

Learning Models for Forecasting Hospital Resource Utilization for COVID-19 Patients in Canada

Jianfei Zhang

University of Alberta

Harini Pathak

University of Alberta

Anne Snowden

University of Windsor

Russell Greiner (✉ rgreiner@ualberta.ca)

University of Alberta

Research Article

Keywords: Forecasting Hospital, TCN models, versus

Posted Date: October 21st, 2021

DOI: <https://doi.org/10.21203/rs.3.rs-962002/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Learning Models for Forecasting Hospital Resource Utilization for COVID-19 Patients in Canada

Jianfei Zhang^{1,2}, Harini Sanjay Pathak¹, Anne Snowden^{3,4}, and Russell Greiner^{1,2,*}

¹Department of Computing Science, University of Alberta, Edmonton, Alberta, Canada

²Alberta Machine Intelligence Institute (Amii), Edmonton, Alberta, Canada

³Odette School of Business, University of Windsor, Windsor, Ontario, Canada

⁴Supply Chain Advancement Network in Health (SCAN Health), Windsor, Ontario, Canada

*Corresponding author email: rgreiner@ualberta.ca

ABSTRACT

COVID-19 pandemic has overwhelmed health systems and hospital capacity in Canada. Hospitals are facing a crisis-level shortage of critical supplies (*e.g.*, hospital and ICU beds) and equipment (*e.g.*, ventilators). This motivates our need for models that can accurately forecast regional demand for hospital resources. This work aims to create predictive models that can use Canada COVID-19 data and pandemic-related factors to accurately forecast 5 quantities – three related to hospital resource utilization (*i.e.*, number of hospital beds, ICU beds, and ventilators that will be needed by COVID-19 patients) and two to the pandemic progress (*i.e.*, number of COVID-19 cases and COVID-19 deaths) – several weeks in advance. We developed a machine learning method that can use information (*i.e.*, resource utilization, pandemic progress, population mobility, weather condition, and public policy) currently known about a region since March 2020, to learn 116 temporal convolutional network (TCN) models every week; each used for forecasting the weekly average of one of these 5 quantities in Canada (respectively, in six specific provinces) for each of the next 4 weeks – *e.g.*, on 20 Nov 2020, forecasting the average number of ICU beds that will be needed by COVID-19 patients in the region of Ontario for the week ending on 18 Dec 2020. We compared our method, versus other standard models, on the COVID-19 data and hospital resource data, on the tasks of predicting these 116 values, every week from Oct 2020 to July 2021. Experimental results show that our 4640 TCN models (each forecasting a regional target for a specific future time, on a specific date) can produce accurate forecasts of demand for every hospital resource (*i.e.*, hospital beds, ICU beds, and ventilators) and pandemic progress (*i.e.*, number of COVID-19 cases and COVID-19 deaths) for each week from 2 Oct 2020 to 2 July 2021. Compared to other state-of-the-art predictive models, our TCN models yield the lowest mean absolute percentage error (MAPE). We developed and validated an accurate COVID-19 forecasting method based on the TCN models that can effectively forecast the hospital resource utilization and pandemic progress for Canada and for each of six provinces.

1 Introduction

Background As of 15 Oct 2021, the COVID-19 has infected more than 1.6 million and killed more than 28K people in Canada¹. The rapid increases in patient volumes during waves of this pandemic overwhelmed health systems and hospital capacity in many regions (*e.g.*, Ontario). Health system leaders and decision makers were unable to anticipate or predict the surge in cases during the pandemic, nor were they able to accurately prepare for the critical supplies (*e.g.*, hospital and ICU beds) and equipment (*e.g.*, ventilators) necessary to care for patients infected with the virus. There is now an urgency for creating predictive tools to accurately predict the risk for surge in demand for health services to inform leadership decisions that enable health organizations to advance preparedness efforts to enable the delivery of quality care, in a timely manner – *e.g.*, to adequately prepare hospital beds to meet patient care demands proactively (*e.g.*, four weeks in advance). Hence, many previous projects have attempted to forecast regional demand for critical care resources² – *e.g.*, personal protective equipment (PPE)³ and hospital beds⁴ in Ontario. However, there remain significant gaps in research on the hospital capacity to effectively respond to sudden surge in demand for patient care during the pandemic.

Tasks In this work, we make multiple target-region-horizon-specific forecasts involving the following 5 targets, 7 regions, and 4 time horizons (summarized in Table 1):

- Five forecasting **targets**: forecasting the number of hospital beds, ICU beds, and ventilators required to accommodate COVID-19 patients, in addition to the number of COVID-19 cases and COVID-19 deaths that is highly relevant to the surge in demand for hospital resources.
- Seven forecasting **regions**: due to the uneven spread of the disease across the Canadian provinces, it is useful to accurately forecast hospital resource utilization not only for all of Canada, but also for each province individually – here, we forecast

Table 1. The 5 targets, 7 regions, and 4 horizons of our forecasting interest.

Target	Region (Abbreviation)	Horizon (days in the future)
Hospital beds	Canada (CA)	1-week (day 1 – day 7)
ICU beds	Alberta (AB)	2-week (day 8 – day 14)
Ventilators	British Columbia (BC)	3-week (day 15 – day 21)
Cases	Manitoba (MB)	4-week (day 22 – day 28)
Deaths	Ontario (ON)	
	Québec (QC)	
	Saskatchewan (SK)	

each target for Canada and for the six major provinces: Alberta (AB), British Columbia (BC), Manitoba (MB), Ontario (ON), Québec (QC), and Saskatchewan (SK).

- Four forecasting **horizons**: forecasting each regional target for each of the next 4 weeks (*i.e.*, 1-,2-,3-,4-week, respectively) because forecasting over these time horizons is needed for adjusting hospital capacity in advance.

