had reached herd immunity, encouraging complacency and insufficient preparation; however, on May 4, 2021, there was still a reported rolling average of 378,000 cases a day, which overwhelmed hospitals and health workers and resulted in a national health crisis [5]. Thus, to a health care system, an essential tool is a model that provides short- and long-range forecasting of the number of COVID-19-positive patients who will be admitted. This COVID-19 hospital census plays a central role in planning decisions that frequently require considerable lead time, such as increasing staff, creating physical beds and rooms, and procuring vital equipment (eg, ventilators and personal protective equipment).

Prior research has demonstrated the utility of forecasting hospital demands (eg, hospital admissions, intensive care unit census, and hospital overall census) using univariate time-series models such as the autoregressive integrated moving average (ARIMA), the seasonal autoregressive integrated moving average (SARIMA), and exponential smoothing [6-8]. Another approach is to use ensemble-based modeling. For example, a hybrid of a SARIMA model and a nonlinear autoregression artificial neural network model has been used to forecast hospital admissions [9]. In another example, two separate models, a time-series model for hospital admission and a patient-level logistic regression model for hospital discharge, were combined to predict the hospital census [10]. While these examples demonstrate the powerful potential of univariate time-series and ensemble modeling, neither incorporate factors inherent to

the behavior of the pandemic, which may serve as important leading indicators of hospital census, especially at times when infection rates become increasingly dynamic (eg, on the approach or descent of peak infection prevalence). To incorporate pandemic indicators into modeling requires recognition that such indicators are typically nonstationary. Consequently, while a stationary multivariate time-series model, called vector autoregression (VAR), has been successfully employed to forecast emergency department patient census by including other hospital resource indicators [11], it cannot be used in this situation. Rather, our problem will require nonstationary multivariate time-series models like the vector error correction model (VECM).

Recently, VECM has been used to forecast the demand for intensive care units during the COVID-19 pandemic by including hospital admission as a leading indicator [12]. Although hospital admission is a natural choice as a leading indicator, it has a short period of lead time (ie, hours to days) and thus, limited predictive power. A more powerful indicator for planning purposes would lead by days to weeks. We have previously used VECM to forecast COVID-19 hospital census using leading indicators from Google relative search volumes for COVID-19 testing-related terms combined with the number of people flagged as having possible COVID-19 when using an internet-based virtual health screening bot [13]. However, these COVID-19 indicators, which are based on symptoms, have limitations. For example, the symptoms of COVID-19 cannot be easily separated from other common conditions, such as the seasonal flu, and search patterns may change due to other external factors over time.

During the COVID-19 pandemic, many papers have been devoted to developing predictive models for the volume of new cases (ie, infection incidence) using various methods from time-series analyses [14-16] to advanced machine learning [17,18]. However, virtually no effort was focused on developing statistical models linking infection incidence to hospitalization. Because hospital admission typically follows the symptoms or exposure that may provoke a person to be tested by roughly 1 week, we hypothesize that at a local population level, infection incidence rates may have a stable relationship with and serve as a reliable leading indicator for the COVID-19 hospital census. In this paper, our main objective is to explore whether the local COVID-19 infection incidence and the COVID-19 hospital census can be successfully incorporated within a VECM to delivery satisfactory 7-days-ahead forecast performance and examine the application of this model to scenario-based long-term forecasting. From our experience, since there can be systematic changes due to the day of the week in a hospital time series, we will need to account for weekly seasonal effects and examine implications on short-term resource planning.

Methods

Time-Series Data


Atrium Health is a large, integrated health care system operating in North Carolina, South Carolina, and Georgia. In this paper, the COVID-19 hospital census (census) refers to the daily aggregate number of beds occupied by patients with COVID-19 at midnight across the subset of 11 Atrium Health hospitals in the greater Charlotte metropolitan area of North Carolina, plus a virtual hospital (Atrium Health Hospital at Home). The virtual hospital uses telemedicine to treat patients who require only a minimal level of care. The local COVID-19 infection incidence (incidence) is the aggregate daily count of new COVID-19-positive cases from 11 local counties belonging to the Cities Readiness Initiative (CRI) region, as designated by the North Carolina Department of Health and Human Services. The CRI region roughly approximates the market catchment area of these hospitals.

Using STL (seasonal and trend decomposition using Loess) time-series decomposition [19], we observed that the two time series had multiplicative weekly seasonality. We transformed both time series to achieve additive seasonality and linearize their relationship. The usual log transformation was applied to incidence. For operational purposes, the health system had previously decided to place an upper bound of 1000 patients with COVID-19 on the hospital time-series range, so we applied the following constrained log transformation so that the back-transformed census forecasts would satisfy the constraint:

$$f(x) = \log\left(\frac{x}{1000 - x}\right)$$

The forecast model described in the following sections was developed for these transformed time series. Figure 1 shows a plot of transformed census and incidence on a standardized scale for the period from May 15 to December 5, 2020. To affirm the association between the two transformed time series, we computed the Pearson cross-correlations between census and values of incidence at lags 0, -1, ..., -21.

Figure 1. Scaled time series for COVID-19 hospital census and local COVID-19 infection incidence in the Cities Readiness Initiative region for the period from May 15 to December 5, 2020. Transformed census (blue) and incidence (red) are linearly standardized to the 0-100 scale.



VECM

A VECM is a vector autoregressive model used for nonstationary multivariate time series and accounts for stable long-run relationships, that is, cointegration, between the time series. A $k \times 1$ time-series vector y_t is said to be cointegrated if there is at least one nonzero $k \times 1$ vector β_i , such that the linear combination $\beta_i^T y_t$ is trend-stationary. If r such linearly

independent vectors β_i ($i=1, \dots, r$) exist, we say y_t is cointegrated with cointegration rank r [20].

Following Pfaff [20], we first describe the VAR representation of order p of the VECM:

$$y_t = \Pi_1 y_{t-1} + \dots + \Pi_p y_{t-p} + \mu + \Phi D_t + \epsilon_t$$

for time $t=1, \dots, T$, where Π_i (for $i=1, \dots, p$) are $k \times k$ coefficient matrices of the lagged series at lag i , μ is a $k \times 1$ vector of constants, D_t is a 6×1 vector of weekly seasonal indicators, Φ is a $k \times 6$ coefficient matrix for seasonal indicators, and ϵ_t is a $k \times 1$ vector of random errors.

