

A SIMULATION MODEL AND DASHBOARD FOR PREDICTING COVID-19 BED REQUIREMENTS

Yin-Chi Chan
Kaya Dreesbeimdiek
Ajith Kumar Parlikad
Tom Ridgman

Institute for Manufacturing
University of Cambridge
17 Charles Babbage Rd
Cambridge, CB3 0FS, UK

Nicholas J. Matheson
Ben Warne

Cambridge Institute of Therapeutic
Immunology & Infectious Disease
University of Cambridge
Jeffrey Cheah Biomedical Centre
Cambridge Biomedical Campus, Puddicombe Way
Cambridge, CB2 0AW, UK

Denise Franks

Cambridge University Hospitals NHS Foundation Trust
Cambridge Biomedical Campus, Hills Road
Cambridge, CB2 0QQ, UK

ABSTRACT

The Covid-19 pandemic has placed extraordinary amounts of stress upon public hospitals globally. This paper describes a simulation model for estimating hospital bed demand based on generated scenarios. Statistical tools were also developed for generating these scenarios, in particular, for fitting distributions to patients' lengths-of-stay and for predicting the number of daily arrivals of Covid-19 patients. A web dashboard has been created for ease of use. The simulation model and statistical tools have been used to estimate Covid-related bed demand at an NHS hospital in the East of England.

1 INTRODUCTION

The coronavirus disease 2019 (Covid-19) pandemic has placed enormous stress on healthcare systems worldwide. For example, in the United Kingdom, the combination of Covid-19 and existing pressures on the National Health Service (NHS) has led to record waits for elective surgery, due to reductions in capacity across the entire British healthcare system as resources are rerouted towards Covid-19 care. Many public healthcare systems in the United Kingdom and abroad have also experienced high rates of staff burnout and resignations since the beginning during the pandemic (Murat et al. 2020; Sharifi et al. 2020; Nashwan et al. 2021), and a perception among healthcare workers that they have been let down due to a lack of preparation by those in charge. Sprung et al. (2020) bluntly state: "The poor outcomes observed in many disasters are often the result of inaction and poor implementation of the necessary measures required to prevent or mitigate the impact of disasters."

Although Covid-19 is no longer an emergent crisis in most parts of the world, the lessons gained during the pandemic can help bolster public healthcare systems against future epidemics and pandemics. For example, while conversion of regular wards (Sagy et al. 2021) or even non-clinical space (Locke et al. 2021) to isolation wards was a common practice during the Covid-19 pandemic, such measures require

time and therefore accurate forecasting of hospital demand in order to make such ward conversion decisions in advance. Furthermore, there are limits beyond which general ward or ICU expansion is not safe: based on experiences gained during SARS, Gomersall et al. (2006) suggested that the maximum safe increase in ICU capacity during a pandemic is 50-100%. Therefore, the ability to predict pandemic progression in advance is important to determine the ability of a healthcare system to meet the increased patient demands and what additional preventive measures (i.e., non-pharmaceutical interventions) are required to prevent healthcare system collapse during a pandemic. This motivates our development of a simulation model for estimating hospital bed demand caused by Covid-19 or a similar disease in the future.

1.1 Background and Related Work

There are many simulation models for estimating the progression of infectious diseases such as Covid-19 and the resultant load on healthcare systems (Currie et al. 2020). The most common of these are based on compartmental modelling, in which the number of individuals in each state is evaluated over time using a system of ordinary or stochastic differential equations (Brauer et al. 2019). However, these models are continuous-state, meaning they do not define a whole patient as the smallest possible unit of flow. Therefore, for operations management at the hospital level, discrete-event simulation (DES) modelling can be preferable instead. Due to each individual being defined in the simulation separately, DES makes it easier to model heterogeneous behaviour such as different hospital length-of-stay (LoS) distributions and ICU admission probabilities by age groups or other demographic factors. Additionally, the stochastic nature of such models makes it simple to observe the variability of a model through repeated simulations. Existing simulation models for Covid-19 include Weissman et al. (2020), Bovim et al. (2021), Melman et al. (2021), and Warde et al. (2021).

To model hospital demand caused by Covid-19, one must first model the volume of Covid-19 patient arrivals and their LoSs. Methods for modelling Covid-19 arrivals to a hospital include epidemiological models and direct curve-fitting. For example, Weissman et al. (2020) and Warde et al. (2021) modelled bed demand based on extended SIR epidemiological models, with parameters specifying the hospitalisation rate of Covid-19 patients and the percentage of these hospital admissions requiring ICU care. In contrast, examples using direct curve-fitting include that of Garcia-Vicuna et al. (2020), who used a Gompertz growth model of the form $F(t) = ae^{-b \exp(-ct)}$ to model the cumulative number of hospital admissions, and Palomo et al. (2020), who used a Gamma function to model daily admissions. This latter approach, where a curve is fitted directly to hospital admission or occupancy numbers, allows for a simple method of generating simulated arrivals to a hospital without the need for studying infection dynamics, as in epidemiological models.

Palomo et al. (2020) performed a queueing-theory analysis of a healthcare system with an arrival rate $\lambda(t)$ of Covid-19 patients following a gamma curve and showed that the timing and mean height of the peak bed demand could be computed analytically by solving a set of fixed-point equations. Furthermore, the distribution of the bed demand at any given point in time follows a Poisson distribution. Remarkably, this result holds for general LoS distributions, i.e., an $M_t/G/\infty$ queue in Kendall notation. However, this analytic approach does not allow for multiple classes of patients, and does not model interactions between multiple hospital wards, e.g., a GIM and an ICU. Therefore, the approach cannot be used to forecast bed demand in most actual hospital settings.