For each target, we focus on the weekly average and therefore define each forecasting task as predicting the 7-day average data, where 7-day average means the average daily number over a 7-day period, *e.g.*, the 7-day average number of new cases for 7 Jan is the average of the number of daily new cases for 1 Jan through 7 Jan. Consequently, the 1-,2-,3-,4-week forecasts are weekly average for 1-week (*i.e.*, day 1 [“tomorrow”] – day 7), 2-week (*i.e.*, day 8 – day 14), 3-week (*i.e.*, day 15 – day 21), and 4-week (*i.e.*, day 22 – day 28), respectively – always made on a Friday (*N.b.*: The forecasting models are trained on Fridays; see subsection 2.3). Hence, each target-region-horizon-specific forecast means the forecast of the **target** in a **region** for a specific **horizon**. Figure 1 shows an example of **ICU-ON-4** forecast – *i.e.*, the forecast of the weekly average number of **ICU** beds (**target**) needed in **Ontario** (**region**) for **4-week** (**horizon**): the **ICU-ON-4**(20 Nov 2020) model built on Friday 20 Nov 2020 predicts that Ontario may need 211 ICU beds for the week of 12-18 Dec (*i.e.*, the 7-day average number for 18 Dec is 211.), which is 4 weeks later. Figure 1 shows that this prediction is close: on average, 261 ICU beds were actually occupied in Ontario by COVID-19 patients for that week. The learned model made this forecast based on the task-specific data – the number of hospital beds and ICU beds, snowfall, and the value of other factors over the 21 days of 20 Nov 2020 (*i.e.*, 31 Oct - 20 Nov). (Table 3 presents the actual 24 factors used for this forecast, which are discussed later.) Besides this 4-week forecast number of ICU beds in Ontario, we also used other learned models to make ICU-ON-1-,2-,3 forecasts on 20 Nov 2020. (See discussion below around Figure 2.) In fact, we have 116 tasks on 20 Nov 2020 for 5 targets, 7 regions, and 4 horizons – *e.g.*, Hospital-MB-1, Ventilators-CA-2, Deaths-QC-3. Here, for each forecasting horizon, we have 29 tasks (*i.e.*, 6 provinces × 4 tasks/province + 5 Canada forecasting tasks) every week, where the provincial number of ventilators is not available.

Methodology Forecasting the resource utilization (as well as pandemic progress) is especially challenging as the spread of the disease has been highly variable over both time and region. Hence, on a weekly basis, we use different models, each with its specific individual input factors, to forecast each regional target for each time horizon – *e.g.*, the model and its input (which includes the number of cases and hospital beds in Alberta) used for the Hospital-AB-1 forecast is different from the ICU-ON-4 model and its input (*i.e.*, the number of hospital and ICU beds, snowfall, etc., in Ontario); see Figure 1.

- **Data**: Overall, we build a specific model for each forecast based on task-specific factors, which are a specific subset of the following 27 factors:
 - the 3 resource utilization factors, including the number of hospital beds, ICU beds, and ventilators occupied by COVID-19 patients in a region, *n.b.*: the provincial number of ventilators is not available,
 - the 2 pandemic progress factors (the regional number of COVID-19 cases and COVID-19 deaths),
 - other factors (see Table 3) for each region, including 5 population mobility factors, 5 weather condition factors, and 12 public policy factors.
- **TCN Models**: We built *temporal neural network* (TCN)⁵ models for the following reasons:
 - As shown in Figure 1, the input data is a set of time series, such as the number of ICU beds and deaths, the snowfall, and others. Consequently, we can view each of our forecasting tasks as a multivariate time-series forecast problem⁶ – utilizing multiple factors (*i.e.*, multiple time series) to predict a single future target value. Of course, we need to learn these models – here from the same type of data; see Table 3. As a deep learning model, TCN has been proven to perform effectively in

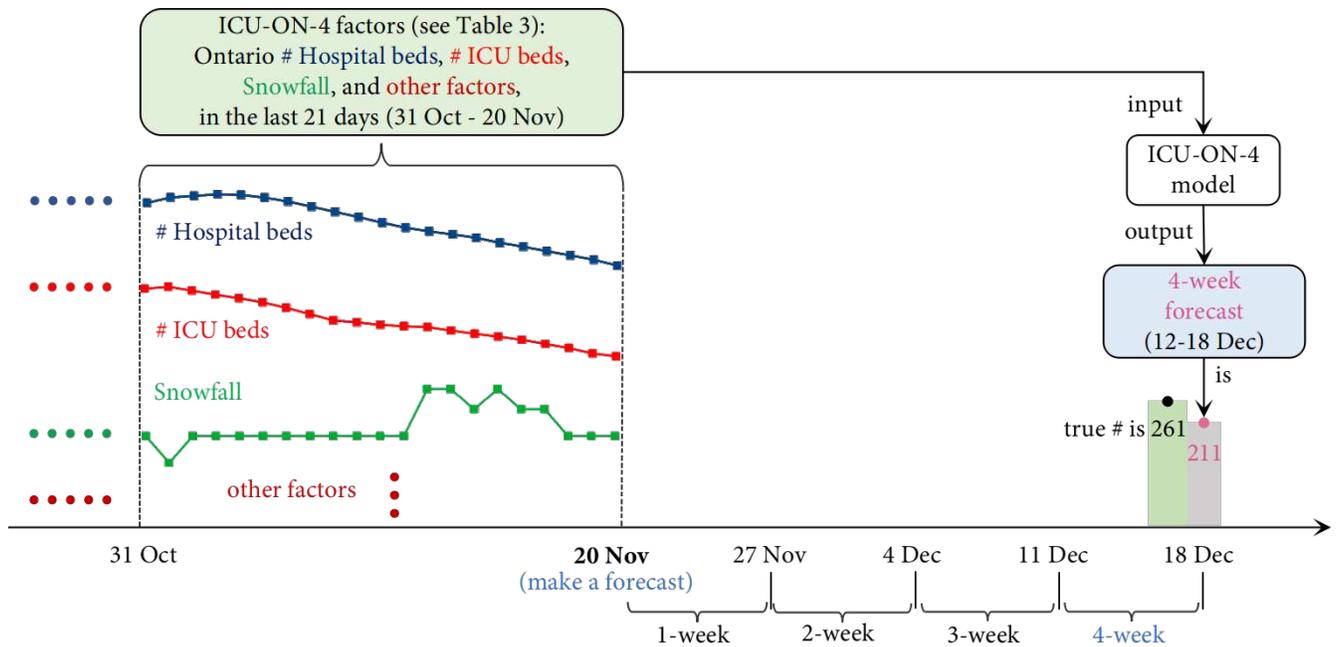


Figure 1. An example of target-region-horizon-specific forecast: forecasting the weekly average number of ICU beds (**target**) needed in Ontario (**region**) for 4-week (**horizon**).

many multivariate time-series modeling and forecasting tasks, *e.g.*, stock prediction⁷, energy forecasting⁸, and traffic flow forecasting⁹.

- TCN models benefit from their special neural network architecture (which is shown in Supplemental Figure S1), which allows this system to easily look far back the past values and explore predictive patterns from the time series. Bai *et al.*⁵ prove empirically that TCNs work better than many state-of-the-art temporal neural network models (*e.g.*, the LSTM model¹⁰) in time-series forecast.
- **Other Models:** To understand the effectiveness and reliability of the TCN models, we compare errors (defined below) yielded by TCN models, to the errors by the four widely used predictive models:
 - The “susceptible exposed infectious removed” (SEIR) compartmental model, which is an epidemiological model that is often used to describe the spread of a disease¹¹⁻¹³
 - The “autoregressive integrated moving average” (ARIMA) model, which can capture a suite of different standard temporal structures in COVID-19 time series data¹⁴
 - The “extreme gradient boosting” (XGBoost) method, which can effectively identify patterns in temporal data and often provides fairly accurate COVID-19 predictions⁴
 - The “long short-term memory” (LSTM) neural network, which is an obvious candidate for analyzing COVID-19 time series due to its notable success in sequence modeling problems¹⁰

Experiments We built 116 TCN models every week based on the weekly data to forecast the 5 targets (the weekly average number of hospital beds, ICU beds, ventilators, cases, and deaths) in each of the 7 regions (Canada and six provinces) for each of the 4 forecasting horizons (*i.e.*, 1-, 2-, 3-, 4-week thereafter). Hence, for the 40 weeks between 2 Oct 2020 and 2 July 2021, we conduct total 4640 (40 weeks × 116 tasks/week) experiments (see the previous discussion around Tasks). The experimental results below demonstrate that the TCN models can produce more accurate forecasts for each task, *i.e.*, with lower *Mean Absolute Percentage Error* (see MAPE defined by below Equation 1), in comparison with the competing models.