The VECM specification can be formulated as an algebraic rearrangement of the VAR representation as:

$$\Delta y_t = \Pi y_{t-1} + \Gamma_1 \Delta y_{t-1} + \dots + \Gamma_{p-1} \Delta y_{t-p+1} + \mu + \Phi D_t + \epsilon_t$$

where Δy_t is a $k \times 1$ vector of the differenced series $y_t - y_{t-1}$, $\Pi = -(I - \Pi_1 - \dots - \Pi_p)$, and $\Gamma_i = -(\Pi_{i+1} + \dots + \Pi_p)$ (for $i=1, \dots, p-1$).

The model has the following assumptions:

- Assumption 1: The components of y_t are at most $I(1)$, that is, an integrated of order 1
- Assumption 2: $0 \leq r = \text{rank}(\Pi) \leq k$
- Assumption 3: ϵ_t are identically and independently distributed $N(0, \Sigma)$ random vectors with covariance matrix Σ .

We now discuss the implications of the assumptions. For assumption 2, if $r=k$, then it can be shown that the VECM becomes a standard VAR model. If $r=0$, then Π is the zero matrix and there is no cointegration relationship between the series. The VECM then becomes a VAR model for differenced time series. If $0 < r < k$, then Π can be factored into $\Pi = \alpha \beta^T$, where α and β are both $k \times r$ matrices. From assumption 1, the differenced series Δy_t , and its lags $\Delta y_{t-1}, \dots, \Delta y_{t-p+1}$ are stationary. It follows that $\Pi y_{t-1} = \alpha \beta^T y_{t-1}$, as well as $\beta^T y_{t-1}$, also called the error correction term, is (trend-)stationary, depending on the specification of the deterministic components. The r linearly independent columns of β are the cointegrating vectors, and the rank r is equal to the cointegration rank of the system of time series.

Estimation and Inference

The VECM was specified and fitted with the steps below.

First, to choose the order p of the VAR representation, we fitted a VAR model to the data and made the decision based on the Akaike information criterion (AIC) [21].

Second, we determined the number of cointegration relationships ($r=0$ or $r=1$) using the Johansen trace test [22].

Third, we needed to decide where to place the constant μ in the model. One option was to leave μ as shown previously to account for linear trend in the data. Another option was to restrict $\mu = \alpha \rho$. The constant would be absorbed into the cointegration relationship as an intercept, and the data would not exhibit linear trend.

We made our decision about whether to restrict μ based on a likelihood ratio test for linear trend, as described elsewhere [23,24].

Fourth, we used maximum likelihood estimation to fit the model, reported parameter estimates, the corresponding T tests, and the

omnibus F tests with a significance level of .05, following Johansen [23].

Finally, we computed the 7-days-ahead forecasts and the 80% forecast intervals. Once the forecasts of the transformed census were made with the VECM, they were back-transformed to the original scale of census. We created 80% forecast intervals for the transformed census using a bootstrap procedure [25]. Then, the lower and upper bound of the forecast intervals were also back-transformed.

The model was fitted to the data between May 15 and December 5, 2020. All the data analysis was done using R statistical software, version 4.0.3 (R Core Team). The implementation of the VECM was done with the *tsDyn*, *vars*, and *urca* R packages. Since there were no packages to make bootstrapped forecast intervals for the VECM, we coded our own implementation. The data and code used in the data analysis are publicly available on GitHub [26].

Model Diagnostics

We examined the omnibus F tests to look for signs of lack of fit and also performed the multivariate Portmanteau test for the existence of serial correlation in the errors. Autocorrelation function and cross-correlation function plots were also generated for visual inspection. We performed the univariate and multivariate Jarque-Bera normality test on the errors [27] and also checked whether the cointegration relationship was stable, that is, stationary, using the Augmented Dickey-Fuller (ADF) test [28] and the Kwiatkowski-Phillips-Schmidt-Shin (KPSS) test [29]. Finally, we checked the stability of the estimated VAR representation. To do so, we looked at the companion matrix of the VAR representation and checked whether the maximum eigenvalue modulus was strictly smaller than 1, which, if true, would imply the stability of the VAR representation [30]. We also generated a trace plot of the maximum eigenvalue modulus, where the model was repeatedly fitted on a daily rolling basis, to check for the consistency of this value over time.

Forecast Performance

We used mean absolute percentage error (MAPE) to evaluate the 7-days-ahead forecasts of census:

$$MAPE = \frac{100}{7} \sum_{i=1}^7 \left| \frac{F_i - A_i}{A_i} \right|$$

where F_i is the forecast value and A_i is the actual value.

In order to approximate the sampling distribution of MAPE, we performed time-series cross-validation. From June 16 to November 28, 2020, for each day, we iteratively fitted the model, made 7-days-ahead forecasts, and computed the MAPE. Eventually, we obtained 166 values of MAPE, plotted the distribution, and computed the median as well as the 95th percentile. We will consider a median MAPE below 10% to be satisfactory, based on the practical effect of a peak surge on bed capacity at our health care system.

Scenario-Based Long-Term Forecasting

Leading up to and at the peak of infection prevalence, there can be high anxiety and uncertainty about how much more incidence

