

Machine learning-based short-term forecasting of COVID-19 hospital admissions using routine hospital patient data

Martin S. Wohlfender^{a,b,c,*}, Judith A. Bouman^{a,c}, Olga Endrich^{d,e,f}, Alban Ramette^{c,h}, Alexander B. Leichtle^{c,i}, Guido Beldi^{c,g}, Christian L. Althaus^{a,c}, Julien Riou^{a,c,j}

^a*Institute of Social and Preventive Medicine, University of Bern, Bern, Switzerland*

^b*Graduate School for Cellular and Biomedical Sciences, University of Bern, Bern, Switzerland*

^c*Multidisciplinary Center for Infectious Diseases, University of Bern, Bern, Switzerland*

^d*Medical Directorate, Inselspital, Bern University Hospital, Bern, Switzerland*

^e*Institute of Clinical Chemistry, Bern University Hospital, Bern, Switzerland*

^f*Bern Center for Precision Medicine, University of Bern, Bern, Switzerland*

^g*Department for Visceral Surgery and Medicine, Bern University Hospital, Bern, Switzerland*

^h*Institute for Infectious Diseases, University of Bern, Bern, Switzerland*

ⁱ*Center of Laboratory Medicine, Bern University Hospital, Bern, Switzerland*

^j*Department of Epidemiology and Health Systems, Unisanté, Center for Primary Care and Public Health & University of Lausanne, Lausanne, Switzerland*

Abstract

During the COVID-19 pandemic, the field of infectious disease modeling advanced rapidly, with forecasting tools developed to track trends in transmission dynamics and anticipate potential shortages of critical resources such as hospital capacity. In this study, we compared short-term forecasting approaches for COVID-19 hospital admissions that generate forecasts one to five weeks ahead, using retrospective electronic health records. We extracted different features (e.g., daily emergency department visits) from an individual-level patient dataset covering six hospitals located in the region of Bern, Switzerland from February 2020 to June 2023. We then applied five methods – last-observation carried forward (baseline), linear regression, XGBoost

*Corresponding author

Email address: martin.wohlfender@unibe.ch (Martin S. Wohlfender)

and two types of neural networks – to time series using a leave-future-out training scheme with multiple cutting points and optimized hyperparameters. Performance was evaluated using the root mean square error between forecasts and observations. Generally, we found that XGBoost outperformed the other methods in predicting future hospital admissions. Our results also show that adding features such as the number of hospital admissions with fever and augmenting hospital data with measurements of viral concentration in wastewater improves forecast accuracy. This study offers a thorough and systematic comparison of methods applicable to routine hospital data for real-time epidemic forecasting. With the increasing availability and volume of electronic health records, improved forecasting methods will contribute to more precise and timely information during epidemic waves of COVID-19 and other respiratory viruses, thereby strengthening evidence-based public health decision-making.

Keywords:

COVID-19, hospital admissions, forecasting, local level, machine learning, electronic health records, wastewater

1. Introduction

The COVID-19 pandemic has highlighted the need for reliable infectious disease monitoring and forecasting systems. As SARS-CoV-2 spread globally, researchers, healthcare professionals, public health authorities, and governments undertook extensive efforts to mitigate its impact and control transmission dynamics. A key priority was to ensure that hospital capacity, particularly in intensive care units, was not exceeded. Whenever hospital capacities were exceeded, hospitals were forced to implement crisis care standards, including treatment protocol classifications that prioritized patients with the highest probability of survival, often leading to delayed or reduced care for other patients. This also resulted in the postponement of elective procedures, in increased stress and burnout among healthcare workers, and in higher mortality rates due to limited access to critical resources such as ICU beds and ventilators (Anderegg et al., 2022; Didriksson et al., 2022). In such situations, short-term forecasts aimed at anticipating new hospital admissions a few weeks in advance can be invaluable for public health decision makers and hospital management.

Researchers worldwide have applied numerous approaches to forecast the spread and impact of SARS-CoV-2 in different settings based on various types of data. Due to increasing digitization, substantially more data reflecting various aspects of the pathogen were collected during the COVID-19 pandemic compared to the historic major infectious disease outbreaks. Examples of such data are wearable or smartphone sensor data (Grantz et al., 2020), viral genome sequences (Shu and McCauley, 2017; Furuse, 2021; CDC, 2024; Hodcroft et al., 2025), viral load measurements in wastewater (Morvan et al., 2022; Jahn et al., 2022), and electronic health records from hospitals and medical practices (Qian et al., 2021). A wide variety of methods have been used to produce forecasts, including mechanistic models (e.g., deterministic or stochastic compartmental models, agent-based models), statistical time series models (e.g., ARIMA, exponential smoothing, regression) and machine learning methods (e.g., tree-based models, neural networks) (Kraemer et al., 2025). In both the United States and Europe, groups of scientists developed standardized forecasting pipelines for COVID-19 cases, hospital admissions, and deaths in different geographic regions (Cramer et al., 2021; ECDC, 2021, 2023). This allowed the combination of multiple models from different groups into an ensemble forecast with a single cone of uncertainty.

Hospital capacity is an important indicator when planning public health

38 interventions during major outbreak of an infectious disease. Therefore, mod-
39 els that provide estimates of expected admissions to hospitals on a national
40 or local level in the coming weeks can be of great benefit for taking de-
41 cisions on the introduction of public health measures. One approach has
42 been to systematically test for infection with SARS-CoV-2 all patients hos-
43 pitalized for elective procedures, which outperformed state-based data in
44 predicting the local clinical burden (Covello et al., 2021). Furthermore, more
45 detailed hospital data like ICU admission and discharge or ambulance service
46 and emergency unit notes have been used for predicting COVID-19-related
47 hospital admissions within a region (Qian et al., 2021; Ferté et al., 2022).
48 Augmenting data extracted from electronic health records with exogenous
49 variables like weather or mobility data also lead to more accurate forecasts
50 of local COVID-19 related hospital admissions compared to using hospital
51 data alone Ferté et al. (2022); Zhang et al. (2022); Klein et al. (2023).

52 In this study, we compared the performance of different machine learning
53 models to forecast the number of COVID-19 hospital admissions based on
54 routinely collected electronic health records (EHR) and wastewater data. We
55 hypothesized that quantities such as the occupancy of a hospital’s emergency
56 ward, vital signs of hospital patients such as fever, or measurements of viral
57 load in wastewater have high predictive power for short-term forecasting of
58 COVID-19-related hospital admissions and lead to more accurate predictions
59 than relying on the number of hospital admissions in the previous days alone.
60 First, we extracted candidate variables that could have high predictive power
61 for the spread of SARS-CoV-2 from a large individual patient-level EHR
62 dataset from six hospitals in the Bern region, Switzerland, in the period from
63 February 2020 to June 2023. Second, we trained different machine learning
64 models with different combinations of features on the data to forecast the
65 number of COVID-19 hospital admissions up to five weeks in advance. Third,
66 we evaluated the performance of the models in comparison to a baseline
67 model across different forecasting setups.