2 Methodology

2.1 Data Collection

During the period between 2 Oct 2020 and 2 July 2021, we collect known data (since 21 March 2020) about the daily statistics on the 27 factors (presented in Table 2), including 3 resource utilization factors, 2 pandemic progress factors, 5 population

Table 2. The 27 factors in terms of resource utilization, pandemic progress, population mobility, weather condition, and public policy. The numeric value for each factor, for each time and each region, is the average of the 7 daily values, associated with the 7th day.

Category	Total # of factors	Factors
Resource utilization	3	Hospital beds ICU beds Ventilators
Pandemic progress	2	Cases Deaths
Population mobility	5	Retail and recreation Groceries and pharmacies Parks Transit stations Residential
Weather condition	5	Average temperature Rainfall Relative humidity Dew point Snowfall
Public policy	12	School closing Workplace closing Cancel public events Restrictions on gatherings Public transport closing Stay at home Restrictions on internal movement International travel Public information campaigns Testing policy Contact tracing Facial coverings

mobility factors, 5 weather condition factors, and 12 public policy factors. The resource utilization and pandemic progress are our targets and therefore their historical data are collected for forecasting their future values. For example, we forecast the *future* number of deaths based on the *currently known* number of deaths, as well as other factors. We also collect those other factors – the mobility, weather, and policy factors – as they have effects on the average contact rates and therefore on the resource utilization and pandemic progress.

- **Resource utilization:** The daily number of hospital beds and ICU beds occupied by COVID-19 patients in Canada and in each province are provided by the geographic information system ArcGIS¹⁵. The daily number of ventilators is only available for Canada and released by the Public Health Agency of Canada¹.
- **Pandemic progress:** The daily number of new COVID-19 cases, and new COVID-19 deaths, in Canada and in each province, is provided by ArcGIS¹⁵.
- **Population mobility:** The mobility data, available from the community mobility reports¹⁶, provide insights into the daily community movement trends over time by geography, across different categories of places, such as retail and recreation, groceries and pharmacies, parks, transit stations, and residential.
- **Weather condition:** The weather data from NOAA¹⁷ include daily meteorological information about Canada and provinces: average temperature, rainfall, relative humidity, dew point, and snowfall. (Note the only field with missing values is for snowfall; we fill in the missing values via *k*-nearest-neighbor imputation¹⁸, where *k*=3.)
- **Public policy:** The policy data provided by the University of Oxford¹⁹ are measures of government daily responses to COVID-19, which are recorded on ordinal or continuous scales for 12 policy areas; see Table 2.

Each factor has a numeric value each day from 21 March 2020 to the current date (the time when we make the forecast – e.g., 20 Nov in Figure 2). Each forecast is made based on the past values of the task-specific factors (which will be discussed in subsection 2.2.2, Table 3, and Table 4).

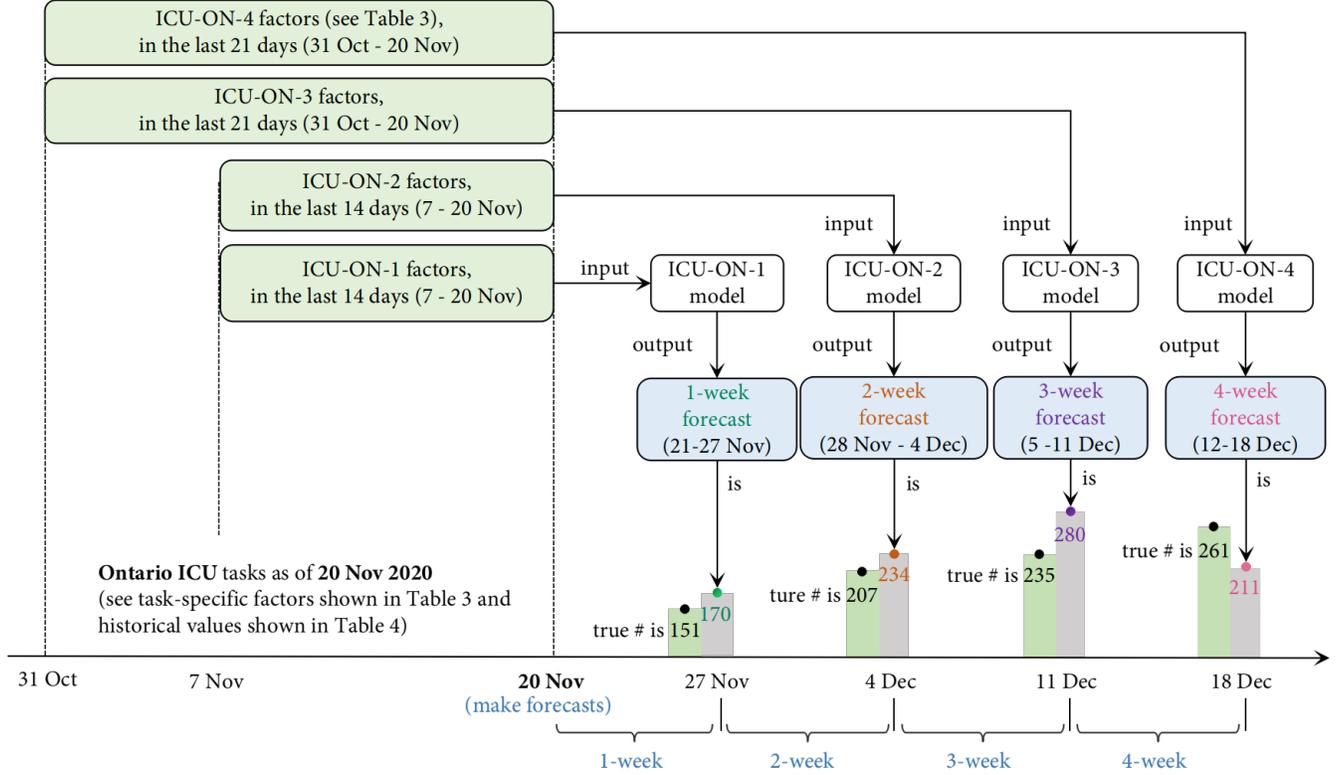


Figure 2. ICU-ON-1, ICU-ON-2, ICU-ON-3, ICU-ON-4 forecasts made on 20 Nov 2020.