1.2 Contributions of This Paper

In this paper, we describe a discrete event simulation (DES) model for estimating the demand for general internal medicine (GIM) and intensive care unit (ICU) beds in a hospital during an infectious disease wave, e.g., one caused by Covid-19. This tool was developed in collaboration with Cambridge University Hospitals NHS Trust (CUH) and was used by the Trust for scenario analysis and demand forecasting over the course of the Covid-19 pandemic. Model inputs include historical and/or hypothetical daily Covid

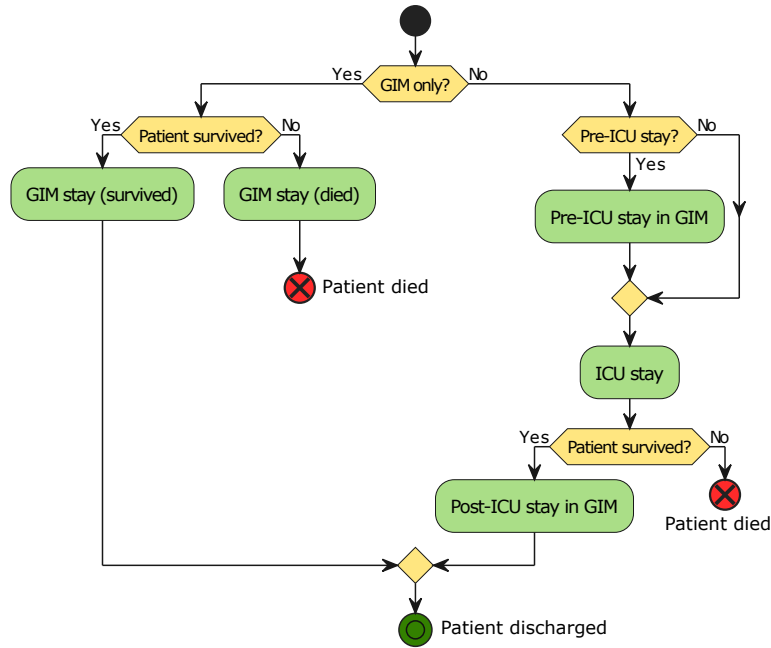


Figure 1: Flowchart depicting patient flow through the hospital in our Covid-19 simulation model.

patient arrival numbers, clinical features such as GIM and ICU LoS, and a threshold in days beyond which a GIM patient may be moved from an isolation (“red”) to a regular (“green”) ward bed. The outputs of the model contain the estimated daily GIM and ICU occupancy levels over the course of the simulation run. Statistical findings discovered during the course of our simulation model’s development include that the number of daily arrivals in each Covid-19 wave can be approximated using a Gamma curve, while patient LoSs can be approximated using Weibull distributions.

The Python code for our simulation model and dashboard is available at <https://gitlab.developers.cam.ac.uk/ycc39/covid-multidash-public>.

2 SIMULATION MODEL

Our simulation model is based on the model of patient flow by Melman et al. (2021) and depicted graphically in Fig. 1. Patients are divided into classes based on whether they were admitted to the ICU or the GIM only, and by whether they died in hospital. For simplicity, the LoS of ICU patients is not separated into the “survived” and “died” cases. There are thus five unique processes in our simulation model:

- GIM-only stay (patient survived)
- GIM-only stay (patient died in GIM)
- Pre-ICU stay in GIM
- ICU stay
- Post-ICU stay in GIM

The process of fitting LoS distributions to these five processes is described in Section 2.3.

2.1 Data Preprocessing

To more accurately model hospital bed demand caused by Covid-19 patients only, the following preprocessing rules were applied to estimate the portion of patient stays attributable to Covid-19:

1. If a patient is a readmitted patient, and their first positive sample was not collected during their *initial* stay, that patient is removed from the analysis.
2. If a patient stayed in the ICU and was discharged to the GIM before acquiring Covid-19, that ICU stay is ignored.
3. Current patients are assigned a “discharge time” equal to the data collection time, but are also marked as censored data.
4. The start timestamp of hospital-acquired cases is set to the time of the first positive Covid-19 testing result, while the start timestamp for community-acquired cases is the time of admission.
5. If a patient was readmitted to the hospital within 15 days:
 - If their initial stay was less than 0.3 days, only their readmission stay is retained and their start timestamp is adjusted accordingly.
 - If their readmission stay was less than 0.3 days, only their initial stay is retained and their end timestamp is adjusted accordingly.
 - If both stays exceed 0.3 days, their readmission stay is shifted forward in time to create a single continuous stay, and their end timestamp is adjusted accordingly.

Note that the source data only contains information for up to one readmission, and does not contain information on any ICU readmissions.

2.2 Modelling the Patient Arrival Process

Following the example of Palomo et al. (2020), we modelled the daily number of arrivals of Covid-19 patients to CUH using a Gamma function. It was found that a scaled Gamma function, denoted $f(t)$ below, fits well to the number of daily Covid-19 patient arrivals to CUH:

$$\tilde{f}(t) = \frac{t^{a-1} e^{-t/a}}{\Gamma(a) b^a}, \quad (a \geq 1)$$

$$f(t) = c \tilde{f}(t),$$

where $t = 0$ denotes the start of the Covid-19 wave being modelled and a and b are shape and scale parameters. The mode of $\tilde{f}(t)$ is at $\hat{t} = (a - 1)b$; thus, $c = N/\tilde{f}(\hat{t})$ is a vertical scaling factor such that the peak patient arrival rate is N patients per day.

Figure 2 shows the daily number of Covid-19 patient arrivals to CUH from Dec 2021 to Feb 2023, with Gamma curves fitted to five different arrival waves. The results demonstrate that Gamma curves can be used to fit the waves of Covid-19 hospital admission quite well. However, the fitted waves overlap; thus care must be taken when using such curves to forecast future Covid-19 admissions, which only model a single wave at a time. Furthermore, the figure demonstrates that the actual number of daily Covid-19 admissions is quite random despite following a Gamma trend. Therefore, hospitals need to be able to admit a larger number of admissions in a single day than modelled by a trendline.

Once the number of patient arrivals $f(t)$ is determined for each day, i.e. $t = 0, 1, 2, \dots$, the hour of each patient’s arrival within the day is modelled using a discrete distribution, with the probability within each hour based on the empirical distribution. Finally, the time of the arrival within the hour is modelled using a uniform distribution.