68 **2. Data and Methods**

69 *2.1. Forecasting setup*

70 This study aimed to validate and compare methods for forecasting the
71 weekly number of COVID-19 hospital admissions up to five weeks in ad-
72 vance, using routinely collected hospital data from the previous days. As a

73 case study, we drew on electronic health records (EHR) data from six hospi-
74 tals belonging to the Insel Gruppe network, all located in the region of Bern,
75 Switzerland, collected from 25 February 2020, the day the first COVID-19
76 case was detected in Switzerland (FOPH, Federal Office of Public Health,
77 2020), and 30 June 2023 (full study period). We adopted a retrospective
78 approach by applying five forecasting models to historical time series data –
79 where outcomes are already known – enabling a comparison of the perfor-
80 mance of each method. We employed a leave-future-out strategy incorporat-
81 ing 12 separate test datasets, each covering a period of two to four months.
82 We selected the 12 cut-off points based on peaks and valleys of the daily time
83 series of COVID-19 hospital admissions in the next seven days (Supplemen-
84 tary Figure S1 A and B of Appendix A). The training datasets contained all
85 data collected before the respective cut-off point (Supplementary Figure S1 C
86 and D of Appendix A). The target week was defined as the sliding seven-day
87 window for which hospital admissions were forecast with the trained models
88 (Figure 1). We systematically varied the forecasting horizon k (i.e., the gap
89 between the last day of observed data and the start of the target week) and
90 the lookback window p (i.e., the number of past days of data included in the
91 model). More formally, placing ourselves at time t , we used data from days
92 $\{t-p, t-p+1, \dots, t-1\}$ to forecast the number of COVID-19 hospital admis-
93 sions during the target week $\{t+k, t+k+1, \dots, t+k+6\}$. In our analysis,
94 we used the following sets $k = \{0, 7, 14, 21, 28\}$ and $p = \{7, 14, 21, 28, 35\}$.
95 We did not consider forecasting horizons beyond five weeks as transmission
96 dynamics – and the many factors that influence them – are likely to shift
97 rapidly within that period (Holmdahl and Buckee, 2020).

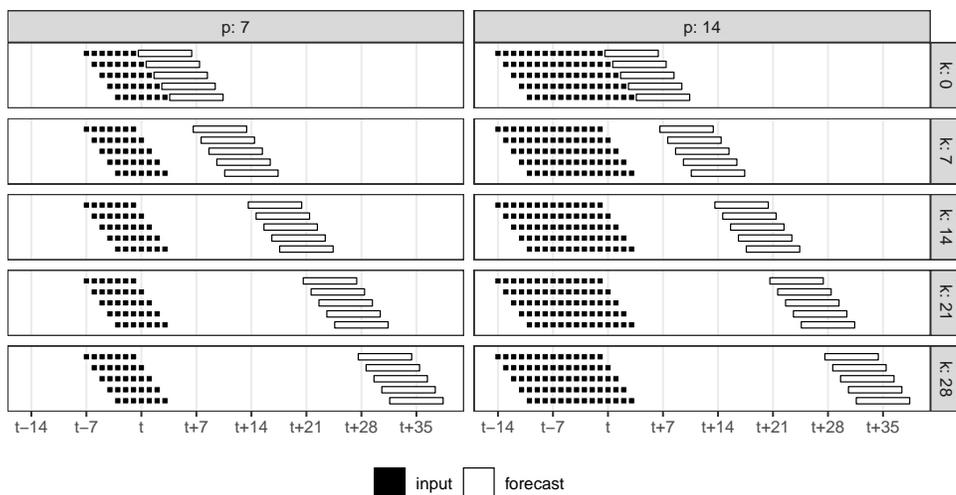


Figure 1: **Forecasting setups.** The forecasting horizon k corresponds to the gap between the last day of observed data and the start of the target week. The lookback window p is the number of past days of data included in the model. Day t corresponds to the start of a testing period.

98 *2.2. Electronic health records data*

99 We obtained individual-level electronic health records (EHR) from the In-
 100 sel Gruppe hospital network (inselgruppe.ch) in the canton of Bern, Switzer-
 101 land. During the study period, this hospital network comprised Bern Univer-
 102 sity Hospital, which is one of the five first-level university general hospitals of
 103 Switzerland, as well as five other hospitals (Aarberg, Belp, Münsingen, Rig-
 104 gisberg and Tiefenau) that are second-level general hospitals (FOPH, Federal
 105 Office of Public Health, 2023). In 2023, about 57,000 inpatients and 900,000
 106 outpatients were treated at Insel Gruppe hospitals (Inselgruppe, 2023). The
 107 full dataset covers the period from 1 January 2014 to 30 June 2023. It con-
 108 tains personal information about patients (e.g., age and sex), details of their
 109 hospital stay (e.g., dates of admission and discharge, hospital ward), as well
 110 as various clinical and laboratory measurements (e.g., body temperature,
 111 blood pressure, C-reactive protein [CRP] concentration). In addition, diag-
 112 noses of inpatients were recorded using ICD10 codes (WHO, 2019). These
 113 codes were assigned after discharge by trained medical coders based on the
 114 clinical documentation – including medical doctors’ notes, laboratory results,
 115 and imaging reports. This process is primarily done for administrative and

116 financial purposes, but can be leveraged for epidemiological monitoring (De-
117 mont et al.).

118 *2.3. Wastewater data*

119 In addition to EHR, we included measurements of the concentration of
120 SARS-CoV-2 RNA in wastewater. We used wastewater samples collected
121 daily at the Sensetal Laupen treatment plant between 16 November 2021
122 and 30 June 2023 (partial study period) as part of a wastewater surveillance
123 program coordinated by the Swiss Federal Institute of Aquatic Science and
124 Technology (Eawag) and the Swiss Federal Office of Public Health (FOPH)
125 (Eawag, 2021). As process control identified a possible underestimation by
126 approximately 30% of the SARS-CoV-2 viral load in wastewater during sum-
127 mer 2022, there is a five-week interruption in the data between 13 July and
128 16 August 2022. This plant covers approximately 62,000 people living in an
129 area west of the city of Bern, overlapping with the of the catchment area
130 of the Insel Gruppe hospital network. Samples were stored on-site at 4°C
131 and transported in batches to a laboratory for concentration, nucleic acid ex-
132 traction, and quantification using qPCR. Further details on the wastewater
133 sample laboratory procedures are available elsewhere (Huisman et al.).

134 *2.4. Data processing*

135 These raw data were processed to create 25 daily time series (Figure 2)
136 in three steps. First, we identified COVID-19-related hospital admissions
137 using the ICD10 code U07.1 (“*COVID-19, virus identified*” WHO (2019))
138 and created daily time series. We then smoothed this time series using a
139 seven-day moving sum to reduce day to day fluctuations and focus on the
140 actual trend of the time series. The entry at day t of the smoothed daily
141 time series corresponds to the total number of COVID-19-related hospital
142 admissions during days t to $t + 6$. This time series was used as the target
143 variable in all models. Furthermore, we stratified COVID-19 hospital admis-
144 sions into five age groups: Ages ≤ 4 , $5 - 14$, $15 - 29$, $30 - 64$ and ≥ 65
145 years. Second, we created several other daily time series to be used as fea-
146 tures in the models. These included the number of patients seeking care at
147 the emergency department of Bern University Hospital, from both patients
148 that were admitted to another hospital ward afterwards as well as patients
149 discharged directly. Next, we identified the daily number of hospital admis-
150 sions including a diagnosis belonging to one of five ICD10 chapters (R, I, E,
151 J or Z), belonging to one of five ICD10 categories (E87, J12, J96, I10, N18)

152 or including one of five specific ICD10 codes (J12.8, I10.90, J96.00, Z22.8, or
153 B33.8) (details about each code are available in Table 1). These chapters,
154 categories and codes were selected on the basis of the frequency with which
155 they appear together with the ICD10 code U07.1 in patients' diagnoses. We
156 also determined the daily number of inpatients admitted to hospital with
157 fever ($\geq 38.5^\circ\text{C}$) and the daily number of inpatients admitted to hospital
158 with a high CRP concentration ($\geq 50\text{ mg/l}$). Third, we processed SARS-
159 CoV-2 wastewater concentration data by 1) normalizing measurement using
160 the flow of wastewater on the sampling day as in common practice (Huisman
161 et al.), and 2) filling missing values using linear interpolation. Note that
162 wastewater data were only available for a shorter time period, referred to as
163 the partial study period in the following. From these 25 times series, we built
164 9 feature sets for the full study period referred to by letters A to I (without
165 wastewater) and 3 additional feature sets for only the partial study period
166 referred to by letters J to L (with wastewater) (Tables 2 and 3).