2.2 Forecasting Task

2.2.1 Outputs

Every week, we learned 116 different forecasting models. Figure 2 (which extends Figure 1) presents the ICU-ON-1,-2,-3,-4 forecasts made by the four models based on the data known on 20 Nov 2020. Each of these four different models uses the known values of its own factors (shown in Table 3) and specific past values (shown in Table 4) to forecast the number of ICU beds that will be needed in Ontario, which is predicted to be 170 for 1-week (21 – 27 Nov), 234 for 2-week (28 Nov – 4 Dec), 280 for 3-week (5 – 11 Dec), and 230 for 4-week (12 – 18 Dec). The true number of ICU beds for these 4 weeks is 151, 207, 235, and 261, respectively.

- **MAPE:** For each task, we perform the weekly forecasts (on every Friday) for the 40 weeks from 2 Oct 2020 to 2 July 2021. Then, we compute the mean absolute percent error (MAPE)²⁰ of the 40 forecasts for each task. (Note smaller MAPE values are better.) Letting T_i and F_i^H (H indicates the forecasting horizon – here, $H \in \{1, 2, 3, 4\}$ for 1-,2-,3-,4-week) be respectively the true and the forecast target (e.g., number of ICU beds) for the i^{th} ($i \in \{1, \dots, 40\}$) week, then the MAPE of the 40 H -week forecasts (each predicted H week(s) in advance) is

$$\text{MAPE}(40 \text{ target-region-horizon-specific forecasts}) = \text{MAPE}\left(\left(F_1^H, T_1\right), \dots, \left(F_{40}^H, T_{40}\right)\right) = \frac{1}{40} \sum_{i=1}^{40} \left| \frac{F_i^H - T_i}{T_i} \right| \times 100. \quad (1)$$

- **Example:** Figure 3 shows six ICU-ON-4 forecasts made from 20 Nov (which is the example shown in Figure 1) to 18 Dec 2020. Every forecast for a time point is made 4 weeks before that time. The MAPE of these six 4-week forecasts is

$$\begin{aligned} \text{MAPE}(6 \text{ ICU-ON-4 forecasts}) &= \text{MAPE}\left(\left(F_1^4, T_1^4\right), \dots, \left(F_6^4, T_6^4\right)\right) \\ &= \text{MAPE}\left(\left(211, 261\right), \dots, \left(405, 383\right)\right) = \frac{1}{6} \left(\frac{|211 - 261|}{261} + \dots + \frac{|405 - 383|}{383} \right) = 7.9\%. \end{aligned}$$

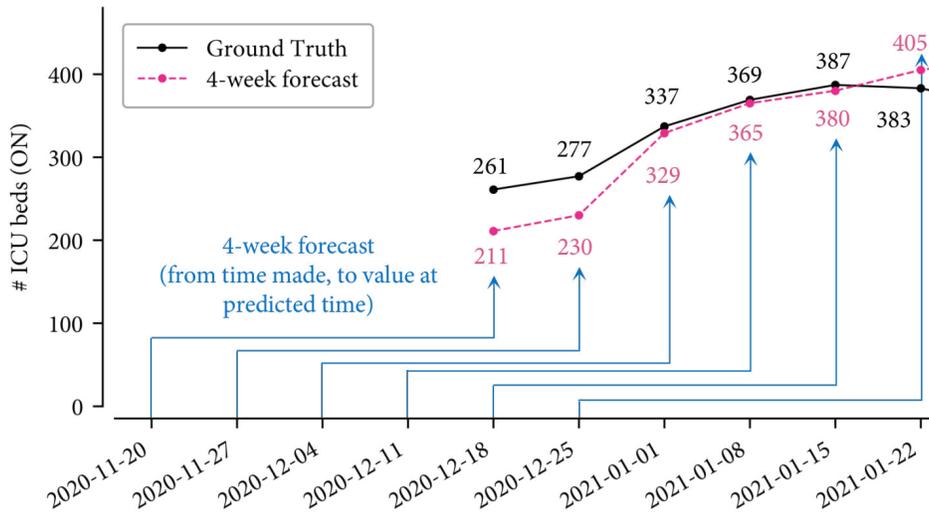


Figure 3. The true and ICU-ON-4 forecast number of ICU beds between 18 Dec 2020 and 22 Jan 2021.

2.2.2 Inputs

Each of the 116 different forecasting models uses as input a different set of factors – *i.e.*, specific factors and their specific past values. (We present the performance of TCN models using various sets of factors and different past values in Supplemental Figure S2.)

Table 3. The factors used for various target-region-horizon-specific forecasting tasks.

Forecasting Task			Factors used for Forecasting Task							
Target	Region	Horizon	Resource			Pandemic		Mobility	Weather	Policy
(# of)	(7 regions)	(week)	Hosp.	ICU	Vent.	Cases	Deaths	5 factors	5 factors	12 factors
Hosp.	for each	1 2, 3, 4	✓ ✓			∨ ∨		†	†	†
ICU	for each	1 2, 3, 4	‡ ‡	✓ ✓			•	†	†	†
Vent.	for each	1 2, 3, 4	‡ ‡		✓ ✓		•	†	†	†
Cases	for each	1 2, 3, 4				✓ ✓		†	†	†
Deaths	for each	1, 2 3, 4		★			✓ ✓			

Task-specific factors: For every target, we decide to include that target factor (*e.g.*, we use the *currently known* number of deaths in the current week, when estimating the *future* number of deaths) and some other factors (shown in Table 3), rather than all the factors (shown in Table 2) for every forecasting task. We select these *task-specific* factors based on our prior knowledge and causal assumption about these factors:

- We assume the change in mobility, weather, and policy will affect *future* (*i.e.*, 1 week later) pandemic progress and hospital resource demand – *e.g.*, adjusting the policy *international travel restrictions* may not alter the trend of resource utilization until 1 week later. Hence, we use these factors for all 2-,3-,4-week forecasts (see † in Table 3, the death forecast excluded), rather than for 1-week forecast.
- For each time series, we assume that the preceding and future values are highly correlated. Hence, to forecast the future target values (*e.g.*, number of ICU beds), we use currently known target values (*e.g.*, currently known number of ICU beds), see ✓ in Table 3.
- We assume the five target factors are causally related.