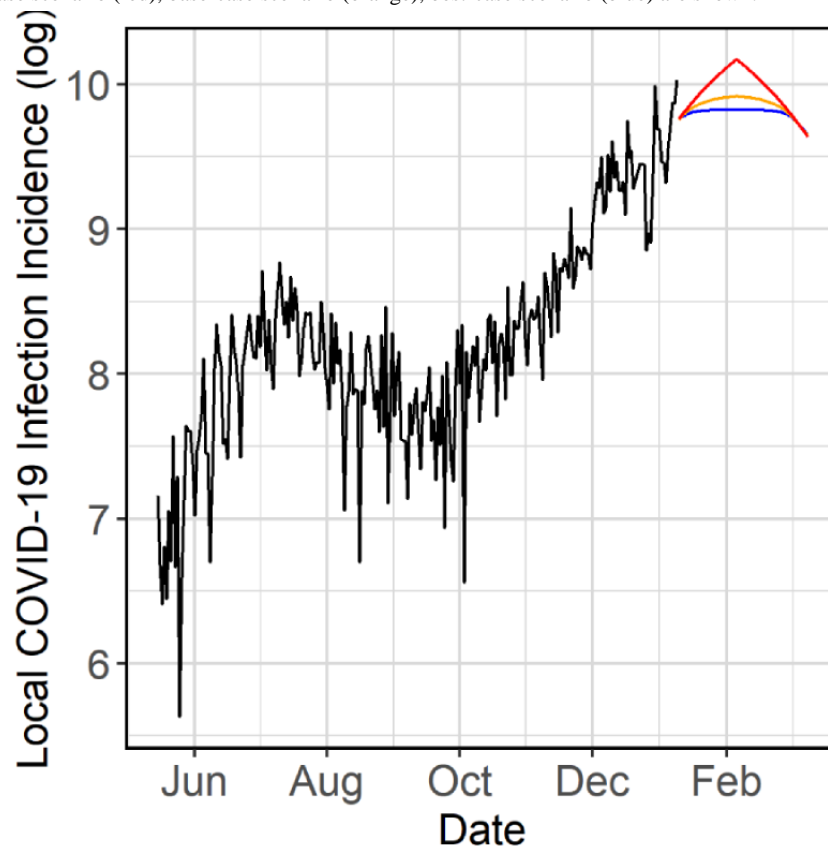
and, in particular, census may increase. Furthermore, traditional univariate time-series models may give linear forecasts for census that do not accurately represent pandemic behavior. However, cointegration allows for census forecasts that leverage subtle, but critical, changes in incidence (eg, concavity). This suggests, if not necessitates, the forecasting of census under different pandemic scenarios. For resource planning, hospital leaders will want to understand the implications associated with a worst-case scenario.

For our health care system, besides routine 7-days-ahead census forecasts, we also deployed our model for 60-days-ahead census forecasts, considering 3 different scenarios of what could happen with incidence (ie, best case, base case, and worst case). On January 9, 2021, we expected the winter surge to reach peak infection prevalence around February 5, 2021, based on an extension of an epidemiological model called the susceptible-infected-removed model [31]. While peak infection

incidence typically leads peak infection prevalence, in the absence of definitively knowing either peak date, we took a conservative approach and linearly extrapolated incidence with a positive trend up to the expected pandemic peak. The severity of a scenario was controlled by a trend-dampening parameter [32]. After the peak, the descent path was initially symmetric to its ascent and then eventually became linear (Figure 2).

Using our model refitted on January 9, 2021, with an increased capacity of 1250 patients, we generated forecasts iteratively forward for 60 days using the past census forecasts together with projected incidence under each scenario. To account for uncertainty in future census and incidence, we also simulated 1000 conditional sample paths of the two time series under each scenario using the bootstrap procedure mentioned earlier and computed the 10th and 90th percentile at each horizon to obtain the 80% forecast intervals.

Figure 2. The 60-day projected local COVID-19 infection incidence in the Cities Readiness Initiative region on the log scale, as of January 9, 2021. Past values (black), worst-case scenario (red), base-case scenario (orange), best-case scenario (blue) are shown.



Ethical Review

Our research protocol was submitted to the Atrium Health Institutional Review Board (IRB) prior to execution, and the study was deemed exempt from IRB oversight. In compliance with HIPAA (Health Insurance Portability and Accountability Act) regulations, individual patient information was not disclosed, and all data have been deidentified and reported as aggregates. The procedures set out in this protocol, pertaining to the conduct, evaluation, and documentation of this study, were designed to ensure that the investigators abide by Good

Clinical Practice guidelines and under the guiding principles detailed in the Declaration of Helsinki.

Results

Estimation and Inference

Our model was specified as a VECM with 7 lags in its VAR representation ($p=7$), 1 cointegration relationship ($r=1$), and a restricted constant parameter μ so that the series would not have linear trend. The AIC scores of VAR models with a varying number of lags from 2 to 14 were inconclusive. However, we

found that 7 lags were sufficient to account for all the correlation in the data, as evidenced by the autocorrelation function and cross-correlation function plots of the residuals (Figure 3). The Johansen trace test indicated that there was 1 cointegration relationship (significant at 1%, based on tabulated critical values). Finally, the likelihood ratio test for linear trend indicated that there was no linear trend in the data ($P=.32$). Furthermore, the restricted model had a lower AIC score than the unrestricted model (the AIC scores were -1519 and -1516 , respectively).

The output from the maximum likelihood estimation showed that the cointegration relationship, that is, the error correction term, had a significant negative effect on census change ($P<.001$); no significant effect was observed for incidence change ($P=.26$) (Table 1). The long-run cointegration relationship was estimated as:

$$ect_{t-1} = census_{t-1} - 0.8013incidence_{t-1} + 7.8266$$

where ect_{t-1} was the (lagged) error correction term. Table 1 also shows that past changes in census and incidence also had meaningful effects on current census change. Past census changes had significant effect at lag 2 ($P=.002$). Past incidence changes had significant effects at lag 1 ($P=.005$), lag 2 ($P=.04$), lag 4 ($P=.02$), lag 5 ($P=.03$), and lag 6 ($P=.02$).

From Table 2, there were some significant seasonal effects, that is, differences in both census and incidence changes among days of the week. Compared to Thursday, census change was higher on Monday and incidence change was lower on Sunday, with significant differences ($P=.01$ and $P=.002$, respectively).

Figure 3. Autocorrelation functions and cross-correlation functions of the residuals: (A) census residuals, (B) lagged census residuals and incidence residuals, (C) census residuals and lagged incidence residuals, and (D) incidence residuals.