167 *2.5. Models*

168 We applied several supervised machine learning models, each based on
169 a different algorithm, to forecast COVID-19 hospital admissions. We se-
170 lected last observation carried forward (LOCF) to serve as the baseline for
171 performance comparison. Four models were evaluated : (1) a simple lin-
172 ear regression (LR) model (using base R `lm` function), (2) a recurrent neu-
173 ral network (RNN) (using Python library Keras (Chollet et al., 2015)), 3)
174 a long short-term memory (LSTM) neural network model (Hochreiter and
175 Schmidhuber, 1997) (using Python library Keras (Chollet et al., 2015)), and
176 4) a gradient boosting model (XGBoost) (XGBoost community, 2025) (us-
177 ing R package `xgboost` (Chen et al., 2024)). For both RNN and LSTM, we
178 used a grid-search strategy to optimize the architecture of the network, the
179 activation function and several hyperparameters (144 combinations each).
180 For XGBoost, we evaluated 864 combinations of hyperparameters, including
181 maximal tree depth. In all cases, the optimization of hyperparameters was
182 based on the root mean square error (RMSE) between forecasts and obser-
183 vations in the test set. A complete list of all hyperparameters for all models
184 is included in Supplementary Table S1 of Appendix A. Model forecasts of
185 COVID-19 hospital admissions were directly taken as forecasts in the case
186 of XGBoost, while for the RNN and LSTM models forecasts were averaged
187 over 50 independent runs (i.e., the forecasts correspond to an ensemble mean

188 taken sample-wise across 50 independent runs of the same model with differ-
189 ent random seeds).

190 *2.6. Evaluation of model performance*

191 We used a summary score based on RMSE to evaluate the predictive per-
192 formance of the different models to forecast COVID-19 hospital admissions in
193 comparison to the baseline model LOCF across a range of experimental con-
194 ditions. For a collection of forecasts, we first determined for each the RMSE
195 between forecast and observed values in the respective test set. Second, we
196 divided the obtained number by the RMSE resulting from the forecast of the
197 baseline model LOCF in the same conditions. Finally, we aggregated these
198 ratios into a single number by computing their geometric mean. We included
199 a more formal definition of the summary score in Chapter 1.5 of Appendix A.
200 This metric was computed for every combination of model and feature set,
201 separately for the full (without wastewater) and the partial study period
202 (with wastewater). As an additional metric for these collections of forecasts,
203 we computed the percentage of forecasts that achieved a lower RMSE than
204 the forecast of the baseline model LOCF with the same forecasting horizon k .
205 We also computed the summary score within additional levels of stratification
206 (e.g., for each combination of k and p) to identify which models performed
207 best across different conditions. The summary score provided a clear and in-
208 terpretable measure of performance: values below 1 indicate that on average
209 the model forecasts considered are more accurate than the forecasts of the
210 baseline model LOCF, while values above 1 suggest inferior performance.

211 *2.7. Data and code availability*

212 All code written in R and Python as well as some data and results files
213 are publicly available in the GitHub repository (github.com/mwohlfender/hospital_admission_forecasting). Due to data protection regulations we can
214 not make the full hospital dataset publicly available, but only in aggregated
215 form.
216

217 **3. Results**

218 Between 25 February 2020 and 30 June 2023, we identified 6,038 COVID-
219 19-related inpatient admissions, i.e. hospital stays of at least one night with
220 ICD10 code U07.1, in 6 hospitals in the canton of Bern, Switzerland (Table 1
221 and Figure 2). 389 patients (6.4%) were 0 – 4, 108 patients (1.8%) 5 –

222 14, 220 patients (3.6%) 15 – 29, 1840 patients (30.5%) 30 – 64 and 3481
223 patients (57.7%) were at least 65 years old. 527 of 717 COVID-19-related
224 inpatient admissions of patients below the age of 30 occurred between 1
225 January 2022 and 30 June 2023. The peaks in the number of visits at the
226 emergency ward of Bern University Hospital in mid-march 2020 and in late
227 October 2020 reflect rapid increases in COVID-19 cases in Switzerland. The
228 Omicron variant did not lead to a distinct increase of the number of patients
229 seeking care at the emergency ward of Bern in December 2022 or January
230 2023. In autumn 2020, the trend in COVID-19-related hospital admissions
231 coincided with that of hospital admissions with ICD10 codes J12.8 (“*Other*
232 *viral pneumonia*”) and B33.8 (“*Other specified viral diseases*”). A similar
233 pattern was observed in the first half of 2022 with ICD10 code Z22.8 (“*Carrier*
234 *of other infectious diseases*”). The wastewater data showed generally similar
235 trends as the COVID-19-related hospital admissions time series.

Table 1: **Summary characteristics of model variables.** Abbreviation, definition and use in feature sets of all input variables extracted from electronic health records (EHR) and wastewater data. Sum, mean, minimum and maximum are taken across all days of the full study period for the EHR data and across all days of the partial study period for the wastewater data. The unit of the variables extracted from EHR data is the number of new hospital admissions fulfilling a certain criterion per day. The unit of the viral load in wastewater samples is the number of SARS-CoV-2 RNA copies per 100,000 people in the collection area and day.

Variable	Details	Feature sets	Sum	Daily Mean	Daily Min	Daily Max
Admissions	COVID-19-related hospital admissions (any age)	A-I and K-L	6,038	4.9	0	31
Age 0-4	COVID-19-related hospital admissions (age 0-4)	B	389	0.3	0	6
Age 5-14	COVID-19-related hospital admissions (age 5-14)	B	108	0.1	0	4
Age 15-29	COVID-19-related hospital admissions (age 15-29)	B	220	0.2	0	4
Age 30-64	COVID-19-related hospital admissions (age 30-64)	B	1,840	1.5	0	13
Age 65+	COVID-19-related hospital admissions (age 65+)	B	3,481	2.8	0	18
Emergency	Patients seeking treatment at Bern University Hospital Emergency Department	C, I and L	200,895	164.4	74	348
ICD10 R	Symptoms, signs and abnormal clinical and laboratory findings, not elsewhere classified	D	68,246	55.8	18	99
ICD10 I	Diseases of the circulatory system	D	97,842	80.1	22	144
ICD10 E	Endocrine, nutritional and metabolic diseases	D	86,092	70.5	17	132
ICD10 J	Diseases of the respiratory system	D	39,121	32.0	7	65
ICD10 Z	Factors influencing health status and contact with health services	D	102,570	83.9	20	165
ICD10 E87	Other disorders of fluid, electrolyte and acid-base balance	E	19,577	16.0	1	39
ICD10 J12	Viral pneumonia, not elsewhere classified	E	2,745	2.2	0	21
ICD10 J96	Respiratory failure, not elsewhere classified	E	9,877	8.1	0	27
ICD10 I10	Essential (primary) hypertension	E	44,497	36.4	7	72
ICD10 N18	Chronic kidney disease	E	27,847	22.8	2	49
ICD10 J12.8	Other viral pneumonia	F, I and L	2,426	2.0	0	21
ICD10 I10.90	Essential hypertension, unspecified without indication of hypertensive crisis	F, I and L	40,909	33.5	6	69
ICD10 J96.00	Acute respiratory failure, not elsewhere classified	F, I and L	5,361	4.4	0	17
ICD10 Z22.8	Carrier of other infectious diseases	F, I and L	1,344	1.1	0	17
ICD10 B33.8	Other specified viral diseases	F, I and L	808	0.7	0	11
Fever	Highest body temperature measurement on day of admission at least 38.5 degrees Celsius	G, I and L	16,589	13.6	1	56
CRP	Highest CRP concentration measurement on day of admission at least 50 mg/l	H, I and L	28,726	23.5	3	50
Wastewater	SARS-CoV-2 RNA copies per 100,000 people living in the collection area of the Sensetal Laupen wastewater treatment plant and day	J, K and L	4.0×10^{15}	6.8×10^{12}	2.3×10^{10}	3.7×10^{13}