- The projected demand of hospital beds for 1-,2-,3-,4-week is always dependent on current number of cases because a large number of cases may cause more hospitalizations in next 1-,2-,3-,4-week, respectively (see \vee in Table 3).
- The projected demand of ICU beds and ventilators for 1-,2-,3-,4-week relies on the current hospitalizations (see \ddagger), where the 1-week demand is also affected by the number of deaths (see \bullet in Table 3), as the ICU beds and ventilators would be available soon after the death of patients in ICU or on ventilators.
- Considering ICU patients are usually at high risk of death, we assume the death occurs within 2 weeks after entering ICU. Therefore, we use the number of occupied ICU beds for forecasting number of deaths in 1 or 2 weeks (see \star in Table 3), but instead use the number of cases for 3-week and 4-week forecasts (see \circ in Table 3), as we assume that a few of the patients who were infected, will die 3 or 4 weeks later.

Task-specific values: Given a learned TCN model, it uses the values of all the *task-specific* factors in the past few days to make a forecast (*i.e.*, for each of the 5 targets and for each of the 7 regions), as shown in Table 4:

- To make a 1-week forecast for each regional target, the learned 1-week TCN model uses the values in the past 14 days of the prediction time – *e.g.*, to make the ICU-ON-1 forecast shown in Figure 2, the trained ICU-ON-1(20 Nov) model uses the values of the ICU-ON-1-specific factors in the past 14 days (7 - 20 Nov). Similarly, the learned 2-week TCN model also uses the values of the task-specific factors in the past 14 days, for 2-week forecast.
- To make a 3-week forecast for each regional target, the learned 3-week TCN model uses the values in the past 21 days of the prediction time – *e.g.*, to make the ICU-ON-3 forecast shown in Figure 2, the trained ICU-ON-3(20 Nov) model uses the values of the ICU-ON-3-specific factors in the past 21 days (31 Oct - 20 Nov). Similarly, the learned 4-week TCN model also uses the past 21-day values of the task-specific factors, for 4-week forecast.

In the learning process, the TCN model explores the predictive patterns – *i.e.*, the relationship between the past factor values (*i.e.*, the factor values in the past 2 or 3 weeks) and future target values (*i.e.*, the target values in the next 1, 2, 3, and 4 weeks). (Note Bai *et al.*⁵ demonstrated that using additional past factor values – *e.g.*, the values in past 10 weeks – increases the computation cost but does not further improve the model’s predictive ability in a time series forecasting task.)

Table 4. The values of task-specific factors used as input for different forecasting tasks.

Target	Region	Horizon (week)	#of past values of the task-specific factors
for each	for each	1, 2	values in the past 14 days
		3, 4	values in the past 21 days

2.3 Training

At each prediction time (*i.e.*, every Friday), we need to train 116 TCN models, each for a target-region-horizon-specific forecasting task. As TCN is a supervised model, it cannot directly learn from time-series data – which means we must first transform the given time series into supervised data – *i.e.*, a set of (*input,output*) pairs that the TCN-learner (*i.e.*, training algorithm) can use to produce a model. Afterwards, given the input – *i.e.*, the values of the relevant factors specified in Table 3 and Table 4 – this model forecasts the output.

An example of (*input,output*) pair To train the ICU-ON-4(20 Nov 2020) model shown in Figure 1, we use the (*input,output*) pairs as of Friday 20 Nov 2020 as training instances; for each (*input,output*) pair, the time lag between *input* and *output* is 4 weeks. Here, the last training *output* is the number of ICU beds for 20 Nov 2020 (which is 128) and the corresponding training *input* is a linearized version of the matrix (shown in Equation 2) composed of 24 factors (*i.e.*, the number of hospital beds and ICU beds, snowfall, facial coverings, etc.) for each of the 21 days from 3 Oct to 23 Oct 2020 – hence, the *input* is a $504 = 24 \times 21$ element vector (*i.e.*, 24 factors \times 21 values/factor).

$$\begin{array}{l}
 \# \text{ Hospital beds} \\
 \# \text{ ICU beds} \\
 \text{Snowfall} \\
 \vdots \\
 \text{Facial coverings}
 \end{array}
 \begin{pmatrix}
 145 & 153 & \dots & 263 & 265 \\
 34 & 36 & \dots & 71 & 73 \\
 10.16 & 10.16 & \dots & 28.45 & 16.93 \\
 \vdots & \vdots & \dots & \vdots & \vdots \\
 1.0 & 1.0 & \dots & 1.0 & 1.0
 \end{pmatrix}
 \begin{array}{l}
 3 \text{ Oct} \\
 4 \text{ Oct} \\
 \dots \\
 22 \text{ Oct} \\
 23 \text{ Oct}
 \end{array}
 \quad (2)$$

This $(input, output) = ([145, 153, \dots, 265, 34, 36, \dots, 73, 10.16, \dots, 16.93, \dots, 1.0, \dots, 1.0], 128)$ pair is a single labeled training instance; other ICU-ON-4 labeled training instances correspond to other dates – e.g., there is another pair whose input is a 504-tuple and output is 123, for Ontario ICU beds for 19 Nov, and so forth, for 196 dates (i.e., every day from 8 May 2020 through 20 Nov 2020). (N.b.: While we only make forecasts for Friday, our training data corresponds to forecasts made on all days of the week) Together, these pairs form the training set used for training the ICU-ON-4(20 Nov) model. We then learn a TCN model from this training data (which involves 196 instance pairs, where the input of each pair is 504-dimensional). Note this specific learned TCN model will only make the ICU-ON-4 forecast for 18 Dec, given the input – the element vector in terms of the same task-specific factors for 21 days from 31 Oct to 20 Nov 2020 (see Figure 1). (N.b.: this ICU-ON-4(20 Nov) model differs from the other ICU-ON-4 models built to make predictions for different times – e.g., it is different from ICU-ON-4(27 Nov), etc.) Here, we use a sliding window method²¹ – e.g., for ICU-ON-4 forecast, we use a 21-day input window and a 1-day output window (which is 4 weeks apart from the input window) to slice each time series: the two windows move forward simultaneously over the ICU-ON-4 factors and generate the $(input, output)$ pairs.

Training and Validation After the above data transformation, we divide the known data (i.e., training instances) into a training set and a disjoint validation set, which are used to learn and optimize a TCN model. We learn a model on the training set, i.e., using the $(input, output)$ pairs to estimate the model parameters (i.e., the weights of a neural network; see the discussion around Supplemental Figure S1). In general, this learning process involves a set of hyperparameters, such as the network depth and kernel size (which will be discussed later). We seek the hyperparameters that produce the most accurate model that can make accurate forecasts for the validation set. The validation mimics the out-of-sample forecasting scenario and reduces the change that the optimized model will overfit²² when it makes a forecast given the new input. Motivated by our recently developed COVID-19 forecast method – LaPoFaPo²³ – we use the “most recent” 10% of the training instances as validation data – e.g., in the example shown in Figure 1, of the data known as of 20 Nov 2020, we use the $(input, output)$ pairs whose outputs correspond to the dates before 23 Oct 2020 as training data and the pairs whose outputs correspond to the 28 days between 23 Oct and 20 Nov 2020 as the validation data.