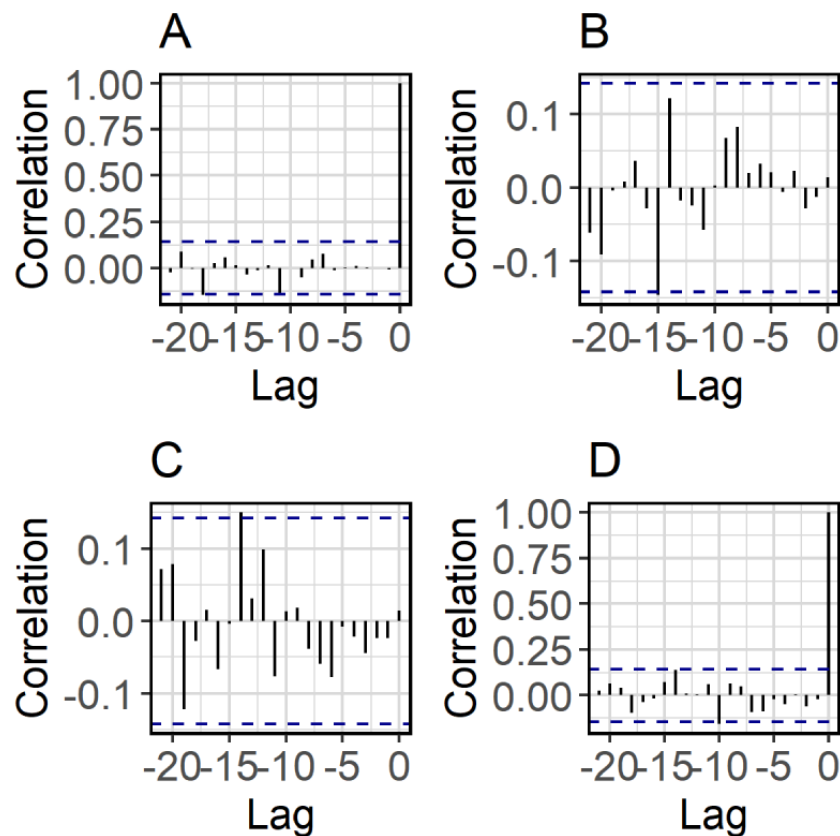


Table 1. Parameter estimates and *T* tests for nonseasonal effects.

Predictor	$\Delta Census_t$			$\Delta Incidence_t$		
	Estimate	<i>T</i> statistics	<i>P</i> value	Estimate	<i>T</i> statistics	<i>P</i> value
ect_{t-1}	-0.1265	-5.6993	<.001	-0.1216	-1.1323	.26
$\Delta Census_{t-1}$	-0.0489	-0.7143	.48	0.5487	1.6555	.10
$\Delta Incidence_{t-1}$	-0.0665	-2.8222	.005	-0.9808	-8.6067	<.001
$\Delta Census_{t-2}$	-0.2220	-3.2277	.002	-0.0614	-0.1844	.85
$\Delta Incidence_{t-2}$	-0.0532	-2.0881	.04	-0.6955	-5.6431	<.001
$\Delta Census_{t-3}$	-0.0700	-0.9949	.32	0.0643	0.1890	.85
$\Delta Incidence_{t-3}$	-0.0472	-1.9094	.06	-0.6428	-5.3755	<.001
$\Delta Census_{t-4}$	-0.0785	-1.1224	.26	0.9769	2.8871	.004
$\Delta Incidence_{t-4}$	-0.0567	-2.4165	.02	-0.5564	-4.8999	<.001
$\Delta Census_{t-5}$	-0.0499	-0.7140	.48	-0.0792	-0.2341	.82
$\Delta Incidence_{t-5}$	-0.0465	-2.1907	.03	-0.4589	-4.4634	<.001
$\Delta Census_{t-6}$	0.0077	0.1107	.91	0.4533	1.3404	.18
$\Delta Incidence_{t-6}$	-0.0373	-2.4015	.02	-0.2384	-3.1739	.002

Table 2. Parameter estimates and *T* tests for day-of-the-week effects, in comparison with Thursday being the reference.

Predictor	$\Delta Census_t$			$\Delta Incidence_t$		
	Estimate	<i>T</i> statistics	<i>P</i> value	Estimate	<i>T</i> statistics	<i>P</i> value
Friday	-0.0213	-1.1120	.27	0.0095	0.1024	.92
Saturday	0.0083	0.3980	.69	-0.1528	-1.5176	.13
Sunday	0.0030	0.1330	.89	-0.3340	-3.0744	.002
Monday	0.0585	2.6205	.01	-0.1939	-1.7950	.07
Tuesday	0.0291	1.3896	.17	-0.1284	-1.2655	.21
Wednesday	-0.0037	-0.1895	.85	0.0343	0.3672	.71

Model Diagnostics

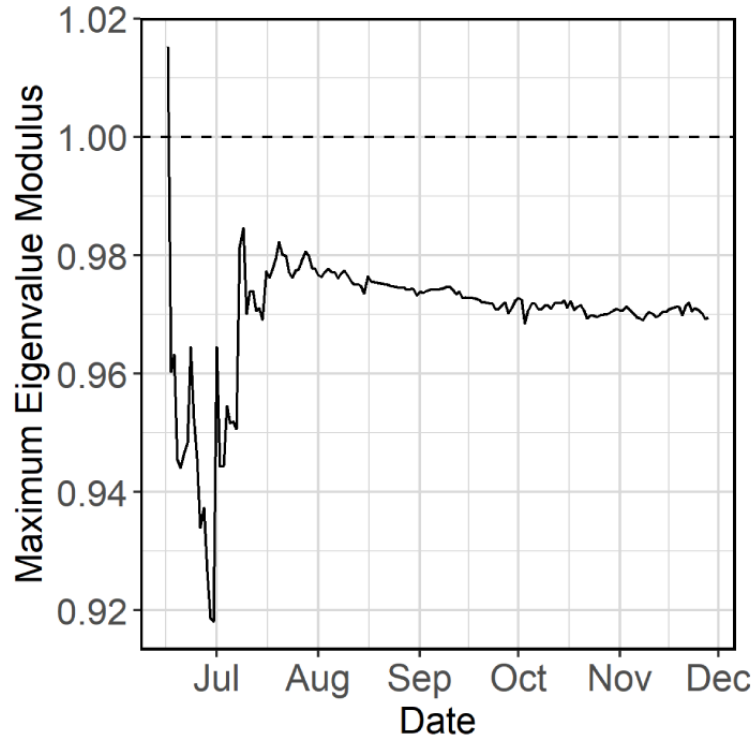
The omnibus *F* tests were significant for both census ($P<.001$) and incidence ($P<.001$) components.