2.3 Estimation of Patient LoS and ICU Admission Probability

We used the *reliability* Python package (Reid 2023) to fit distributions to patient LoS data. The Weibull distribution was found to provide a good fit to patient LoS. The Weibull distribution is parameterised in *reliability* as $X \sim \text{Weibull}(\alpha, \beta, \gamma)$ with a cumulative distribution function of

$$F(t) = 1 - e^{-\left(\frac{t-\gamma}{\alpha}\right)^\beta}, \quad t \geq \gamma.$$

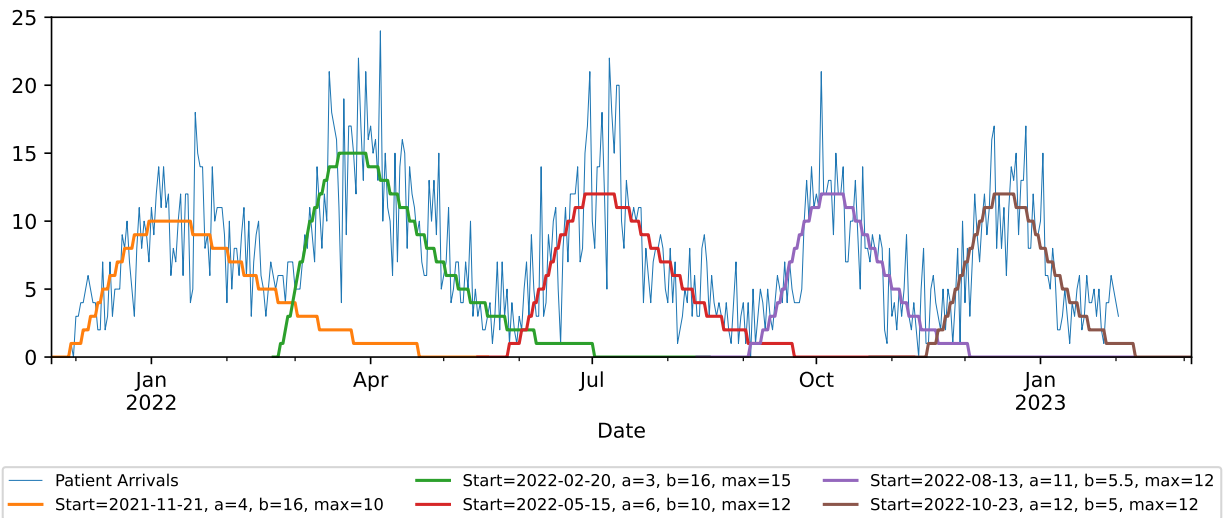


Figure 2: Daily Covid-19 arrivals to CUH, with five superimposed Gamma fits. Note that the “Start” dates refer to the zero-dates of the fitted Gamma functions, which may be considerably before the date of the first patient arrival in the fitted curves.

Table 1: Log-likelihood scores for various distributions as fitted to patient LOSs, for patients admitted March–May 2022. Rankings are shown in parentheses.

	GIM (survived)	GIM (died)	Pre-ICU	ICU	Post-ICU
Weibull	-8163.64 (2)	-2048.32 (1)	-207.449 (2)	-765.906 (2)	-507.992 (3)
Lognormal	-8158.31 (1)	-2083.41 (4)	-197.148 (1)	-756.628 (1)	-507.884 (2)
Gamma	-8179.8 (3)	-2049.1 (2)	-215.39 (3)	-768.092 (3)	-507.573 (1)
Exponential	-8207.86 (4)	-2051.44 (3)	-242.884 (4)	-768.759 (4)	-508.707 (4)

In general, patient LoS parameters and ICU admission probability (i.e., the probability that a patient will enter the ICU at any point during their hospital stay) for each Covid-19 wave was estimated using patient data from the previous wave. In contrast with earlier iterations of our simulation model, it was decided to compute a single LoS distribution for each type of patient stay, without stratification by age cohort. This is due to the low number of patients in certain cohorts, e.g., paediatric patients who died in GIM. Similarly, the current Python model does not split ICU patients by survival outcome, whereas this was performed for certain waves towards the beginning of the Covid-19 pandemic.

2.3.1 Numerical Examples

Table 1 shows log-likelihood scores for the Weibull, Gamma, lognormal, and exponential distributions for each type of patient LoS. The Weibull distribution is demonstrated to be a good compromise for all the LoS types considered. Note that the *reliability* package adds optional location parameters to all the distributions considered; thus shifting the entire distribution along the x -axis.

Figure 3 shows probability plots for fitted Weibull distributions for patient LoSs, for patients admitted between March and May 2022. The results demonstrate good fit for GIM-only patients when the LoS is sufficiently large (at least one day), such patients having more effect on bed occupancy demand in the hospital model. The results also show moderately good fit for the pre-ICU, ICU, and post-ICU stays of ICU patients, again with better fit when ignoring the shortest stays.