Settings for TCN Learning We implement the TCN in TensorFlow²⁴, with a 0.01 learning rate and the sigmoid activation²⁵. Each TCN model is trained for 500 epochs by using the Adam optimizer²⁶. We employ the Bayesian optimizer²⁷ to choose a value from {2, 3, 4, 5, 6, 7} for kernel size (i.e., how many different kernel weights used in the TCN, see Supplemental Figure S1 for more details), a value from {2, 4, 8, 16} for dilation rate (i.e., d in Supplemental Figure S1), and a dropout rate in the range [0.01, 0.1], so that the TCN model can lead to the lowest loss on the validation data. Here, the dropout implements a regularization that can improve the model performance (see the results shown in Supplemental Figure S2c).

Other Models and Their Settings Additionally, we train SEIR¹³, ARIMA¹⁴, XGBoost⁴, and LSTM¹⁰ models every week to forecast resource utilization in Canada. For SEIR, we set the average incubation period and average infection period as 5 days and 10 days, respectively. We set the number of autoregressive terms, nonseasonal differences, and lagged forecast errors for ARIMA as 4, 0, and 1, respectively. For XGBoost, we set maximum depth to 2, learning rate to 0.2, tree estimators to 150, observation fraction and column fraction to 0.9. We train the LSTM model with two hidden layers, each including 20 fully-connected neurons, for 500 epochs via the Adam optimizer²⁶ with a 0.1 learning rate. To match our TCN system, these four competitors used the task-specific factors (shown in Table 3) for training and forecasting. LSTM uses the past values shown in Table 4.

3 Results

TCN’s forecasts Figure 4 shows the forecast numbers (dotted lines) of hospital beds, ICU beds, and ventilators utilized by COVID-19 patients (as well as cases and deaths) during the period between 2 Oct 2020 and 2 July 2021, in comparison with their true numbers (solid lines). We see that TCN can accurately forecast resource utilization in advance: for any target in any region, the forecasts are close to the true values, throughout the forecasting period. As expected, we see that forecasting further in the future (e.g., 4-week) is more challenging than forecasting the near future (e.g., 1-week). We also present the forecasts for six provinces in Supplemental Figures S3-S8.

MAPE Table 5 presents the MAPEs of the forecast number of hospital beds, ICU beds, ventilators, cases, and deaths (respectively called ‘Hosp.’, ‘ICU’, ‘Vent.’, ‘Cases’, ‘Deaths’) in Canada and in each of the six provinces. (Note that we could not forecast the provincial number of ventilators.) The low MAPEs indicate that TCN can effectively offer health providers accurate forecasts of resource utilization and therefore allow them to precisely prepare the needed hospital resources in advance.

Comparison between TCN and other models In Figure 5, we compare the MAPEs yielded by our TCN models and the four competitors (LSTM, SEIR, XGBoost, and ARIMA). Our models consistently outperform all these competitors for all forecasting tasks. In particular, compared to LSTM, TCN achieves an average of 10% MAPE decrease (across the four horizons) for