The Portmanteau test did not show sufficient evidence that the errors were autocorrelated ($P=.19$). From the residual autocorrelation function and cross-correlation function plots, the correlations were within the 95% confidence band (Figure 3). The Jarque-Bera normality tests failed to reject the normality null hypothesis for the census errors ($P=.71$) but did for incidence ($P<.001$). Specifically, the incidence residuals were moderately left-skewed. The Jarque-Bera multivariate test also rejected the multivariate normality null hypothesis ($P<.001$).

The Augmented Dickey-Fuller test for stationarity of the error correction term rejected the unit root null hypothesis at the 10% significance level but failed to reject the null hypothesis at the 5% significance level (based on tabulated critical values). The KPSS test failed to reject the stationarity null hypothesis ($P=.10$). Examination of the time plot of the predicted error correction term showed no obvious departure from stationarity.

The companion matrix of the VAR representation had a maximum eigenvalue modulus of 0.97, strictly less than 1. Although this value was close to 1, the trace plot showed that this value had been slowly declining and below 1 across time when the model was fitted repeatedly in a daily rolling basis from June 16 to November 28 (Figure 4).

Figure 4. Trace plot of the maximum eigenvalue modulus for the period from June 16 to November 28, 2020.



Forecast Performance

We obtained the approximate sampling distribution of the out-of-sample MAPE from the time-series cross-validation (Figure 5). The typical value (median) of MAPE was 5.9% and the 95th percentile of MAPE was 13.4%. For the sake of comparison, the corresponding values from an ARIMA model

using the COVID-19 hospital census only were 6.6% and 14.3%. Additionally, after fitting the data from May 15 to December 5, we forecasted the census out to 7 days. Subsequently, the actual values were accurately forecasted with a MAPE of 1.9% and were all within the 80% bootstrapped forecast intervals (Figure 6).

Figure 5. Distribution of the 7-days-ahead mean absolute percentage error from the time-series cross-validation for the period from June 16 to November 28, 2020. Median (blue) and 95th percentile (red) are shown.


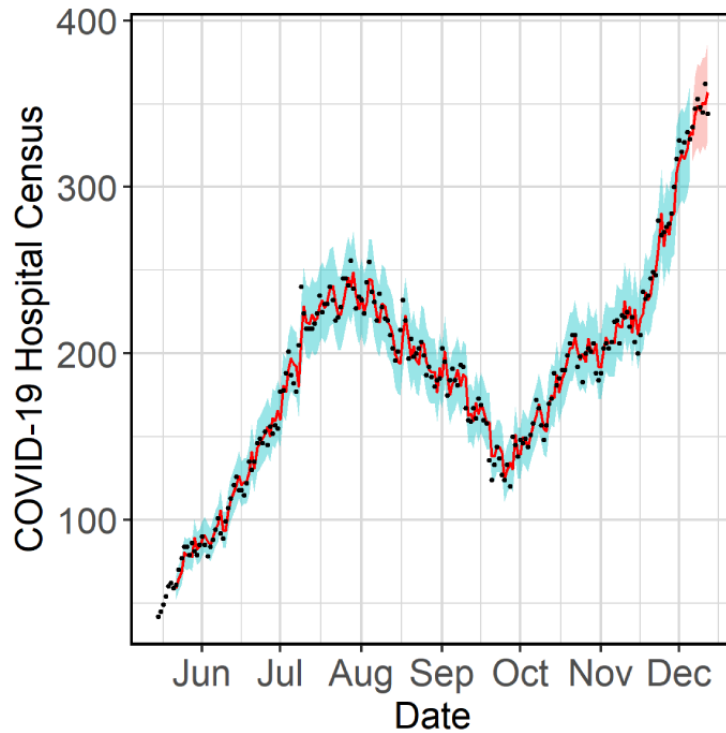


Figure 6. One-step-ahead in-sample and 7-days-ahead out-of-sample predictions for COVID-19 hospital census in the Cities Readiness Initiative region. True values (black), in-sample and out-of-sample predictions (red line), 95% prediction intervals (blue band), 80% forecast intervals (red band) are shown. The model is fitted on data from May 15 to December 5, 2020.




Scenario-Based Long-Term Forecasting

In all scenarios, due to cointegration, census followed corresponding concave trajectories with peaks occurring approximately 2 to 3 weeks later than incidence depending on

the scenario. In the worst-case scenario, census was projected to peak on February 16, 2021 (11 days later than incidence), with approximately 850 patients at the 80% forecast interval upper bound (Figure 7).

Figure 7. Worst-case-scenario, 60-day forecasts for COVID-19 hospital census in the Cities Readiness Initiative region, as of January 9, 2021. Past values (black), forecasts (red line), and 80% forecast intervals (red band) are shown.



Discussion

Principal Results

Our VECM provides a very good fit to the data and outperforms models with no or other leading indicators. Significant omnibus *F* tests showed that the model fit was better than that of a reduced VECM representation with no predictors (ie, a bivariate random walk model). When we examined model diagnostics, there was no sign of any serious departure from model assumptions. From the Portmanteau test, the errors were not different from white noise (ie, the errors do not exhibit serial correlation). Although the normality assumption (for incidence) was not met, the asymptotic properties of our estimation and hypothesis tests in the VECM would not be affected [33]. To address the possible effect of this violation on the forecast intervals, we implemented a bootstrap procedure for the forecast intervals. Both the ADF test and KPSS test showed reasonable evidence that the long-run relationship was stable. With the maximum eigenvalue modulus of the VAR representation consistently below 1 across time, the model itself was quite stable. Examining the day-of-the-week effects, we observed a higher increase in census at the beginning of the week. This agrees with our observations of hospital operations and suggests higher resource allocation when starting the week, as is also reflected in the forecasts (Figure 7). In terms of forecast performance, the VECM yielded a smaller MAPE, in terms of the median and the 95th percentile, when compared to an ARIMA model using the COVID-19 hospital census only. Our VECM also performed better than another VECM that uses two internet-based leading indicators (median MAPE of 10.5%), albeit on time domains that were partially overlapping [13].

The long-run relationship plays a crucial role in the model. Our model results show how future census responds to perturbations in the long-run cointegration relationship in the direction that would preserve the stability of the relationship. For instance, if incidence increases significantly and drives the error correction term below 0, the next-day census will tend to increase so that the error correction term will move back toward 0. Compared to short-run relationships between census change and past changes in incidence and census, the long-run relationship effect is also strongly significant and is a major driver in the model.