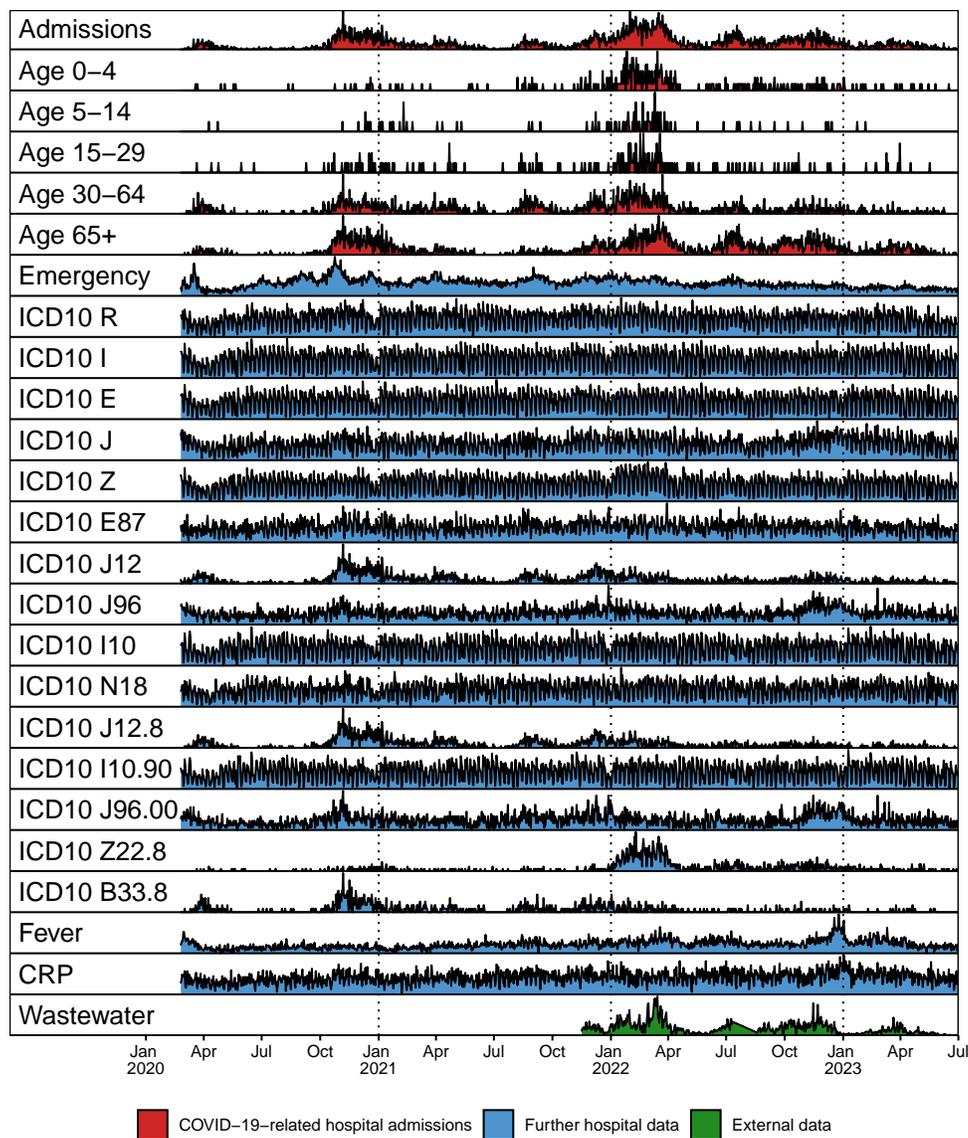


Figure 2: **Temporal profile of model variables.** Variables are extracted from electronic health records and wastewater data. All time series are normalized and on a daily level.

236 We combined the outputs of models with varying forecasting horizons
237 to generate forecasts of the weekly number of COVID-19-related hospital
238 admissions up to five weeks ahead. Examples of such forecasts are presented
239 in Figure 3. We found that no single combination of model and feature set
240 consistently produced the most accurate forecasts. The precision of forecasts,
241 as measured by the RMSE between forecasts and observations in the test
242 set, varied substantially across time periods, models, feature sets, lookback
243 windows and forecasting horizons. Overall, forecast precision improved as
244 more data became available for model training. The largest discrepancies
245 between forecasts and observed values were observed during periods with
246 rapid increases of COVID-19-related hospital admissions.

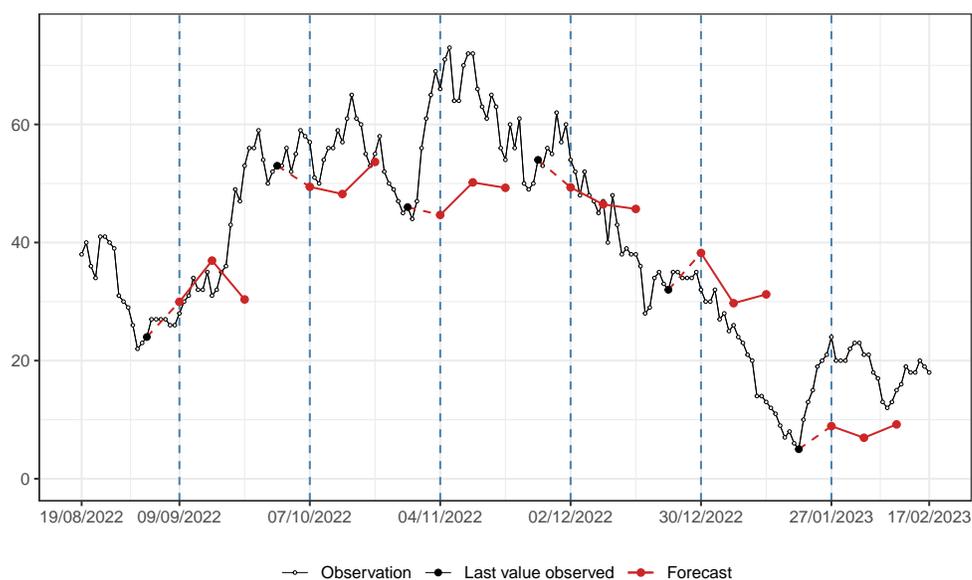


Figure 3: **Examples of forecasts of weekly COVID-19-related hospital admissions to three weeks ahead during autumn 2022.** Empty dots correspond to the number of COVID-19-related hospital admissions in the next seven days. For six dates set at regular intervals of four weeks, forecasts are generated using XGBoost for each of the following three weeks, based on COVID-19-related hospital admissions of the last 28 days.

247 Overall, all models except LR outperformed baseline when trained exclu-
248 sively on counts of past COVID-19-related hospital admissions (feature set
249 A), as assessed by both the summary score and the proportion of forecasts
250 with lower RMSE than baseline (Table 2). The reductions in summary score

251 were relatively modest (0.91 for RNN, 0.90 for LSTM and 0.88 for XGBoost),
252 with XGBoost achieving the lower summary score and the most consistent
253 performance across all combinations of lookback window, forecasting hori-
254 zon, and train-test split (outperforming baseline in 76% of cases). With the
255 exception of XGBoost, which maintained stable performance, all models per-
256 formed worse on average on the partial study period compared to the full
257 study period (Table 3). This decline in performance was even more pro-
258 nounced for LR and LSTM model.

259 Including additional features beyond past COVID-19-related hospital ad-
260 missions did not lead to any substantial improvement in the average summary
261 score for any model on the full study period. For all feature sets except the
262 number of COVID-19-related hospital admissions combined with the num-
263 ber of patients seeking care at the emergency ward (feature set C), XGBoost
264 consistently outperformed both the baseline and all other models on the full
265 study period (Table 2). RNN and LSTM performed similarly to baseline, only
266 showing noticeable improvement when trained on the number of COVID-19-
267 related hospital admissions alone (feature set A) or in combination with the
268 number of patients seeking care at the emergency ward (feature set C). There
269 was no feature set that enabled LR to produce more accurate forecasts than
270 the baseline. The performance of LR was particularly poor when multiple
271 features were added.