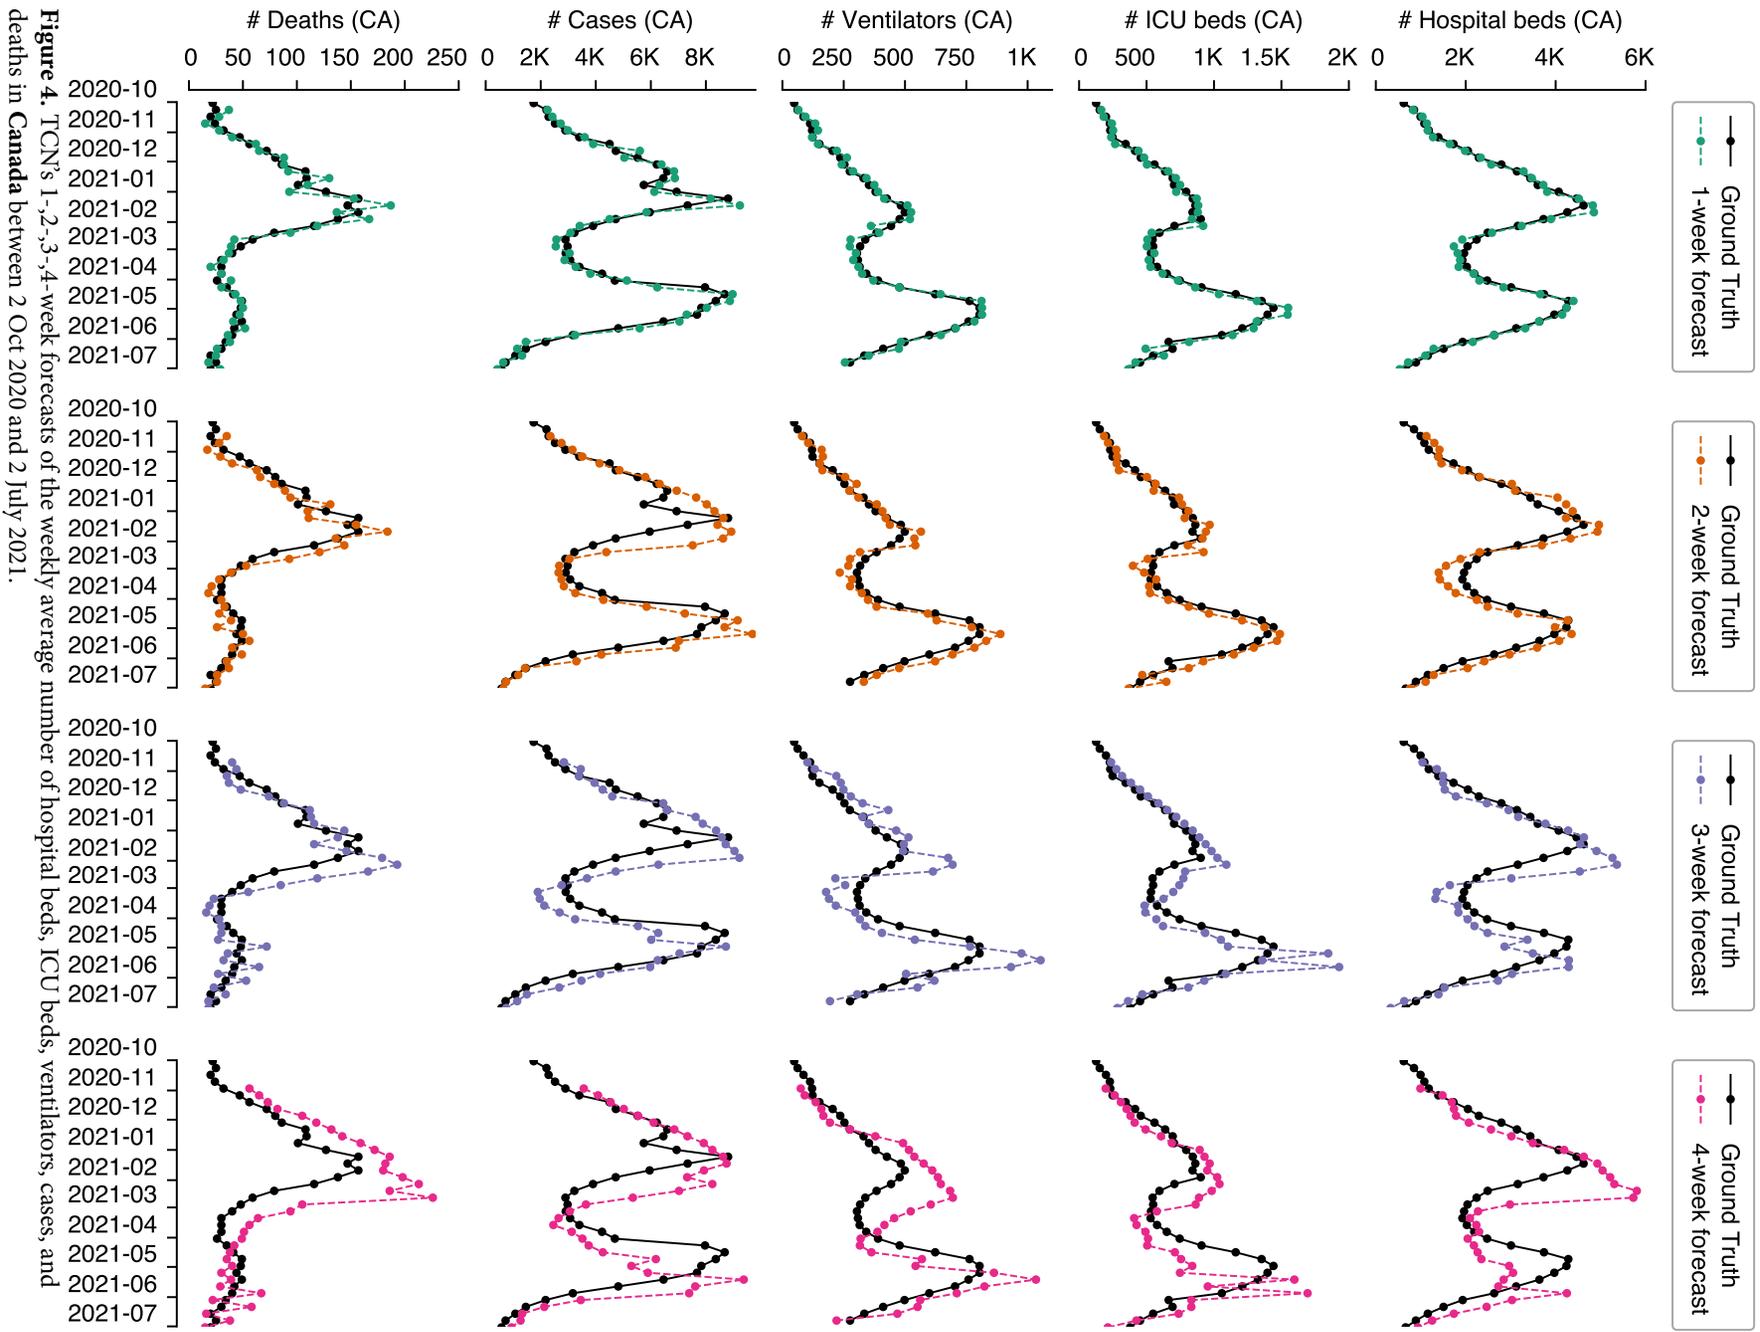


Table 5. MAPEs (%) of TCN's 1-,2-,3-,4-week forecasts for **Canada** (CA) and **six provinces** (AB, BC, MB, ON, QC, SK) between 2 Oct 2020 and 2 July 2021.

Target		Hosp.	ICU	Vent.	Cases	Deaths	Hosp.	ICU	Vent.	Cases	Deaths
Horizon		1-week					2-week				
Region	CA	6.0	8.38	6.55	10.21	17.72	14.23	13.75	13.29	19.38	22.51
	AB	8.53	16.39	-	15.67	45.55	18.81	22.14	-	29.61	46.08
	BC	8.27	13.33	-	12.50	40.61	16.81	21.65	-	26.51	34.00
	MB	14.68	23.33	-	19.07	39.15	24.88	28.66	-	34.31	48.73
	ON	29.35	9.45	-	12.30	32.77	19.05	16.93	-	20.55	28.60
	QC	8.92	13.89	-	14.22	21.81	18.09	18.80	-	22.43	27.79
	SK	14.89	27.57	-	18.90	38.14	26.56	30.00	-	23.82	39.21
Horizon		3-week					4-week				
Region	CA	22.55	20.07	27.48	29.98	34.67	32.42	27.26	33.26	37.72	56.53
	AB	24.68	34.79	-	42.69	61.77	37.13	45.80	-	52.62	97.42
	BC	25.13	33.18	-	28.53	53.42	29.85	28.70	-	37.72	75.64
	MB	31.62	46.23	-	37.78	37.68	52.28	50.36	-	66.11	100.81
	ON	29.97	22.88	-	32.13	52.12	38.20	28.76	-	27.14	51.88
	QC	24.39	37.89	-	35.66	45.76	33.60	35.04	-	42.95	84.08
	SK	30.23	35.65	-	39.33	58.05	35.33	45.23	-	40.84	79.19

forecasting the number of hospital beds, 4% decrease for ICU beds, and 5% decrease for ventilators. These comparisons reveal the effectiveness of our TCN models.

More results yielded by TCN's variants We also explore other variants of the TCN learning algorithm with different settings – different model inputs (which were discussed in subsection 2.2.2) – *e.g.*, using all factors (rather than the task-specific factors shown in Table 3) and different number of earlier weeks for each factor (not only the values in the past 14 or 21 days, as shown in Table 4) – and different model hyperparameters – *e.g.*, changing the dropout rate. The results of these empirical studies (shown in Supplemental Figure S2) for Canada during the period between 2 Oct 2020 and 2 Jul 2021 demonstrate that our settings for TCN models are effective.