We observed that local infection incidence led the hospital census by about 2 weeks. The cross-correlations between incidence and census were uniformly high, between 0.7 and 0.8 at different lags, but the highest correlation was at lag 14. Clinically, we know that after someone is diagnosed with SARS-CoV-2, it can take several days before they become sick enough to be hospitalized. During the summer 2020 wave of the pandemic, incidence peaked 18 days earlier, on July 10, than when census peaked, on July 28. In the model, we also saw that past incidence changes at multiple lags have statistically significant effects on census. While previous studies have focused on other types of leading indicators [12,13], our model results and our observations demonstrate that local infection incidence can be a very effective leading indicator for COVID-19 hospital census.

Applying the model to scenario-based forecasting in a health care system is an important method for long-term forecasting when approaching an infection prevalence peak and helps determine the potential for resource capacity to be exceeded under a worst-case scenario. There are several advantages to our approach. With a scenario-based and epidemiologically informed approach, the VECM produces realistic, nonlinear, long-range trajectories of census. In contrast, an ARIMA model can have an upward linear trajectory even as we approach and arrive at the infection prevalence peak because it is agnostic to incidence. Hence, the VECM fit with scenario-based incidence will provide better accuracy since it is more reflective of pandemic behavior. Additionally, when the concern is a specific scenario, our approach is particularly useful at minimizing long-range forecast uncertainty, since the bootstrapped sample paths are constrained to fluctuate around the marginalized scenario-based census projection. Without such a constraint, 60-day forecasts can typically have wide forecast intervals that are of no practical utility.

Our study has mathematically ascertained the stable long-run relationship, that is, cointegration, between the COVID-19 hospital census and the local infection incidence, and we have developed a statistical incidence-based model to forecast the COVID-19 hospital census. In comparison, prior COVID-19 hospital capacity planning models that make use of infection incidence data rely on simplified assumptions about the incidence-census relationship. For example, in the COVID-19 Hospital Impact Model for Epidemics (CHIME) at the University of Pennsylvania [34], the ratio between hospital admissions and infection incidence is a scenario parameter defined by the user and is not time varying.

Limitations

Although our model has been thoroughly developed, it is not free of limitations. First, it is possible that we may lose the stable long-run relationship at some point in the future, either because it has run its course or due to structural changes in the time series. For instance, in the latter case, inadequate community-based testing might suddenly underestimate the actual local infection incidence, and there may be a level shift in the relationship that would have to be accounted for by a modified VECM [35,36]. In other cases, more complex structural changes may arise and be challenging to model. Second, in the future, other regions may find that the ratio between asymptomatic and symptomatic cases fluctuates considerably over time. Because case severity affects the time to hospitalization, this situation may require model revision. A potential remedy is to include both the number of asymptomatic and symptomatic cases as two leading indicators with census in a VECM in the hopes that some cointegration exists among the three variables. Third, it is relatively more difficult to fit a VECM. For univariate models such as ARIMA and exponential smoothing, well-developed R packages exist for automated model specification and estimation. With the VECM, more deliberate modeling decisions and careful checking of assumptions need to be made to fit a reliable model. Finally, the inclusion of seasonal effects in our model requires that the seasonality is deterministic. However, another health care system may find that their time-series data have stochastic seasonality

or multiple deterministic seasonality. If seasonality is not important, we potentially may resolve this by simply deseasonalizing the series. Otherwise, it may be possible to account for this with more advanced parameterization of the seasonal effects.

Conclusions

The construct presented here provides a framework in the context of a health care system for incorporating other leading indicators that may yield further increases in forecasting performance. For instance, the VECM that uses internet-based leading indicators [13] could potentially be improved by including incidence. It is also possible to incorporate other nested hospital-related time series, such as the number of intensive care units and the number of ventilators, into the

VECM if there was a need to simultaneously forecast other resources. Additionally, a VECM could be a valuable candidate for a model-averaged ensemble. This can be particularly useful if the ensemble consists only of agnostic univariate time-series models.

We have shown that infection incidence can be successfully tethered with hospital census in a multivariate time-series model to achieve accurate forecasting of COVID-19 hospital census. When coupled with scenario-based forecasting, the model helped our leaders evaluate resource capacity against different possible peak resource demands. In hindsight, our analyses correctly assured our leaders of our capability to handle a worst-case scenario, alleviated uncertainty, and effectively guided long-term planning of adequate staffing, bed capacity, and equipment supplies through the pandemic.

Authors' Contributions

HMN prepared the original draft. HMN and PJT were involved in study conceptualization, statistical analysis, and review and editing of the manuscript. ADM supervised the study and contributed to the review and editing of the manuscript.

Conflicts of Interest

ADM is an administrative member of iEnroll LLC.