272 Adding measurements of SARS-CoV-2 viral load in wastewater to the
273 feature set led to noticeable improvement of the performance of XGBoost on
274 the partial study period (summary score 0.73 and improvement over baseline
275 in 97% of the cases, Table 3). Using other feature sets, XGBoost led to
276 slightly more precise forecasts on the partial study compared to the full
277 study period. On the contrary, the other models generally performed worse
278 on the partial study period than on the full study period. This drop in
279 accuracy could be substantial, for instance for forecasts generated with LSTM
280 using past COVID-19-related hospital admissions combined with counts of
281 inpatients with high CRP value (feature set H, average summary score of
282 1.51 compared to 0.96).

Table 2: **Summary of model performance for the full study period (25 February 2020 to 30 June 2023)**. For each model and feature set, the summary score was computed as the geometric mean of the ratios of the root mean square error (RMSE) over the baseline RMSE across all combinations of forecasting horizon k , lookback window p , and train-test split. The proportion of forecasts where the RMSE is lower than the baseline is shown in parentheses.

Feature set	LOCF	LR	RNN	LSTM	XGBoost
A (COVID-19-related hospital admissions)	1	1.08 (48 %)	0.91 (69 %)	0.90 (68 %)	0.88 (76 %)
B (A + COVID-19-related hospital admissions by age group)	1	1.69 (19 %)	0.96 (58 %)	1.01 (55 %)	0.87 (75 %)
C (A + emergency)	1	1.17 (51 %)	0.89 (70 %)	0.89 (58 %)	0.90 (69 %)
D (A + ICD10 chapters)	1	1.51 (28 %)	0.98 (57 %)	0.93 (56 %)	0.86 (77 %)
E (A + ICD10 categories)	1	1.68 (16 %)	1.00 (55 %)	0.95 (58 %)	0.87 (75 %)
F (A + ICD10 codes)	1	1.82 (17 %)	1.02 (46 %)	1.02 (48 %)	0.89 (75 %)
G (A + fever)	1	1.16 (39 %)	0.98 (56 %)	0.91 (61 %)	0.87 (75 %)
H (A + CRP)	1	1.28 (37 %)	0.98 (53 %)	0.96 (54 %)	0.87 (76 %)
I (A + emergency + ICD10 codes + fever + CRP)	1	2.02 (12 %)	1.12 (37 %)	1.01 (46 %)	0.91 (71 %)

Table 3: **Summary of model performance for the partial study period (16 November 2021 to 30 June 2023).** For each model and feature set, the summary score was computed as the geometric mean of the ratios of the root mean square error (RMSE) over the baseline RMSE across all combinations of forecasting horizon k , look-back window p , and train-test split. The proportion of forecasts where the RMSE is lower than the baseline is shown in parentheses.

Feature set	LOCF	LR	RNN	LSTM	XGBoost
A (COVID-19-related hospital admissions)	1	1.45 (14 %)	1.16 (43 %)	1.35 (32 %)	0.81 (78 %)
B (A + COVID-19-related hospital admissions by age group)	1	2.26 (2 %)	1.05 (53 %)	1.39 (32 %)	0.80 (79 %)
C (A + emergency)	1	1.02 (45 %)	0.83 (75 %)	1.01 (58 %)	0.89 (69 %)
D (A + ICD10 chapters)	1	2.24 (0 %)	1.46 (24 %)	1.33 (34 %)	0.78 (81 %)
E (A + ICD10 categories)	1	3.06 (3 %)	1.52 (17 %)	1.50 (23 %)	0.82 (67 %)
F (A + ICD10 codes)	1	2.42 (0 %)	1.32 (29 %)	1.41 (26 %)	0.81 (75 %)
G (A + fever)	1	1.77 (4 %)	1.19 (49 %)	1.20 (36 %)	0.75 (90 %)
H (A + CRP)	1	1.73 (8 %)	1.54 (15 %)	1.51 (23 %)	0.81 (77 %)
I (A + emergency + ICD10 codes + fever + CRP)	1	3.50 (0 %)	1.08 (60 %)	1.23 (44 %)	0.82 (74 %)
J (Wastewater)	1	2.04 (0 %)	1.33 (34 %)	1.50 (30 %)	0.74 (89 %)
K (A + wastewater)	1	1.51 (16 %)	1.13 (45 %)	1.38 (32 %)	0.73 (97 %)
L (A + emergency + ICD10 codes + fever + CRP + wastewater)	1	3.77 (2 %)	1.02 (61 %)	1.22 (44 %)	0.75 (86 %)

283 Forecasting performance was highly dependent on forecasting horizon k
 284 and lookback window p , with the ranking of models and feature sets varying
 285 across the values chosen for k and p . At least one model outperformed the
 286 baseline for all combinations of k and p , in both the full and the partial
 287 study periods (Figure 4A). XGBoost was the best-performing model for 12
 288 out of 25 combinations of k and p for the full study period (summary scores
 289 ranging from 0.67 to 0.84) and for 22 out of 25 combinations of k and p during
 290 the partial study period (summary scores ranging from 0.53 to 0.89). For
 291 smaller values of k and p , the RNN and LSTM models outperformed the other
 292 models, particularly for the full study period (summary scores ranging from
 293 0.72 to 0.95). As the forecasting horizon and lookback window increased,
 294 the XGBoost model more frequently achieved the best performance. This
 295 pattern was more pronounced for the partial study period than for the full
 296 study period.

297 The optimal feature set also varied according to experimental conditions.

298 For 20 out of 25 combinations of k and p during the full study period and
299 all 25 combinations of k and p during the partial study period, best per-
300 formance was obtained using a feature set that included additional features
301 besides COVID-19-related hospital admissions (Figure 4B). For the full study
302 period, and for longer horizons for the partial study period, past COVID-19-
303 related hospital admissions combined with counts of patients admitted with
304 fever (feature set G) most frequently achieved the best summary score (sum-
305 mary score ranging from 0.53 to 0.91). Using counts of patients admitted
306 with high CRP (feature set H) and counts of patients seeking care at the
307 emergency ward (feature set C) were also sometimes selected as achieving
308 best performance, especially for longer forecast horizons (three to five weeks
309 ahead). The inclusion of viral load measurements in wastewater samples led
310 to the best results at short forecast horizons (up to three weeks ahead) for
311 the partial study period (summary score ranging from 0.71 to 0.88).