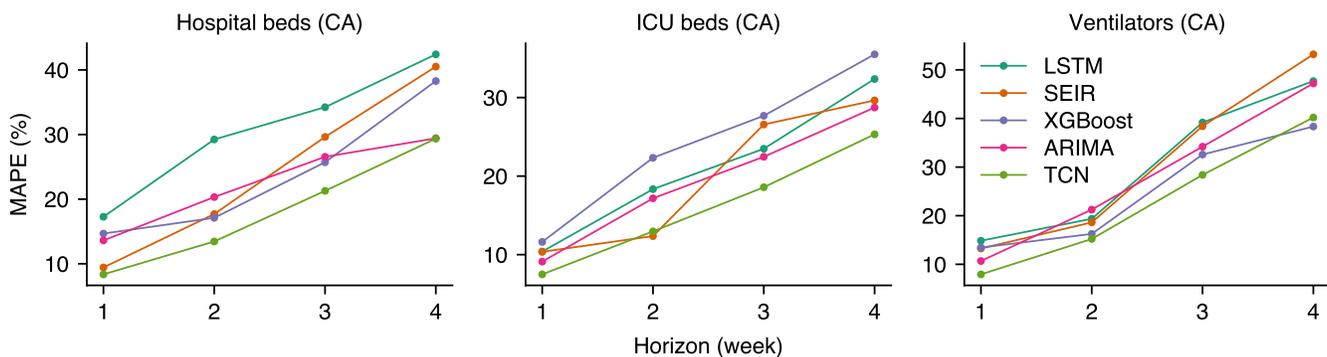


Figure 5. Models' performance, in terms of MAPE (%) of the 1-,2-,3-,4-week forecasts of resource utilization in **Canada**.

4 Conclusions

This paper provided a TCN-based predictive method to forecast the health system capacity required to respond to surge in demand, including hospital beds, ICU beds, and ventilators, which are required to effectively respond to the pandemic. The modeling outcomes reflect the complex interplay of critical factors that can offer decisions makers a predictive and proactive tool that considers data and information about regional mobility, weather, and policy. We have provided a method for effectively learning these TCN models for various target-region, and horizon-specific forecasts, including the weekly average number of hospital beds, ICU beds, ventilators, cases, and deaths in Canada and six provinces (AB, BC, MB, ON, QC, SK) for up-to-four weeks in the future. The numerous experiments demonstrated that our method is more accurate (in terms of MAPE) than four state-of-the-art predictive models. We also demonstrated that our method can accurately forecast the weekly average number of cases and deaths in the future.

References

1. COVID-19 daily epidemiology update (october 15, 2021). <https://health-infobase.canada.ca/covid-19/epidemiological-summary-covid-19-cases.html>.
2. Rahimi, I., Chen, F. & Gandomi, A. H. A review on COVID-19 forecasting models. *Neural Comput. Appl.* 1–11 (2021).
3. Shoukat, A. *et al.* Projecting demand for critical care beds during COVID-19 outbreaks in Canada. *Can. Med. Assoc. J.* **192**, E489–E496 (2020).
4. Snider, B. *et al.* Artificial intelligence to predict the risk of mortality from COVID-19: Insights from a Canadian application. *medRxiv* (2020).
5. Bai, S., Kolter, J. Z. & Koltun, V. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling (2018).
6. Torres, J. F., Hadjout, D., Sebaa, A., Martínez-Álvarez, F. & Troncoso, A. Deep learning for time series forecasting: A survey. *Big Data* **9**, 3–21 (2021).
7. Deng, S. *et al.* Knowledge-driven stock trend prediction and explanation via temporal convolutional network. In *WWW*, 678–685 (2019).
8. Lara-Benítez, P., Carranza-García, M., Luna-Romera, J. M. & Riquelme, J. C. Temporal convolutional networks applied to energy-related time series forecasting. *Appl. Sci.* **10**, 2322 (2020).
9. Zhao, W. *et al.* Deep temporal convolutional networks for short-term traffic flow forecasting. *IEEE Access* **7**, 114496–114507 (2019).
10. Shahid, F., Zameer, A. & Muneeb, M. Predictions for COVID-19 with deep learning models of LSTM, GRU and Bi-LSTM. *Chaos, Solitons & Fractals* **140**, 110212 (2020).
11. Radulescu, A., Williams, C. & Cavanagh, K. Management strategies in a SEIR-type model of COVID-19 community spread. *Sci. Reports* **10**, 1–16 (2020).
12. López, L. & Rodo, X. A modified SEIR model to predict the COVID-19 outbreak in Spain and Italy: simulating control scenarios and multi-scale epidemics. *Results Phys.* **21**, 103746 (2021).
13. Arik, S. Ö. *et al.* Interpretable sequence learning for Covid-19 forecasting. In *NIPS* (2020).
14. Malki, Z. *et al.* ARIMA models for predicting the end of COVID-19 pandemic and the risk of second rebound. *Neural Comput. Appl.* **33**, 2929–2948 (2021).
15. ArcGIS. COVID-19 Resources, Canadian Outbreak At-A-Glance. <https://resources-covid19canada.hub.arcgis.com>.
16. Google. Community Mobility Reports. <https://www.google.com/covid19/mobility>.
17. National Centers for Environmental Information. <https://www.ncei.noaa.gov>.
18. Batista, G. E., Monard, M. C. *et al.* A study of K-nearest neighbour as an imputation method. *HIS* **87**, 48 (2002).
19. Variation in government responses to COVID-19. <https://www.bsg.ox.ac.uk/research/publications/variation-government-responses-covid-19>.
20. Armstrong, J. S. & Collopy, F. Error measures for generalizing about forecasting methods: Empirical comparisons. *Int. journal forecasting* **8**, 69–80 (1992).
21. Chu, C.-S. J. Time series segmentation: A sliding window approach. *Inf. Sci.* **85**, 147–173 (1995).
22. Hawkins, D. M. The problem of overfitting. *J. Chem. Inf. Comput. Sci.* **44**, 1–12 (2004).
23. Ramazi, P. *et al.* Accurate long-range forecasting of covid-19 mortality in the usa. *Sci. Reports* **11**, 1–11 (2021).
24. Tensorflow-TCN. <https://github.com/Baichenjia/Tensorflow-TCN>.
25. Narayan, S. The generalized sigmoid activation function: Competitive supervised learning. *Inf. Sci.* **99**, 69–82 (1997).
26. Kingma, D. P. & Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
27. Wu, J. *et al.* Hyperparameter optimization for machine learning models based on Bayesian optimization. *J. Electron. Sci. Technol.* **17**, 26–40 (2019).

Acknowledgments

This work was partially supported by the Canadian Institute of Health Research (CIHR) Operating Grant: COVID-19 May 2020 Rapid Research Funding Opportunity - Social Policy and Public Health Responses. RG gratefully acknowledges funding from the Natural Sciences and Engineering Research Council of Canada (NSERC) and Alberta Machine Intelligence Institute (Amii). We also acknowledge Google Cloud platform and data providers ArcGIS and Public Health Agency of Canada.

Author contributions

J.Z. designed the modeling system, implemented experiments, and prepared all figures. J.Z. and H.P. performed data collection and preprocessing. J.Z. and R.G. led the interpretation of results and the writing of the manuscript. All authors reviewed the manuscript.

Data availability

All the data used in our analyses are available online, where the links have been presented in the paper.

Competing interests

The authors declare no competing interests.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [Supplement.pdf](#)