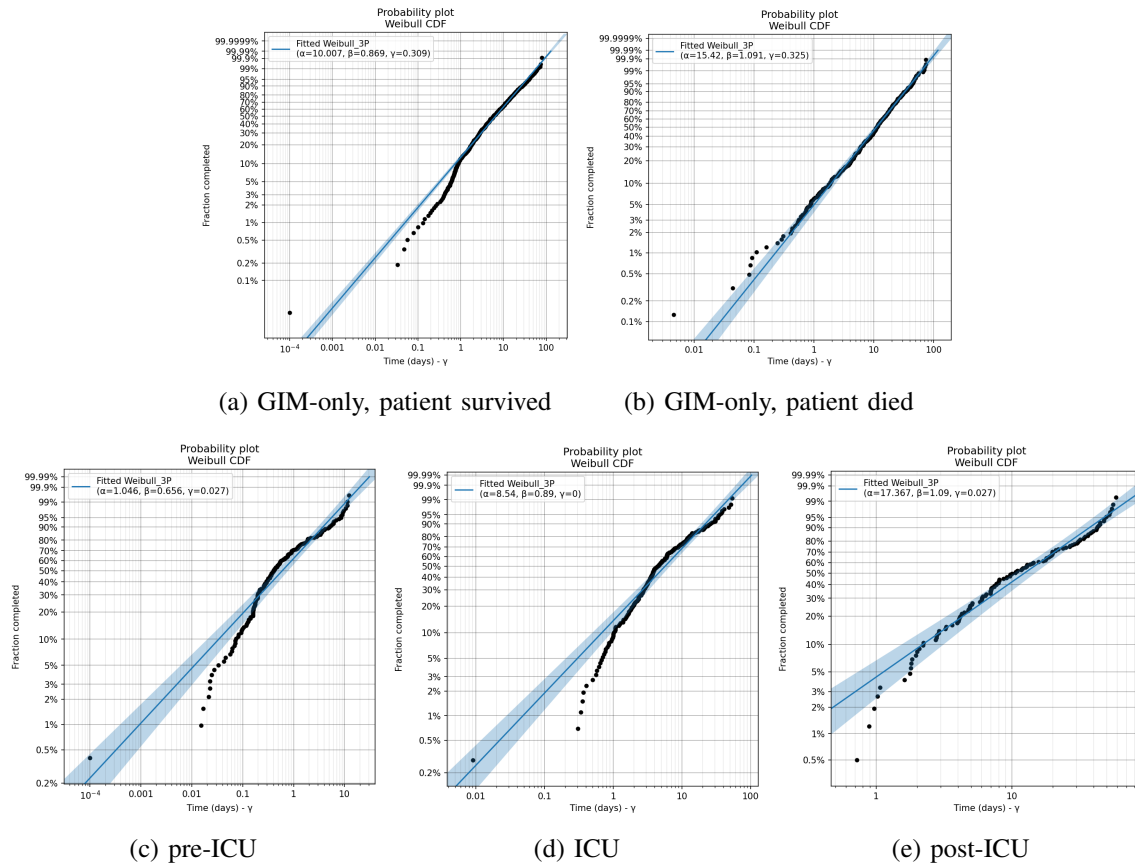


Figure 3: Probability plots for fitted Weibull distributions to patient LoS, for patients admitted March–May 2022.

2.3.2 “Red” GIM Beds and Step-down

During the course of the Covid-19 pandemic, it was recognised that some patients admitted with Covid-19 would require additional hospital stay for other reasons, but should no longer be classified as infectious. Therefore, a rule was added to our model such that Covid-19 patients in the GIM would be “stepped down” from a “red” bed to a “green” bed if their total Covid-19 LoS has exceeded 10 days. Note that this is an approximate approach as actual policy depends on the availability of in-hospital testing and the recovery rate of each patient.

2.4 Discrete-event Simulation with SimPy

Initially, a simulation model was created in Arena, but we eventually switched to a Python implementation simulation using the *SimPy* library, while removing features that were no longer used. The result was a simulation program that runs over one order of magnitude faster than the original Arena model. The *SimPy* library used for the new simulation program is based on Python generators, which yield events that are handled by the main event loop. Patients are added to the simulation environment using `env.process(new_arrival(env))`, where the `new_arrival` function launches either a `icu_process` or a `gim_process` using `yield from`. Note that `icu_process` also includes the pre-ICU and possibly post-ICU stay of an ICU patient in GIM.

To model bed demand, the GIM and ICU are modelled as infinite-capacity *Resources*, with their usage levels recorded in a *pandas* (McKinney 2022) dataframe upon each patient arrival, departure, and

CUH Covid Multi-Dashboard: Scenario Modelling

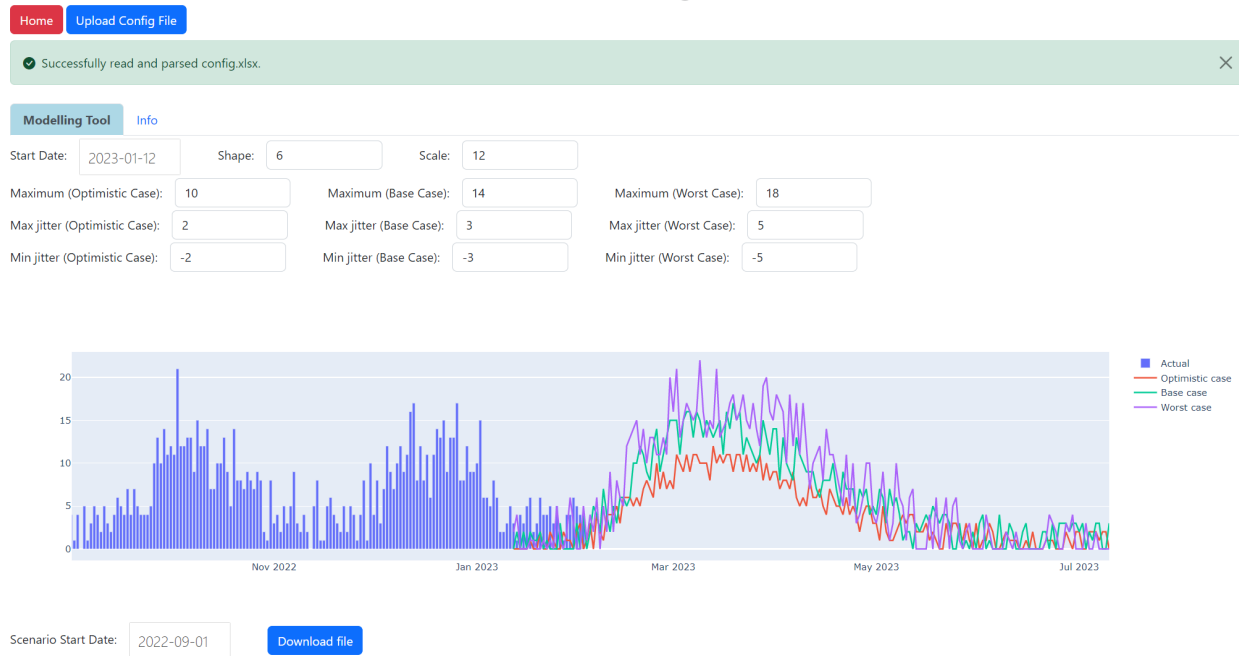


Figure 4: Example screenshot from the Python web application.