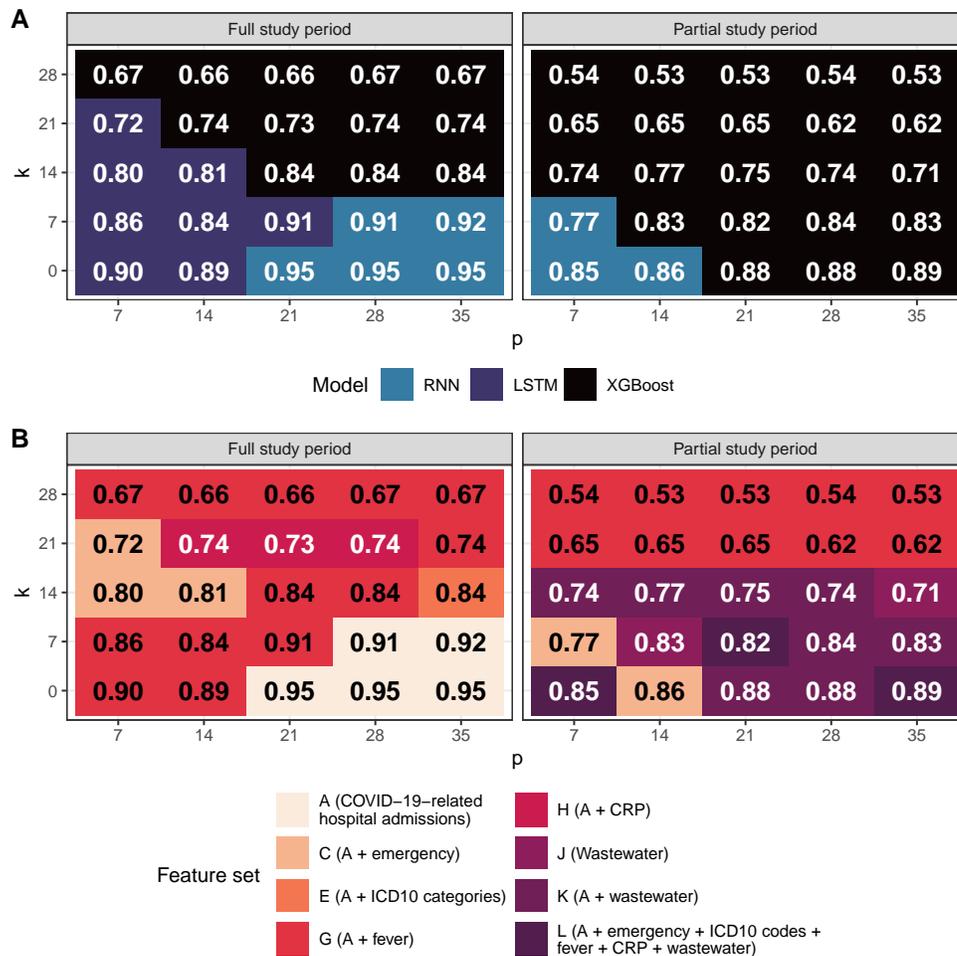


Figure 4: **Best-performing model and feature set for each combination of forecasting horizon k and lookback window p .** A: Model achieving the lowest root mean squared error (RMSE) during the full and the partial study period. B: Feature set achieving lowest RMSE during the full and the partial study period. Numbers indicate the average summary score, computed as the geometric mean of the ratios of RMSE over the baseline RMSE across all train-test splits.

312 4. Discussion

313 In this study, we systematically evaluated and compared the ability of var-
 314 ious machine learning algorithms to forecast the number of weekly COVID-
 315 19 hospital admissions up to five weeks ahead, using different combinations
 316 of variables extracted from EHR from six hospitals in the region of Bern,

317 Switzerland, as well measurements of SARS-CoV-2 viral load in wastewater
318 samples. Across all examined forecasting horizons, we were able to gener-
319 ate forecasts that consistently outperformed the baseline model of LOCF,
320 with greater improvements observed for longer forecasting horizons. Overall,
321 our findings confirm that EHR hold considerable potential for improving the
322 forecasting of infectious disease dynamics.

323 We found that gradient boosting using the XGBoost algorithm outper-
324 formed other models on average across all combinations of forecasting horizon
325 k and lookback window p . This is somewhat surprising as XGBoost was not
326 inherently built for time series forecasting. Still, XGBoost performed better
327 than linear regression and neural networks (RNN and LSTM), particularly
328 for longer forecasting horizons. These findings may be explained by the
329 discrete tree-based approach of XGBoost, leading to a good handling of non-
330 linearities in addition to the reduced risk of overfitting (Park and Ho, 2021).
331 Moreover, XGBoost may have an advantage because of the relative scarcity
332 of data. As we work with daily time series, our models never get a training
333 set containing more than about 1,200 data points, which was further reduced
334 when using longer lookback windows or focusing on the partial study period.
335 In these situations, forecasts generated with linear regression and neural net-
336 work were prone to perform considerably worse than LOCF, while XGBoost
337 remained mostly adequate. This feature makes XGBoost particularly ap-
338 pealing in the early stages of epidemics and for emerging infectious diseases
339 lacking historical data.

340 With regards to variables relevant for forecasting COVID-19-related hos-
341 pital admissions, our findings indicate that relying solely on past admission
342 counts is suboptimal. Complementing these data with additional variables
343 available in EHR such as the number of patients admitted with fever, with el-
344 evated CRP or presenting to emergency care improved forecast performance,
345 particularly at longer forecasting horizons (three to five weeks ahead). Be-
346 sides EHR, our results confirm the transformative potential of incorporat-
347 ing viral load measurements in wastewater in infectious disease forecast-
348 ing (Rankin et al.). Forecasts based on recent wastewater data demonstrated
349 substantially improved performance for shorter horizons (up to three weeks
350 ahead), while EHR-based variables such as fever-related admissions retained
351 a performance advantage at longer horizons (four and five weeks ahead). This
352 pattern likely reflects the temporal lag between infection incidence (captured
353 by wastewater surveillance via fecal shedding) and subsequent hospital ad-
354 missions, which has been estimated to range between 10 and 14 days (Hegazy

355 et al., 2022).

356 Other studies have used similar approaches for hospital admission fore-
357 casting, and found that combining hospital admissions data on the level of a
358 single hospital or aggregated on a regional level with additional health data
359 (e.g., occupancy of emergency units or use of ambulance services) or external
360 data (e.g., mobility or weather data) lead to more accurate forecasts (Ferté
361 et al., 2022; Zhang et al., 2022; Klein et al., 2023). Our results are aligned
362 with their findings, but a quantitative comparison of the precision of the
363 obtained forecasts between studies is difficult due to different available data,
364 study periods and evaluation metrics.

365 The main strength of our work lies in the breadth and thoroughness of
366 the systematic comparison of different models and combinations of features,
367 which leads to a proof-of-concept that routinely collected EHR can indeed
368 provide a solid data basis for an infectious disease forecasting system. This
369 represents a step forward in the development of infectious disease monitoring
370 and forecasting systems relying on data that has not been collected specifi-
371 cally for research purposes. Given that the data can be accessed with little
372 time delays, forecasting could be conducted continuously and provide reli-
373 able estimates of quantities of interest such as new COVID-19-related hos-
374 pital admissions without depending on time-consuming and expensive data
375 collection.

376 Our study comes with several limitations. First, the generalizability of
377 our findings beyond the Insel Gruppe hospital network in the region of Bern
378 remains uncertain. Differences in EHR structures, conventions and formats
379 could make it difficult to replicate our study in other settings. We refrained
380 from requesting access to EHR from other Swiss university hospitals. Sec-
381 ond, we carried out a purely retrospective analysis, and did not implement
382 our forecasting framework in a real-time operational context. Real-time de-
383 ployment would require additional development of data pipelines and in-
384 frastructure. One key obstacle, which we could not influence, was the time
385 lag between hospital admission and the encoding of diagnoses using ICD10
386 codes, which can occur several weeks after discharge. Reducing these delays
387 is essential for enabling the practical application of forecasting approaches
388 such as ours, although as we showed the best-performing features (fever,
389 CRP and emergency ward) do not rely on ICD-10 encoding and are available
390 immediately. Third, from a technical perspective, we did not use a distinct
391 validation set to tune model hyperparameters, instead doing this directly on
392 the testing set. This decision was made in light of limited data availability,

393 as reserving additional data for validation would have reduced the training
394 set. Similarly, we also did not estimate the uncertainty of the model fore-
395 casts. While techniques such as conformal prediction were considered, their
396 application would have required additional data splitting, further reducing
397 the training set. Finally, as with many forecasting approaches, we did not
398 account for changes in transmission dynamics, for example due to shifts in
399 population behavior, the emergence of new variants or increases in the immu-
400 nity level due to vaccination. Future work is needed to develop forecasting
401 methods that can incorporate a broader range of dynamic data sources.

402 The vast amount of routinely collected medical data remains underuti-
403 lized for infectious disease forecasting. Our findings demonstrate that such
404 data, when properly harnessed with modern machine learning approaches,
405 can substantially enhance the accuracy of short-term hospital admission fore-
406 casts. Such forecasts are especially valuable for informing public health pol-
407 icy, enabling healthcare systems to anticipate surges in demand and allocate
408 resources accordingly. As data infrastructures continue to expand, with more
409 and more hospital data becoming available in standardized format and with
410 decreasing delays, the integration of routine clinical and surveillance data
411 into real-time forecasting systems will become more feasible. This paves the
412 way for highly-efficient forecasting tools that can support timely and data-
413 driven responses to emerging infectious disease threats, strengthening overall
414 pandemic preparedness.

