

Real-time forecasting of COVID-19 bed occupancy in wards and Intensive Care Units

Stef Baas · Sander Dijkstra · Aleida Braaksma · Plom van Rooij · Fieke J. Snijders · Lars Tiemessen · Richard J. Boucherie

Received: January 12, 2021/ Accepted: will be inserted by the editor

Abstract This paper presents a mathematical model that provides a real-time forecast of the number of COVID-19 patients admitted to the ward and the Intensive Care Unit (ICU) of a hospital based on the predicted inflow of patients, their Length of Stay (LoS) in both the ward and the ICU as well as transfer of patients between the ward and the ICU. The data required for this forecast is obtained directly from the hospital's data warehouse. The resulting algorithm is tested on data from the first COVID-19 peak in the Netherlands, showing that the prediction is very accurate. The forecast may be visualised in real-time in the hospital's control centre and is used in several Dutch hospitals during the second COVID-19 peak.

Keywords COVID-19 · forecast · bed occupancy · network of infinite server queues · Richards' curve · Kaplan-Meier estimator

1 Introduction

The COVID-19 pandemic impacts people's health, jobs and well-being and puts an enormous strain on healthcare resources. This is also the case in the Netherlands, where hospital resources are under pressure due to the large number of hospitalised COVID-19 patients, which moreover results in reduction of resources available

Stef Baas · Sander Dijkstra · Aleida Braaksma · Richard J. Boucherie
Center for Healthcare Operations Improvement and Research (CHOIR), University of Twente,
Enschede, The Netherlands

Plom van Rooij
Elisabeth-TweeSteden Ziekenhuis, Tilburg, The Netherlands

Fieke J. Snijders
Leiden University Medical Centre, Leiden, The Netherlands

Lars Tiemessen
Rijnstate, Arnhem, The Netherlands

for non-COVID-19 patients [25]. An accurate forecast of the number of COVID-19 patients being hospitalised supports allocation of the resources required for treatment of both COVID-19 and non-COVID-19 patients. This paper presents a mathematical model that provides a real-time forecast of the number of COVID-19 patients admitted to the ward and the Intensive Care Unit (ICU) of a hospital based on the predicted inflow of patients, their Length of Stay (LoS) in both the ward and the ICU as well as transfer of patients between the ward and the ICU. The data required for this forecast is obtained directly from the hospital's data warehouse and the forecast is available to the hospital in real-time.

Forecasting the number of hospitalised COVID-19 patients is required to determine the resource allocation to COVID-19 and non-COVID-19 patients. A COVID-19 patient's medical condition may change rapidly and unexpectedly [20]. As a consequence, it is not possible to accurately forecast the COVID-19 patients' resource requirements many days ahead of time [12]. We therefore focus on forecasting the number of hospitalised COVID-19 patients one to five days ahead of time. In particular, we are interested in forecasting the mean number of patients present and the risk of bed shortage, expressed as the probability that the maximum number of COVID-19 patients in the ward and ICU exceeds a pre-specified safety level. Several studies have demonstrated the Erlang loss model (or $M/G/c/c$ queue) to be suitable for dimensioning of isolated wards (e.g., [5, 28]) and ICUs (e.g., [2, 6]). Such models typically assume constant arrival rates and LoS distributions. In our context, however, the LoSs of patients may vary over time due to improved treatment, arrivals of patients are non-stationary, and patients may transfer between ward and ICU. In such settings with time-varying load, a network of Erlang loss queues can well be approximated by a network of infinite server queues [1, 18], either using a Pointwise Stationary Approximation or a Modified Offered Load Approximation. The advantage of this approximation is that it allows explicit evaluation of performance measures. Hence, for our case, the most suitable model is a network of two infinite server queues with time-varying Poisson arrivals and generally distributed time-varying LoSs. We build upon the results for networks of infinite server queues as presented in [4, 18, 26] to allow for time-varying arrival rates and patient-specific time-varying LoS distributions.

With data on COVID-19 patients becoming more and more available, prediction of the infection rate a few days ahead of time [10, 30], and of the LoS [21] is possible. Predictions of the number of hospitalised COVID-19 patients based on regression methods are, e.g., reported in [7, 8, 17]. The LoS distribution of COVID-19 ICU patients in the United Kingdom is fitted to probability distributions in [27]. Our model combines such predictions to forecast the number of patients residing in the COVID-19 ward and ICU. We have chosen to predict the arrival rates and LoS directly from the hospital's data warehouse as external data does not represent the case mix of the hospital. Our method requires a complete set of time stamps for patient admissions, transfers and discharges that is made available by the participating hospitals. We use a Richards' curve [24] to predict the arrival rates. The Richards' curve is a growth model that can be used to describe the cumulative total number of hospitalised COVID-19 patients, i.e., ward and ICU combined. The Richards' curve was introduced to describe processes in biological systems, but has recently gained popularity in predicting the outbreak of diseases. For instance, the Richards' curve has been successfully applied to predict the daily number of new