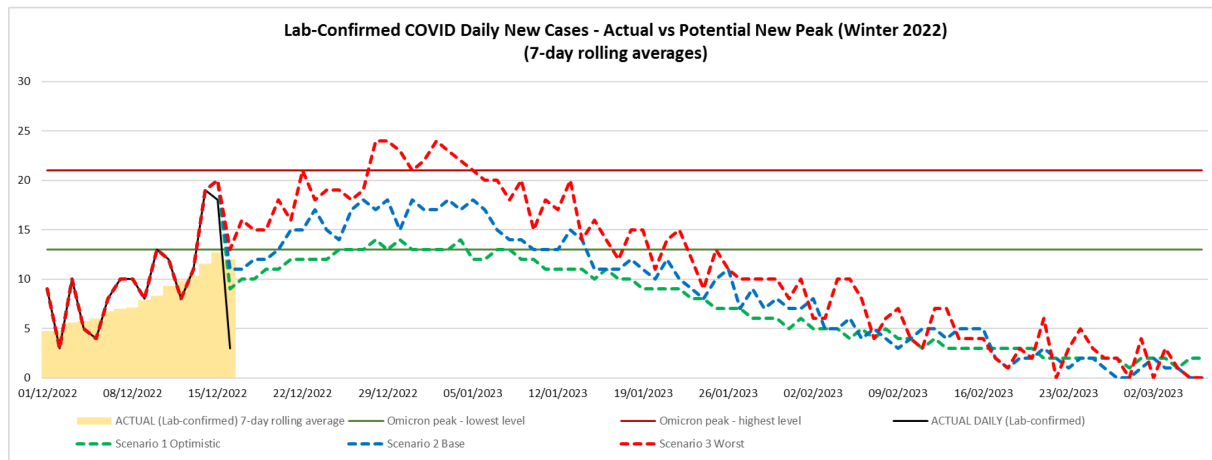
GIM/ICU transfer. These dataframes are later post-processed to generate the maximum patient occupancy for each day of the simulation. The choice to implement GIM and ICU as having infinite capacity is so that we can observe how many beds would be needed if all Covid-19 patients were admitted using the normal pathway. To model red-bed step-down, a virtual `Resource` was created that is seized simultaneously with a GIM bed, but is released at step-down (i.e. when the patient is transferred to a “green” GIM bed).

2.5 Web Interface

To make it easier to use the proposed simulation model, a web interface was created for our DES model implementation, using the Python libraries *Plotly*, *Dash*, and *Flask* (Mayer, Schroeder, and Ward 2022). The application consists of three dashboards for patient analysis, scenario creation, and discrete-event simulation, respectively. The dashboards are designed so that the output files of each dashboard can be used as input to the next.

- The *patient analysis* dashboard takes raw Covid-19 patient data and a range of dates, and returns the number of daily arrivals, their distribution by hour of day, and the parameters of the fitted LoS distributions for the five types of patient stays. Preprocessing is applied to the patient LoS data as outlined in Section 2.1.
- The *scenario modelling* dashboard provides controls for creating simulation scenarios based on Gamma curves for the predicted number of daily patient arrivals. According to the needs of CUH, three scenarios can be produced using the dashboard, each with a different peak number of daily admissions. A screenshot of this dashboard is provided as Fig. 4. By design, the generated scenario numbers are only applied during simulation if actual admission numbers are not available for the given day.
- The *simulation* dashboard runs the simulation model according to the patient and scenario parameters set using the previous two dashboards. The number of simulation runs can be chosen, with the mean, 5th percentile, and 95th percentile of the simulated GIM, ICU, and red-bed occupancy values

Admissions – new scenarios



Sources: Python model 16th December 2022, COVID Modelling Group
CHEQS data to 16th December 2022

COVID and Winter Respiratory Viruses Overview

2

Figure 5: Graph from CUH internal slide pack, 16 Dec. 2022.

plotted against time. If available, actual occupancy numbers can be plotted against the simulated values for validation purposes, as shown in Fig. 7 in Section 3.