415 **5. Acknowledgments**

416 This study is funded by the Multidisciplinary Center for Infectious Dis-
417 eases, University of Bern, Bern, Switzerland. We gratefully acknowledge the
418 Insel Data Science Center (IDSC) (www.idsc.io/en/) for facilitating the ac-
419 cess to the electronic health records from the Insel Gruppe hospital network.
420 Calculations were performed on UBELIX (www.id.unibe.ch/hpc), the HPC
421 cluster at the University of Bern.

422 **Appendix A. Supplementary material**

423 Detailed model description, table of hyperparameters, additional results
424 and figures supporting the main text.

425 References

- 426 Anderegg, N., Panczak, R., Egger, M., Low, N., Riou, J., 2022. Survival
427 among people hospitalized with COVID-19 in Switzerland: a nationwide
428 population-based analysis. *BMC Med* 20, 164. doi:10.1186/s12916-022-
429 02364-7.
- 430 CDC, 2024. SPHERES. URL: [https://www.cdc.gov/advanced-](https://www.cdc.gov/advanced-molecular-detection/php/spheres/index.html)
431 [molecular-detection/php/spheres/index.html](https://www.cdc.gov/advanced-molecular-detection/php/spheres/index.html).
- 432 Chen, T., He, T., Benesty, M., Khotilovich, V., Tang, Y., Cho, H., Chen,
433 K., Mitchell, R., Cano, I., Zhou, T., Li, M., Xie, J., Lin, M., Geng, Y., Li,
434 Y., Yuan, J., 2024. xgboost: Extreme Gradient Boosting. URL: [https:](https://CRAN.R-project.org/package=xgboost)
435 [//CRAN.R-project.org/package=xgboost](https://CRAN.R-project.org/package=xgboost). r package version 1.7.7.1.
- 436 Chollet, F., et al., 2015. Keras. <https://keras.io>.
- 437 Covello, L., Gelman, A., Si, Y., Wang, S., 2021. Routine Hospital-based
438 SARS-CoV-2 Testing Outperforms State-based Data in Predicting Clinical
439 Burden. *Epidemiology* 32, 792–799. doi:10.1097/EDE.0000000000001396.
- 440 Cramer, E.Y., Huang, Y., Wang, Y., Ray, E.L., Cornell, M., Bracher,
441 J., Brennen, A., Castro Rivadeneira, A.J., Gerding, A., House, K.,
442 Jayawardena, D., Kanji, A.H., Khandelwal, A., Le, K., Niemi, J., Stark,
443 A., Shah, A., Wattanachit, N., Zorn, M.W., Reich, N.G., Consortium,
444 U.C.F.H., 2021. The United States COVID-19 Forecast Hub dataset.
445 medRxiv URL: [https://www.medrxiv.org/content/10.1101/2021.11.](https://www.medrxiv.org/content/10.1101/2021.11.04.21265886v1)
446 [04.21265886v1](https://www.medrxiv.org/content/10.1101/2021.11.04.21265886v1), doi:10.1101/2021.11.04.21265886.
- 447 Demont, C., Petrica, N., Bardoulat, I., Duret, S., Watier, L., Chosidow, A.,
448 Lorrot, M., Kieffer, A., Lemaitre, M., . Economic and disease burden
449 of RSV-associated hospitalizations in young children in France, from 2010
450 through 2018. *BMC Infect Dis* 21, 730. URL: [https://doi.org/10.1186/](https://doi.org/10.1186/s12879-021-06399-8)
451 [s12879-021-06399-8](https://doi.org/10.1186/s12879-021-06399-8), doi:10.1186/s12879-021-06399-8.
- 452 Didriksson, I., Leffler, M., Spångfors, M., Lindberg, S., Reepalu, A., Nilsson,
453 A., Cronqvist, J., Andertun, S., Nelderup, M., Jungner, M., Johnsson,
454 P., Lilja, G., Frigyesi, A., Friberg, H., 2022. Intensive care unit burden
455 is associated with increased mortality in critically ill COVID-19 patients.
456 *Acta Anaesthesiol Scand* , 10.1111/aas.14184 URL: [https://www.ncbi.](https://www.ncbi.nlm.nih.gov/pmc/articles/PMC9878196/)
457 [nlm.nih.gov/pmc/articles/PMC9878196/](https://www.ncbi.nlm.nih.gov/pmc/articles/PMC9878196/), doi:10.1111/aas.14184.

- 458 Eawag, 2021. SARS-CoV-2 in Wastewater. URL: [https://www.eawag.ch/
459 en/department/sww/projects/sars-cov2-in-wastewater/](https://www.eawag.ch/en/department/sww/projects/sars-cov2-in-wastewater/).
- 460 ECDC, 2021. European COVID-19 Forecast Hub. URL: [https://
461 covid19forecasthub.eu/index.html](https://covid19forecasthub.eu/index.html).
- 462 ECDC, 2023. RespiCast ECDC Respiratory Diseases Forecasting Hub. URL:
463 <https://respicast.ecdc.europa.eu/>.
- 464 Ferté, T., Jouhet, V., Griffier, R., Hejblum, B.P., Thiébaud, R., Bordeaux
465 University Hospital Covid-19 Crisis Task Force, 2022. The benefit of
466 augmenting open data with clinical data-warehouse EHR for forecasting
467 SARS-CoV-2 hospitalizations in Bordeaux area, France. JAMIA Open 5,
468 ooac086. doi:10.1093/jamiaopen/ooac086.
- 469 FOPH, Federal Office of Public Health, 2020. New Coronavirus 2019-nCoV:
470 first confirmed case in Switzerland. URL: [https://www.bag.admin.
471 ch/bag/en/home/das-bag/aktuell/medienmitteilungen.msg-id-
472 78233.html](https://www.bag.admin.ch/bag/en/home/das-bag/aktuell/medienmitteilungen.msg-id-78233.html).
- 473 FOPH, Federal Office of Public Health, 2023. Kennzahlen der Schweizer
474 Spitäler 2023. URL: [https://www.bag.admin.ch/bag/de/home/zahlen-
475 und-statistiken/zahlen-fakten-zu-spitaelern/kennzahlen-der-
476 schweizer-spitaeler.html](https://www.bag.admin.ch/bag/de/home/zahlen-und-statistiken/zahlen-fakten-zu-spitaelern/kennzahlen-der-schweizer-spitaeler.html).
- 477 Furuse, Y., 2021. Genomic sequencing effort for SARS-CoV-2 by country
478 during the pandemic. International Journal of Infectious Diseases 103, 305–
479 307. URL: [https://www.sciencedirect.com/science/article/pii/
480 S1201971220325571](https://www.sciencedirect.com/science/article/pii/S1201971220325571), doi:10.1016/j.ijid.2020.12.034.
- 481 Grantz, K.H., Meredith, H.R., Cummings, D.A.T., Metcalf, C.J.E., Grenfell,
482 B.T., Giles, J.R., Mehta, S., Solomon, S., Labrique, A., Kishore, N., Buc-
483 kee, C.O., Wesolowski, A., 2020. The use of mobile phone data to inform
484 analysis of COVID-19 pandemic epidemiology. Nature Communications 11,
485 4961. URL: <https://www.nature.com/articles/s41467-020-18190-5>,
486 doi:10.1038/s41467-020-18190-5. publisher: Nature Publishing Group.
- 487 Hegazy, N., Cowan, A., D'Aoust, P.M., Mercier, É., Towhid, S.T., Jia, J.J.,
488 Wan, S., Zhang, Z., Kabir, M.P., Fang, W., Graber, T.E., MacKen-
489 zie, A.E., Guilherme, S., Delatolla, R., 2022. Understanding the dy-
490 namic relation between wastewater SARS-CoV-2 signal and clinical metrics