COVID-19 infection cases in (provinces of) China and countries in Europe [15, 30]. We estimate the distribution of the LoS in both the COVID-19 ward and ICU using a Kaplan-Meier estimator [13]. Analytical evaluation of the exact expressions for the rates of the time-dependent distribution of bed occupancy is prohibitive. Therefore, in our forecasting algorithm we use a Monte-Carlo method to sample from the LoS distribution. We forecast the mean number of patients present and the risk of bed shortage, expressed as the maximum number of COVID-19 patients at the ward and ICU at a number of subsequent days, including the corresponding prediction intervals. In particular, we sample the patient trajectories in the Poisson Arrival Location Model [18] that determines the queue occupancy in our network of infinite server queues. As such, our approach may be viewed as a data-driven approach that predicts the number of hospitalised COVID-19 patients based on estimated arrival rates and LoSs, justified by an underlying queueing model. The algorithm is implemented in **R version 3.6.3** and first tested on data from the first COVID-19 peak for four hospitals in the Netherlands. The forecast is found to be very accurate.

Forecasting the number of hospitalised COVID-19 patients is difficult [12]. As stated above, various approaches exist based on regression methods [7, 8, 10, 17, 30], and on estimation of the LoS [21, 27]. The key to the accuracy of our forecast is that admissions are predicted according to a growth curve, which has strong links with the Susceptible-Exposed-Infected-Recovered (SEIR) model [19], while also taking the joint effect of the arrival rates and the LoS into account via the underlying network of infinite server queues. This enables evaluation of the future evolution of bed occupancy via the trajectories of the Poisson Arrival Location Model, taking into account patient admissions, transfers and discharges. The algorithm currently runs in four hospitals to forecast the number of hospitalised COVID-19 patients during the second peak the Netherlands is currently facing.

This paper is organised as follows. Section 2 presents our modelling assumptions and the network of two infinite server queues that we propose to forecast the number of hospitalised COVID-19 patients. Section 3 describes the statistical forecasting approach used in our method, and Section 4 presents forecasts for our method using data of the first COVID-19 peak in the Netherlands. Occupancy is most easily forecast in a large hospital with a homogeneous patient mix. In Section 4, we present forecasting results for a medium-sized, academic hospital as well as for a number of larger hospitals, and reflect on the accuracy of our forecasts. Finally, Section 5 wraps up the paper with concluding remarks and our aims for further research: extending our model to a regional prediction model including patient transfers between hospitals.

2 Model

Upon COVID-19 infection, some patients develop mild or no symptoms, whereas others develop symptoms that require hospitalisation at either the ward or Intensive Care Unit (ICU) depending on, e.g., the need for artificial respiration [11, 20]. While hospitalised, a patient's condition may worsen, resulting in a transfer from the ward to the ICU or death, or a patient may recover, resulting in a transfer from ICU to ward or a discharge from the ward. Some patients admitted

to the ward have treatment restrictions that prohibit transfer from ward to ICU. As COVID-19 is a relatively new disease, the evolution of a patient's condition and the effect of treatment are still under investigation. A patient's Length of Stay (LoS) at the ward or ICU may depend on patient characteristics such as age, gender, length, BMI and treatment restrictions, may differ considerably between hospitals due to different treatment protocols or differences in case mix, and may also change over time due to, e.g., improved treatment [3, 20, 21]. Therefore, we estimate the distribution of the LoS and the probability of patient transfers between ward and ICU from the data on COVID-19 patients in the hospital's data warehouse. Arrivals of new patients are influenced by the number of infections in the hospital's region, and also by the characteristics of the hospital, e.g., more severely ill patients will be admitted to university medical centres, whereas less ill patients may be treated in local hospitals and may be transferred if their condition worsens. Therefore, we predict the rate of admittance of COVID-19 patients from the hospital's data warehouse record of earlier admissions. Figure 1 depicts the flow of patients in the COVID-19 ward and ICU.

We now develop a network of two infinite server queues with multiple patient-types, time-dependent arrival process, and general and time-dependent LoS that records the number of hospitalised patients. Consider a hospital that admits COVID-19 patients to its ward and ICU. The arrival rate of patients is determined by the number infected in the hospital's region and a consequence of the infection rate in that region. In agreement with arrivals to Emergency Departments, the arrival process may be modelled as a Poisson process with time-dependent rate [29]. Consider a set C of patient characteristics, including, e.g., age, length, BMI,