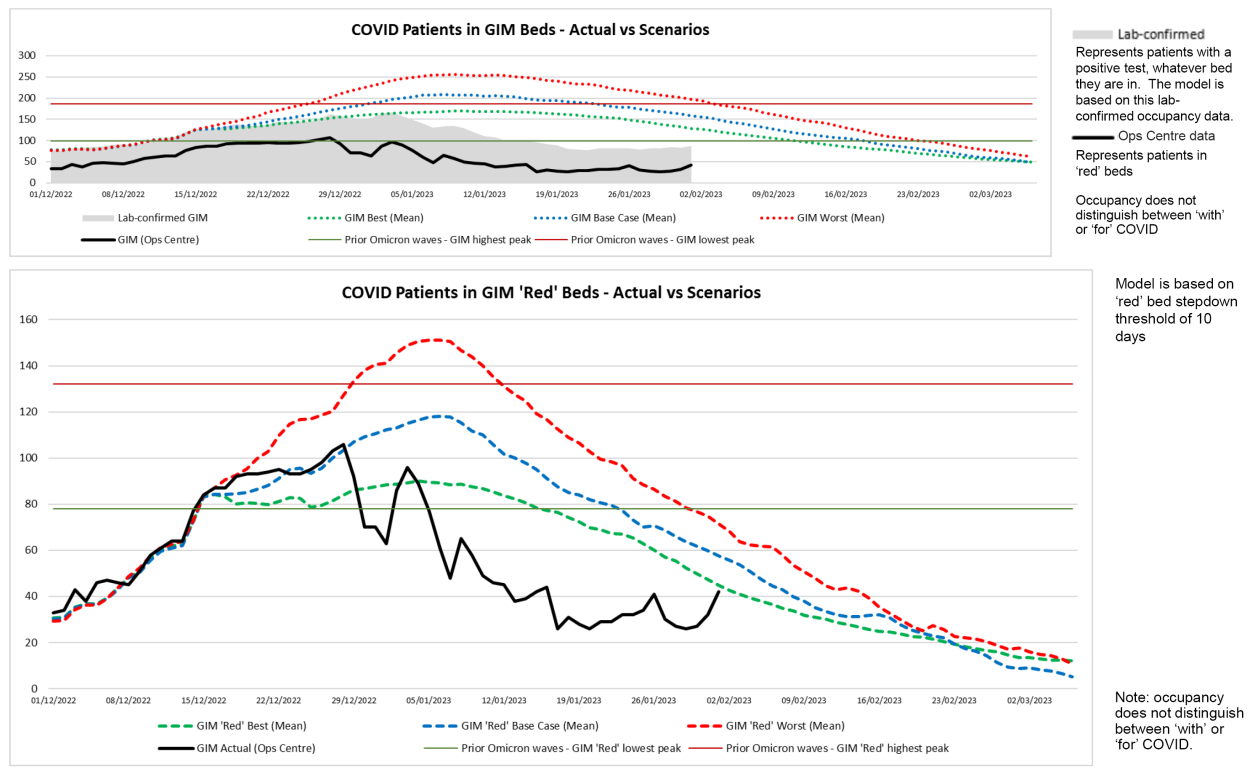
3 USAGE OF THE SIMULATION MODEL AT CAMBRIDGE UNIVERSITY HOSPITALS NHS FOUNDATION TRUST

The simulation model described in this paper has been used to support the Cambridge University Hospitals Trust's (CUH) Covid-19 response. Simulation results for the modelled scenarios assisted CUH to provision suitable number of “red” beds throughout the pandemic and allowed hospital administrators to assess the impact of operational changes such as adjustments to the step-down policy as described in Section 2.3.2, particularly with regard to non-Covid care capacity. The key benefit of the model was to plan for the provisioning of “red” beds during the upswing of each Covid-19 wave, and also to identify the availability of new “green” beds during the downswing as the requirements for “red” bed decreases.

Case study: Winter wave 2022-2023

Figure 2 shows a wave of Covid-19 infections during Dec. 2022 to Jan. 2023. On 15 Dec. 2022, it was determined that the wave should be modelled with a peak of 13, 17, and 21 patients per day for each of the three scenarios, respectively, all “at approximately Jan. 1”. The number of daily arrivals for each of these scenarios is shown in Fig. 5. Note that Fig. 2 suggests an earlier and lower peak than modelled in Fig. 5; part of the reason for this is the large variability of patient admissions from day to day and the difficulty in predicting the height and timing of a wave while it is still on the increase. In particular, the predicted scenarios for this wave were made just after an unusual spike in the number of daily admissions, adding to the level of uncertainty.

Fig. 6 shows the available daily GIM occupancy data as of 1 Feb. 2023. During the period under study, the peak number of daily admissions matched our “optimistic” scenario, but then fell more sharply than predicted; therefore, in late Jan. 2023, the number of daily admissions was much less than any of



Sources: Python Model 16th December 2022, CUH COVID Modelling Group; CHEQS and Ops Centre data to 1st February 2023

COVID and Winter Respiratory Viruses Overview

8

Figure 6: Graphs from CUH internal slide pack, 1 Feb. 2023.

our modelled scenarios. As a result, the Covid-19 occupancy level of the GIM beds was much lower than expected during this period.

Finally, Fig. 7 shows a *retrospective* analysis of the Dec. 2022 – Jan. 2023 wave, using the fitted wave parameters shown in Fig. 2 (brown curve) and patient LoS parameters computed from patients between 1 Sep. and 15 Nov. 2022 (purple curve). The results demonstrate a much closer fit to the actual occupancy levels of the GIM, showing that the accuracy of the simulation model is primarily limited by the accuracy of the modelled scenarios in terms of the predicted peak level and duration of the infection wave. Note that as the simulation model only tracks a single infection wave, the increase in GIM occupancy in mid-February 2023 is not captured by the simulation results. Furthermore, due to low admission numbers, the uncertainty of the ICU occupancy simulation results is much larger than for GIM, relative to the mean occupancy level.

4 CONCLUDING REMARKS

In this paper, we described a discrete event simulation framework for estimating GIM and ICU demand in a hospital for patients with Covid-19, which we have encapsulated within a Python *Flask* application. The framework has been used to support CUH's Covid-19 response. As shown in Section 3, the framework is quite accurate at estimating Covid-19 bed occupancy given the daily number of patient arrivals. However, accurate prediction of daily patient arrivals remains a challenge, thus affecting the accuracy of our simulation framework for forecasting purposes.

Note that in addition to the simulation itself, our framework includes tools for LoS distribution fitting and scenario generation. It was found that Weibull distributions are a good approximation for each type of