References

1. Cheng ZJ, Shan J. 2019 Novel coronavirus: where we are and what we know. *Infection* 2020 Apr;48(2):155-163 [FREE Full text] [doi: [10.1007/s15010-020-01401-y](https://doi.org/10.1007/s15010-020-01401-y)] [Medline: [32072569](https://pubmed.ncbi.nlm.nih.gov/32072569/)]
2. Singhal T. A Review of Coronavirus Disease-2019 (COVID-19). *Indian J Pediatr* 2020 Apr 13;87(4):281-286 [FREE Full text] [doi: [10.1007/s12098-020-03263-6](https://doi.org/10.1007/s12098-020-03263-6)] [Medline: [32166607](https://pubmed.ncbi.nlm.nih.gov/32166607/)]
3. WHO Director-General's statement on IHR Emergency Committee on Novel Coronavirus (2019-nCoV). World Health Organization. 2020 Jan 30. URL: [https://www.who.int/director-general/speeches/detail/who-director-general-s-statement-on-ih-er-emergency-committee-on-novel-coronavirus-\(2019-ncov\)](https://www.who.int/director-general/speeches/detail/who-director-general-s-statement-on-ih-er-emergency-committee-on-novel-coronavirus-(2019-ncov)) [accessed 2020-12-29]
4. Covid-19 Map. Johns Hopkins Coronavirus Resource Center. URL: <https://coronavirus.jhu.edu/map.html> [accessed 2020-12-29]
5. The Lancet. India's COVID-19 emergency. *The Lancet* 2021 May;397(10286):1683. [doi: [10.1016/s0140-6736\(21\)01052-7](https://doi.org/10.1016/s0140-6736(21)01052-7)]
6. Earnest A, Chen MI, Ng D, Sin LY. Using autoregressive integrated moving average (ARIMA) models to predict and monitor the number of beds occupied during a SARS outbreak in a tertiary hospital in Singapore. *BMC Health Serv Res* 2005 May 11;5(1):36 [FREE Full text] [doi: [10.1186/1472-6963-5-36](https://doi.org/10.1186/1472-6963-5-36)] [Medline: [15885149](https://pubmed.ncbi.nlm.nih.gov/15885149/)]
7. Jones S, Thomas A, Evans R, Welch S, Haug P, Snow G. Forecasting daily patient volumes in the emergency department. *Acad Emerg Med* 2008 Feb;15(2):159-170 [FREE Full text] [doi: [10.1111/j.1553-2712.2007.00032.x](https://doi.org/10.1111/j.1553-2712.2007.00032.x)] [Medline: [18275446](https://pubmed.ncbi.nlm.nih.gov/18275446/)]
8. Capan M, Hoover S, Jackson E, Paul D, Locke R. Time Series Analysis for Forecasting Hospital Census: Application to the Neonatal Intensive Care Unit. *Appl Clin Inform* 2017 Dec 16;07(02):275-289. [doi: [10.4338/aci-2015-09-ra-0127](https://doi.org/10.4338/aci-2015-09-ra-0127)]
9. Zhou L, Zhao P, Wu D, Cheng C, Huang H. Time series model for forecasting the number of new admission inpatients. *BMC Med Inform Decis Mak* 2018 Jun 15;18(1):39 [FREE Full text] [doi: [10.1186/s12911-018-0616-8](https://doi.org/10.1186/s12911-018-0616-8)] [Medline: [29907102](https://pubmed.ncbi.nlm.nih.gov/29907102/)]
10. Koestler DC, Ombao H, Bender J. Ensemble-based methods for forecasting census in hospital units. *BMC Med Res Methodol* 2013 May 30;13(1):67 [FREE Full text] [doi: [10.1186/1471-2288-13-67](https://doi.org/10.1186/1471-2288-13-67)] [Medline: [23721123](https://pubmed.ncbi.nlm.nih.gov/23721123/)]
11. Jones SS, Evans RS, Allen TL, Thomas A, Haug PJ, Welch SJ, et al. A multivariate time series approach to modeling and forecasting demand in the emergency department. *J Biomed Inform* 2009 Feb;42(1):123-139 [FREE Full text] [doi: [10.1016/j.jbi.2008.05.003](https://doi.org/10.1016/j.jbi.2008.05.003)] [Medline: [18571990](https://pubmed.ncbi.nlm.nih.gov/18571990/)]
12. Berta P, Paruolo P, Verzillo S, Lovaglio PG. A bivariate prediction approach for adapting the health care system response to the spread of COVID-19. *PLoS One* 2020 Oct 15;15(10):e0240150 [FREE Full text] [doi: [10.1371/journal.pone.0240150](https://doi.org/10.1371/journal.pone.0240150)] [Medline: [33057389](https://pubmed.ncbi.nlm.nih.gov/33057389/)]
13. Turk P, Tran T, Rose G, McWilliams A. A Predictive Internet-Based Model for COVID-19 Hospitalization Census. medRxiv. Preprint posted online November 18, 2020. [FREE Full text] [doi: [10.1101/2020.11.15.20231845](https://doi.org/10.1101/2020.11.15.20231845)]
14. Lynch CJ, Gore R. Short-Range Forecasting of COVID-19 During Early Onset at County, Health District, and State Geographic Levels Using Seven Methods: Comparative Forecasting Study. *J Med Internet Res* 2021 Mar 23;23(3):e24925 [FREE Full text] [doi: [10.2196/24925](https://doi.org/10.2196/24925)] [Medline: [33621186](https://pubmed.ncbi.nlm.nih.gov/33621186/)]