- 491 throughout the pandemic. *Sci Total Environ* 853, 158458. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC9444583/>, doi:10.1016/j.scitotenv.2022.158458.
- 492
493
- 494 Hochreiter, S., Schmidhuber, J., 1997. Long Short-Term Memory. *Neural Comput.* 9, 1735–1780. URL: <https://doi.org/10.1162/neco.1997.9.8.1735>, doi:10.1162/neco.1997.9.8.1735.
- 495
496
- 497 Hodcroft, E.B., Wohlfender, M.S., Neher, R.A., Riou, J., Althaus, C.L.,
498 2025. Estimating Re and overdispersion in secondary cases from
499 the size of identical sequence clusters of SARS-CoV-2. *PLOS Com-*
500 *putational Biology* 21, e1012960. URL: <https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1012960>,
501 doi:10.1371/journal.pcbi.1012960. publisher: Public Library of
502 Science.
503
- 504 Holmdahl, I., Buckee, C., 2020. Wrong but Useful — What Covid-19 Epi-
505 *demiologic Models Can and Cannot Tell Us.* *New England Journal of*
506 *Medicine* 383, 303–305.
- 507 Huisman, J.S., Scire, J., Caduff, L., Fernandez-Cassi, X., Ganesanandamoorthy,
508 P., Kull, A., Scheidegger, A., Stachler, E., Boehm, A.B., Hughes,
509 B., Knudson, A., Topol, A., Wigginton, K.R., Wolfe, M.K., Kohn, T.,
510 Ort, C., Stadler, T., Julian, T.R., . Wastewater-Based Estimation of the
511 Effective Reproductive Number of SARS-CoV-2. *Environmental Health*
512 *Perspectives* 130, 057011. URL: <https://ehp.niehs.nih.gov/doi/10.1289/EHP10050>, doi:10.1289/EHP10050. publisher: Environmental Health
513 Perspectives.
514
- 515 Inselgruppe, 2023. Jahresberichte - Insel Gruppe AG. URL: <https://inselgruppe.ch/de/die-insel-gruppe/organisation/direktion-insel-gruppe/stabsbereiche/kommunikation/publikationen/jahresberichte>.
- 516
517
518
- 519 Jahn, K., Dreifuss, D., Topolsky, I., Kull, A., Ganesanandamoorthy, P.,
520 Fernandez-Cassi, X., Bänziger, C., Devaux, A.J., Stachler, E., Caduff,
521 L., Cariti, F., Corzón, A.T., Fuhrmann, L., Chen, C., Jablonski, K.P.,
522 Nadeau, S., Feldkamp, M., Beisel, C., Aquino, C., Stadler, T., Ort, C.,
523 Kohn, T., Julian, T.R., Beerenwinkel, N., 2022. Early detection and

524 surveillance of SARS-CoV-2 genomic variants in wastewater using CO-
525 JAC. *Nature Microbiology* 7, 1151–1160. URL: <https://www.nature.com/articles/s41564-022-01185-x>, doi:10.1038/s41564-022-01185-
526 x. number: 8 Publisher: Nature Publishing Group.

528 Klein, B., Zenteno, A.C., Joseph, D., Zahedi, M., Hu, M., Copen-
529 haver, M.S., Kraemer, M.U.G., Chinazzi, M., Klompas, M., Vespig-
530 nani, A., Scarpino, S.V., Salmasian, H., 2023. Forecasting hospital-level
531 COVID-19 admissions using real-time mobility data. *Commun Med* 3,
532 1–9. URL: <https://www.nature.com/articles/s43856-023-00253-5>,
533 doi:10.1038/s43856-023-00253-5. publisher: Nature Publishing Group.

534 Kraemer, M.U.G., Tsui, J.L.H., Chang, S.Y., Lytras, S., Khurana, M.P.,
535 Vanderslott, S., Bajaj, S., Scheidwasser, N., Curran-Sebastian, J.L., Se-
536 menova, E., Zhang, M., Unwin, H.J.T., Watson, O.J., Mills, C., Das-
537 gupta, A., Ferretti, L., Scarpino, S.V., Koua, E., Morgan, O., Tegally,
538 H., Paquet, U., Moutsianas, L., Fraser, C., Ferguson, N.M., Topol, E.J.,
539 Duchène, D.A., Stadler, T., Kingori, P., Parker, M.J., Dominici, F., Shad-
540 bolt, N., Suchard, M.A., Ratmann, O., Flaxman, S., Holmes, E.C., Gomez-
541 Rodriguez, M., Schölkopf, B., Donnelly, C.A., Pybus, O.G., Cauchemez,
542 S., Bhatt, S., 2025. Artificial intelligence for modelling infectious dis-
543 ease epidemics. *Nature* 638, 623–635. URL: <https://www.nature.com/articles/s41586-024-08564-w>, doi:10.1038/s41586-024-08564-
544 w. publisher: Nature Publishing Group.

546 Morvan, M., Jacomo, A.L., Souque, C., Wade, M.J., Hoffmann, T.,
547 Pouwels, K., Lilley, C., Singer, A.C., Porter, J., Evens, N.P., Walker,
548 D.I., Bunce, J.T., Engeli, A., Grimsley, J., O'Reilly, K.M., Danon, L.,
549 2022. An analysis of 45 large-scale wastewater sites in England to esti-
550 mate SARS-CoV-2 community prevalence. *Nature Communications* 13,
551 4313. URL: <https://www.nature.com/articles/s41467-022-31753-y>,
552 doi:10.1038/s41467-022-31753-y. number: 1 Publisher: Nature Pub-
553 lishing Group.

554 Park, Y., Ho, J.C., 2021. Tackling Overfitting in Boosting for Noisy Health-
555 care Data. *IEEE Transactions on Knowledge and Data Engineering* 33,
556 2995–3006. URL: [https://ieeexplore.ieee.org/abstract/document/](https://ieeexplore.ieee.org/abstract/document/8933485)
557 [8933485](https://ieeexplore.ieee.org/abstract/document/8933485), doi:10.1109/TKDE.2019.2959988.

- 558 Qian, Z., Alaa, A.M., van der Schaar, M., 2021. CPAS: the UK's national
559 machine learning-based hospital capacity planning system for COVID-19.
560 Machine Learning 110, 15–35. doi:10.1007/s10994-020-05921-4.
- 561 Rankin, N., Saiyed, S., Du, H., Gardner, L.M., . A multi-city COVID-19 fore-
562 casting model utilizing wastewater-based epidemiology. Science of The To-
563 tal Environment 960, 178172. URL: [https://www.sciencedirect.com/
564 science/article/pii/S004896972408330X](https://www.sciencedirect.com/science/article/pii/S004896972408330X), doi:10.1016/j.scitotenv.
565 2024.178172.
- 566 Shu, Y., McCauley, J., 2017. GISAID: Global initiative on
567 sharing all influenza data – from vision to reality. Euro-
568 surveillance 22, 30494. URL: [https://www.eurosurveillance.
569 org/content/10.2807/1560-7917.ES.2017.22.13.30494](https://www.eurosurveillance.org/content/10.2807/1560-7917.ES.2017.22.13.30494), doi:10.2807/
570 1560-7917.ES.2017.22.13.30494.
- 571 WHO, 2019. ICD-10 Version: 2019. URL: [https://icd.who.int/
572 browse10/2019/en/](https://icd.who.int/browse10/2019/en/).
- 573 XGBoost community, 2025. XGBoost: Scalable and Flexible Gradient Boost-
574 ing. URL: <https://xgboost.ai/>.
- 575 Zhang, J., Pathak, H.S., Snowdon, A., Greiner, R., 2022. Learning mod-
576 els for forecasting hospital resource utilization for COVID-19 patients
577 in Canada. Scientific Reports 12, 8751. URL: [https://www.nature.
578 com/articles/s41598-022-12491-z](https://www.nature.com/articles/s41598-022-12491-z), doi:10.1038/s41598-022-12491-
579 z. publisher: Nature Publishing Group.