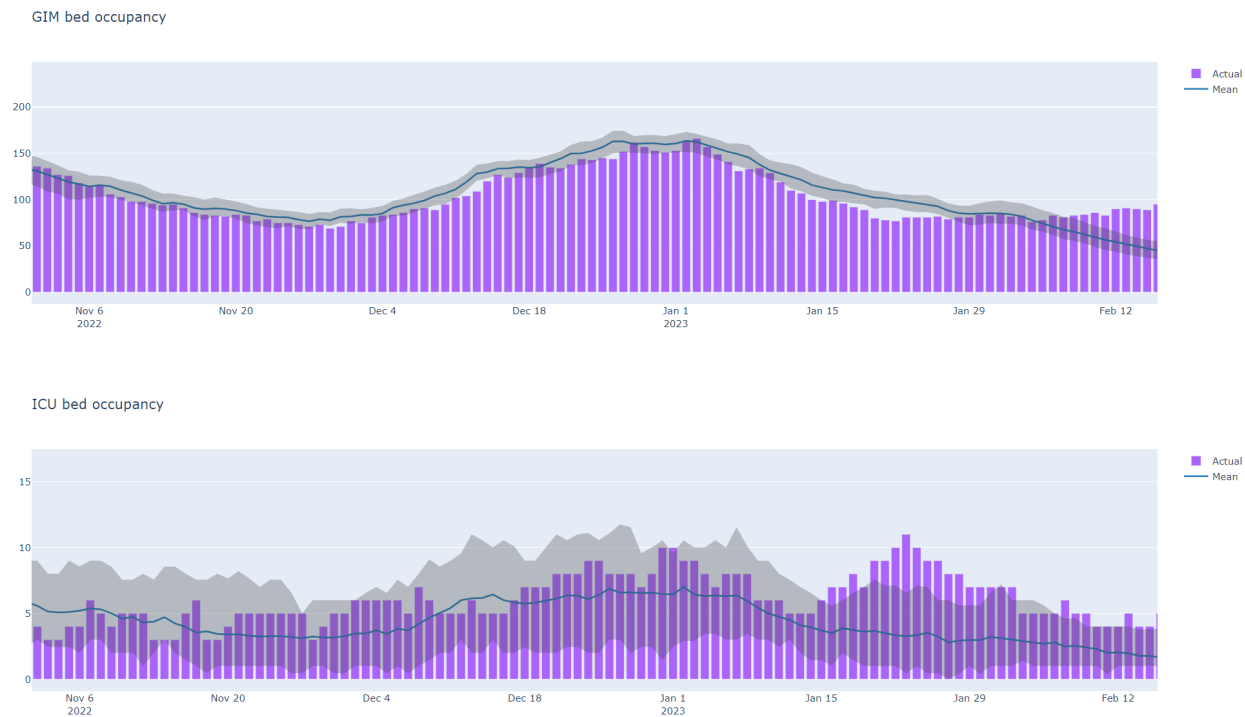


Figure 7: Simulation results for the Winter 2022-2023 Covid-19 wave, using wave and LoS parameters computed retrospectively. The grey bands represent 90% credible intervals of the GIM/ICU bed occupancies.

LoS in the simulation model, as shown in Fig. 3. Furthermore, Gamma functions form a good basis for modelling daily hospital arrivals in each Covid-19 wave.

Potential improvements to the simulation framework and web application include:

- Support for multiple infectious diseases simultaneously. For example, the extended simulation model could be used to model winter influenza and respiratory syncytial virus (RSV) in addition to Covid-19.
- Ability to add a baseline number of daily cases to a wave. Currently, the scenario modelling component of our web application assumes that daily new cases are zero at the start and end of the modelled wave.
- Stratification of patient LoS parameters by age group or other demographic groupings, especially for GIM patients. On the other hand, the volume of ICU patients in our scenarios is typically low and thus further stratification of such patients may not be beneficial. This is because it is more difficult to obtain accurate LoS parameter estimates when the sample size is small.
- Support for other LoS distribution types, e.g. Gamma and lognormal.
- Per-replication randomization of daily arrival numbers when non-zero jitter is applied (see Fig. 4). Currently, the same daily arrival numbers are used for all replications in a simulation run.

ACKNOWLEDGMENTS

The authors thank all members of the University of Cambridge / Cambridge University Hospitals NHS Foundation Trust COVID / Respiratory Viruses Modelling Group for their input regarding the simulation model development and modelled scenarios.

The research presented in this article was funded by Cambridge University Hospitals NHS Foundation Trust.

NJM was supported by the UKRI Medical Research Council (TSF ref. MR/T032413/1) and NHS Blood and Transplant (grant ref. WP15-02).

DATA AVAILABILITY

Aggregate data supporting this study are available from the authors upon reasonable request.