15. Singh RK, Rani M, Bhagavathula AS, Sah R, Rodriguez-Morales AJ, Kalita H, et al. Prediction of the COVID-19 Pandemic for the Top 15 Affected Countries: Advanced Autoregressive Integrated Moving Average (ARIMA) Model. *JMIR Public Health Surveill* 2020 May 13;6(2):e19115 [[FREE Full text](#)] [doi: [10.2196/19115](https://doi.org/10.2196/19115)] [Medline: [32391801](https://pubmed.ncbi.nlm.nih.gov/32391801/)]
16. Zeng C, Zhang J, Li Z, Sun X, Olatosi B, Weissman S, et al. Spatial-Temporal Relationship Between Population Mobility and COVID-19 Outbreaks in South Carolina: Time Series Forecasting Analysis. *J Med Internet Res* 2021 Apr 13;23(4):e27045 [[FREE Full text](#)] [doi: [10.2196/27045](https://doi.org/10.2196/27045)] [Medline: [33784239](https://pubmed.ncbi.nlm.nih.gov/33784239/)]
17. Yeung AY, Roewer-Despres F, Rosella L, Rudzicz F. Machine Learning-Based Prediction of Growth in Confirmed COVID-19 Infection Cases in 114 Countries Using Metrics of Nonpharmaceutical Interventions and Cultural Dimensions: Model Development and Validation. *J Med Internet Res* 2021 Apr 23;23(4):e26628 [[FREE Full text](#)] [doi: [10.2196/26628](https://doi.org/10.2196/26628)] [Medline: [33844636](https://pubmed.ncbi.nlm.nih.gov/33844636/)]
18. Mehta M, Julaiti J, Griffin P, Kumara S. Early Stage Machine Learning-Based Prediction of US County Vulnerability to the COVID-19 Pandemic: Machine Learning Approach. *JMIR Public Health Surveill* 2020 Sep 11;6(3):e19446 [[FREE Full text](#)] [doi: [10.2196/19446](https://doi.org/10.2196/19446)] [Medline: [32784193](https://pubmed.ncbi.nlm.nih.gov/32784193/)]
19. Cleveland R, Cleveland W, McRae J, Terpenning I. STL: A Seasonal-Trend Decomposition Procedure Based on Loess. *Journal of Official Statistics* 1990;6:3-73 [[FREE Full text](#)]
20. Pfaff B. *Analysis of Integrated and Cointegrated Time Series With R*, 2nd Ed. New York, NY: Springer; 2008.
21. Akaike H. A new look at the statistical model identification. *IEEE Trans Automat Contr* 1974 Dec;19(6):716-723. [doi: [10.1109/TAC.1974.1100705](https://doi.org/10.1109/TAC.1974.1100705)]
22. Johansen S. Estimation and Hypothesis Testing of Cointegration Vectors in Gaussian Vector Autoregressive Models. *Econometrica* 1991 Nov;59(6):1551. [doi: [10.2307/2938278](https://doi.org/10.2307/2938278)]
23. Johansen S. *Likelihood-Based Inference in Cointegrated Vector Autoregressive Models*. Oxford; New York: Oxford University Press; 1995.
24. Johansen S, Juselius K. Maximum Likelihood Estimation and Inference on Cointegration - With Applications to the Demand for Money. *Oxford Bulletin of Economics and Statistics* 2009;52:169-210. [doi: [10.1111/j.1468-0084.1990.mp52002003.x](https://doi.org/10.1111/j.1468-0084.1990.mp52002003.x)]
25. Hyndman R, Athanasopoulos G. *Forecasting: Principles and Practice*, 3rd Edition. Melbourne, Australia: OTexts; 2021. URL: <https://otexts.com/fpp3/> [accessed 2021-07-27]
26. Incidence-Census-Model. GitHub. URL: <https://github.com/hmnguye/Incidence-Census-Model> [accessed 2021-07-28]
27. Jarque CM, Bera AK. A Test for Normality of Observations and Regression Residuals. *International Statistical Review / Revue Internationale de Statistique* 1987 Aug;55(2):163. [doi: [10.2307/1403192](https://doi.org/10.2307/1403192)]
28. Said SE, Dickey DA. Testing for unit roots in autoregressive-moving average models of unknown order. *Biometrika* 1984;71(3):599-607. [doi: [10.1093/biomet/71.3.599](https://doi.org/10.1093/biomet/71.3.599)]
29. Kwiatkowski D, Phillips PC, Schmidt P, Shin Y. Testing the null hypothesis of stationarity against the alternative of a unit root. *Journal of Econometrics* 1992 Oct;54(1-3):159-178. [doi: [10.1016/0304-4076\(92\)90104-y](https://doi.org/10.1016/0304-4076(92)90104-y)]
30. Hamilton J. *Time Series Analysis*. Princeton, NJ: Princeton University Press; 1994.
31. Wang L, Zhou Y, He J, Zhu B, Wang F, Tang L. An epidemiological forecast model and software assessing interventions on COVID-19 epidemic in China. *Journal of Data Science* 2020;18(3):409-432. [doi: [10.6339/jds.202007_18\(3\).0003](https://doi.org/10.6339/jds.202007_18(3).0003)]
32. Gardner ES, McKenzie E. Forecasting Trends in Time Series. *Management Science* 1985 Oct;31(10):1237-1246. [doi: [10.1287/mnsc.31.10.1237](https://doi.org/10.1287/mnsc.31.10.1237)]
33. Johansen S. Cointegration: Overview and Development. In: Mikosch T, Kreiß JP, Davis RA, Andersen TG, editors. *Handbook of Financial Time Series*. Berlin, Heidelberg: Springer; 2009:671-693.
34. Weissman GE, Crane-Droesch A, Chivers C, Luong T, Hanish A, Levy MZ, et al. Locally Informed Simulation to Predict Hospital Capacity Needs During the COVID-19 Pandemic. *Annals of Internal Medicine* 2020 Jul 07;173(1):21-28. [doi: [10.7326/m20-1260](https://doi.org/10.7326/m20-1260)]
35. Saikkonen P, Lütkepohl H, Lutkepohl H. Testing for the Cointegrating Rank of a VAR Process with Structural Shifts. *Journal of Business & Economic Statistics* 2000 Oct;18(4):451. [doi: [10.2307/1392226](https://doi.org/10.2307/1392226)]
36. Lutkepohl H, Saikkonen P, Trenkler C. Testing for the Cointegrating Rank of a VAR Process with Level Shift at Unknown Time. *Econometrica* 2004 Mar;72(2):647-662. [doi: [10.1111/j.1468-0262.2004.00505.x](https://doi.org/10.1111/j.1468-0262.2004.00505.x)]

Abbreviations

- ADF:** augmented Dickey-Fuller
- AIC:** Akaike information criterion
- ARIMA:** autoregressive integrated moving average
- CRI:** Cities Readiness Initiative
- HIPAA:** Health Insurance Portability and Accountability Act
- IRB:** Institutional Review Board
- KPSS:** Kwiatkowski-Phillips-Schmidt-Shin
- MAPE:** mean absolute percentage error
- SARIMA:** seasonal autoregressive integrated moving average

SARS: severe acute respiratory syndrome
STL: seasonal and trend decomposition using Loess
VAR: vector autoregressive
VECM: vector error correction model

Edited by T Sanchez; submitted 24.02.21; peer-reviewed by E Mahmoudi, R Gore; comments to author 02.06.21; revised version received 22.06.21; accepted 29.06.21; published 04.08.21

Please cite as:

Nguyen HM, Turk PJ, McWilliams AD

Forecasting COVID-19 Hospital Census: A Multivariate Time-Series Model Based on Local Infection Incidence

JMIR Public Health Surveill 2021;7(8):e28195

URL: <https://publichealth.jmir.org/2021/8/e28195>

doi: [10.2196/28195](https://doi.org/10.2196/28195)

PMID: [34346897](https://pubmed.ncbi.nlm.nih.gov/34346897/)

©Hieu M Nguyen, Philip J Turk, Andrew D McWilliams. Originally published in JMIR Public Health and Surveillance (<https://publichealth.jmir.org>), 04.08.2021. This is an open-access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work, first published in JMIR Public Health and Surveillance, is properly cited. The complete bibliographic information, a link to the original publication on <https://publichealth.jmir.org>, as well as this copyright and license information must be included.