REFERENCES

- Bovim, T. R., A. N. Gullhav, H. Andersson, J. Dale, and K. Karlsen. 2021. "Simulating Emergency Patient Flow During the COVID-19 Pandemic". *Journal of Simulation*:1–15.
- Brauer, F., C. Castillo-Chavez, and Z. Feng. 2019. *Mathematical Models in Epidemiology*. New York, NY: Springer New York.
- Currie, C. S., J. W. Fowler, K. Kotiadis, T. Monks, B. S. Onggo, D. A. Robertson, and A. A. Tako. 2020. "How Simulation Modelling Can Help Reduce the Impact of COVID-19". *Journal of Simulation* 14(2):83–97.
- Garcia-Vicuna, D., F. Mallor, and L. Esparza. 2020. "Planning Ward and Intensive Care Unit Beds for COVID-19 Patients Using a Discrete Event Simulation Model". In *Proceedings of the 2020 Winter Simulation Conference*, edited by K.-H. Bae, B. Feng, S. Kim, S. Lazarova-Molnar, Z. Zheng, T. Roeder, and R. Thiesing, 759–770. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Gomersall, C. D., D. Y. H. Tai, S. Loo, J. L. Derrick, M. S. Goh, T. A. Buckley, C. Chua, K. M. Ho, G. P. Raghavan, O. M. Ho, L. B. Lee, and G. M. Joynt. 2006. "Expanding ICU Facilities in an Epidemic: Recommendations Based on Experience From the SARS Epidemic in Hong Kong and Singapore". *Intensive Care Medicine* 32(7):1004–1013.
- Locke, C. J., B. Koo, S. W. Baron, J. Shapiro, and J. Pacifico. 2021. "Creation of a Medical Ward From Non-Clinical Space Amidst the Covid-19 Pandemic". *Journal of Evaluation in Clinical Practice* 27(4):992–995.
- Mayer, C., A. Schroeder, and A. M. Ward. 2022. *The Book of Dash: Build Dashboards with Python and Plotly*. San Francisco, CA, USA: No Starch Press, Incorporated.
- McKinney, W. 2022. *Python for Data Analysis: Data Wrangling with Pandas, NumPy, and Jupyter*. 3rd ed. Sebastopol, CA, USA: O'Reilly.
- Melman, G. J., A. K. Parlikad, and E. A. B. Cameron. 2021. "Balancing Scarce Hospital Resources During the COVID-19 Pandemic using Discrete-Event Simulation". *Health Care Management Science* 24:356–374.
- Murat, M., S. Köse, and S. Savaşer. 2020. "Determination of Stress, Depression and Burnout Levels of Front-Line Nurses During the COVID-19 Pandemic". *International Journal of Mental Health Nursing* 30(2):533–543.
- Nashwan, A. J., A. A. Abujaber, R. C. Villar, A. Nazarene, M. M. Al-Jabry, and E. C. Fradelos. 2021. "Comparing the Impact of COVID-19 on Nurses' Turnover Intentions before and during the Pandemic in Qatar". *Journal of Personalized Medicine* 11(6):456.
- Palomo, S., J. Pender, W. Massey, and R. C. Hampshire. 2020. "Flattening the Curve: Insights From Queueing Theory". <https://arxiv.org/abs/2004.09645>.
- Reid, Matthew 2023. "MatthewReid854/reliability: v0.8.8". Zenodo. <https://doi.org/10.5281/ZENODO.7660623>.
- Sagy, I., E. Rosenberg, and L. Barski. 2021. "Converting a Standard Internal Medicine Ward Into an Isolation Unit During the COVID-19 Outbreak". *International Journal of Clinical Practice* 75(5):e13979.
- Sharifi, M., A. A. Asadi-Pooya, and R. S. Mousavi-Roknabadi. 2020. "Burnout Among Healthcare Providers of COVID-19; a Systematic Review of Epidemiology and Recommendations". *Archives of Academic Emergency Medicine* 9:e7.
- Sprung, C. L., G. M. Joynt, M. D. Christian, R. D. Truog, J. Rello, and J. L. Nates. 2020. "Crisis Level ICU Triage Is About Saving Lives". *Critical Care Medicine* 49(1):e102–e103.
- Warde, P. R., S. S. Patel, T. D. Ferreira, H. B. Gershengorn, M. C. Bhatia, D. J. Parekh, K. J. Manni, and B. S. Shukla. 2021. "Linking Prediction Models to Government Ordinances to Support Hospital Operations During the COVID-19 Pandemic". *BMJ Health & Care Informatics* 28(1):e100248.
- Weissman, G. E., A. Crane-Droesch, C. Chivers, T. Luong, A. Hanish, M. Z. Levy, J. Lubken, M. Becker, M. E. Draugelis, G. L. Anesi, P. J. Brennan, J. D. Christie, C. W. Hanson, M. E. Mikkelsen, and S. D. Halpern. 2020. "Locally Informed Simulation to Predict Hospital Capacity Needs During the COVID-19 Pandemic". *Annals of Internal Medicine* 173(1):21–28.

AUTHOR BIOGRAPHIES

YIN-CHI CHAN obtained his B.Math. degree at the University of Waterloo in 2010 and his M.Sc. and Ph.D. degrees at City Univ. of Hong Kong, in 2011 and 2017, respectively. He is currently a Research Associate at the Institute for Manufacturing (IFM) at University of Cambridge. His current research interests include discrete-event simulation modelling and analysis of healthcare systems and the use of simulation methods in digital twins applications. His email address is ycc39@cam.ac.uk.

KAYA DREESBEIMDIEK is a doctoral student at the Institute for Manufacturing at the University of Cambridge. Her research focuses on resilient hospital operations management. Prior to starting her PhD, Kaya obtained an MPhil in Industrial Systems, Manufacture and Management from the University of Cambridge. She holds undergraduate degrees in (Technology) Management from Ludwig-Maximilians-University and the Center for Digital Technology and Management at Technical University Munich. Her email is kmd46@cam.ac.uk.

AJITH KUMAR PARLIKAD is Professor of Asset Management at the University of Cambridge. He leads the Asset Management Research Group at the Institute for Manufacturing, and is the principal investigator of the Digital Hospitals research programme funded by Cambridge University Hospitals (CUH) NHS Foundation Trust. His email is aknp2@cam.ac.uk.

TOM RIDGMAN graduated with a degree in Mechanical Engineering from Southampton in 1973. The first half of his career was spent in the automotive industry where he held a number of senior management posts. The second half of his career was spent mostly at the IFM at the University of Cambridge. He taught post graduates manufacturing and operations management and has professional and research interests in standards of competence for professional engineers and engineering education. His email is twr20@cam.ac.uk.

NICHOLAS J. MATHESON is an Associate Professor of Virology at the Cambridge Institute of Therapeutic Immunology and Infectious Disease (CITIID), Honorary Consultant in Infectious Diseases at the Cambridge University Hospitals NHS Foundation Trust (CUHNFT), and Principal Investigator for NHS Blood and Transplant. He qualified in medicine from the University of Oxford in 2003, undertook specialist training in General Medicine and Infectious Diseases in London and Cambridge, and spent 2 years as a Visiting Scientist at the Massachusetts Institute of Technology. His clinical work and research are focused on important human viruses, particularly HIV and SARS-CoV-2. As a clinician scientist, he was closely involved with the COVID-19 pandemic response at CUHNFT, the University of Cambridge and the UK higher education sector. His email is njm25@cam.ac.uk.

BEN WARNE qualified from the University of Cambridge School of Clinical Medicine in 2010. He was awarded a PhD from the University of Cambridge in 2022, and subsequently became an academic clinical lecturer in infectious diseases based in the Cambridge Institute of Therapeutic Immunology and Infectious Disease (CITIID). Ben's research interests include the use of genomic and healthcare data to study the impact of nosocomial transmission of common bacterial and viral infections of public health concern. His email is bw252@cam.ac.uk.

DENISE FRANKS is Assistant Director of Planning and Development at Cambridge University Hospitals NHS Foundation Trust and chairs the COVID / Respiratory Viruses Modelling Group. Her email is denise.franks3@nhs.net.