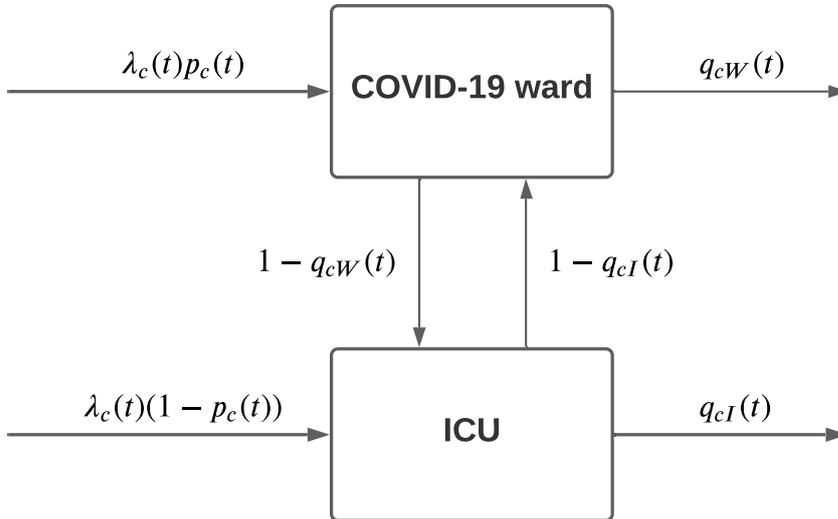


Fig. 1: Patient flows.

and treatment restrictions. Let patients with characteristics $c \in C$ arrive with rate $\lambda_c(t)$, where t denotes time. A fraction $p_c(t)$ of these patients is admitted to the ward, the other patients are admitted to the ICU. Let random variables $L_{cW}(t)$ and $L_{cI}(t)$ denote the LoS of a patient with characteristics $c \in C$ admitted or transferred to the ward and ICU, respectively, at time t . Let $q_{cW}(t)$, $q_{cI}(t)$ denote the probability that a patient admitted or transferred to the ward or ICU at time t is discharged or dies. Then $1 - q_{cW}(t)$ and $1 - q_{cI}(t)$ are the probabilities of transfer from ward to ICU and vice versa upon completion of the LoS. We assume that patients do not interfere with each other, hence that all random variables related to patients' arrival, transfer, and LoS are independent.

Characteristics of a patient's LoS and transfer probabilities are related to the time the patient is admitted. Therefore, we may incorporate these characteristics in the set C . As a consequence, we may model the system as a network of two infinite server queues with multiple job-types, time-varying arrival rates and general LoS distribution. To this end, let $N_{cW}(t)$ and $N_{cI}(t)$ record the number of patients with characteristics c at time t in the ward and ICU, respectively. These random variables have a time-dependent Poisson distribution, for $n_{cW}, n_{cI} = 0, 1, 2, \dots$:

$$\mathbb{P}[N_{cW}(t) = n_{cW}, N_{cI}(t) = n_{cI}] = \frac{\rho_{cW}(t)^{n_{cW}}}{n_{cW}!} \frac{\rho_{cI}(t)^{n_{cI}}}{n_{cI}!} e^{-(\rho_{cW}(t) + \rho_{cI}(t))}, \quad (1)$$

where the means $\rho_{cW}(t)$, $\rho_{cI}(t)$ are, in closed-form, determined by $\lambda_c(t)$, $p_c(t)$, $L_{cW}(t)$, $L_{cI}(t)$, $q_{cW}(t)$ and $q_{cI}(t)$, via the Poisson Arrival Location Model, as integrals over a location function, see [18, Theorem 2.1]. If the LoSs are exponentially distributed with rates μ_{cW} and μ_{cI} at the ward and ICU, then the means $\rho_{cW}(t)$, $\rho_{cI}(t)$ may be obtained from

$$\begin{aligned} \frac{1}{\mu_{cW}} \frac{d\rho_{cW}(t)}{dt} &= \lambda_c(t)(1 - p_c(t)) + \rho_{cI}(t)(1 - q_{cI}(t)) - \rho_{cW}(t), \\ \frac{1}{\mu_{cI}} \frac{d\rho_{cI}(t)}{dt} &= \lambda_c(t)p_c(t) + \rho_{cW}(t)(1 - q_{cW}(t)) - \rho_{cI}(t), \end{aligned}$$

with initial conditions

$$\rho_{cW}(0) = \rho_{cW}^*, \quad \rho_{cI}(0) = \rho_{cI}^*, \quad (2)$$

that reflect the current number of hospitalised patients in the ward and ICU at the starting time 0 of our forecasting period [4]. If the LoS has a general distribution not depending on the arrival time of the patients, and transfer probabilities do not depend on time, the means $\rho_{cW}(t)$, $\rho_{cI}(t)$ are obtained as

$$\rho_{cW}(t) = \mathbb{E} \left[\int_{t-L_{cW}}^t \lambda_{cW}^+(u) du \right], \quad (3)$$

$$\rho_{cI}(t) = \mathbb{E} \left[\int_{t-L_{cI}}^t \lambda_{cI}^+(u) du \right], \quad (4)$$

with

$$\lambda_{cW}^+(t) = \lambda_c(t)p_c(t) + \mathbb{E} \left[\lambda_{cI}^+(t - L_{cI}) \right] (1 - q_{cI}),$$

$$\lambda_{cI}^+(t) = \lambda_c(t)(1 - p_c(t)) + \mathbb{E} \left[\lambda_{cW}^+(t - L_{cW}) \right] (1 - q_{cW}),$$

see [18, Theorem 1.2]. In our network of ward and ICU with time-varying LoS and transfer probabilities, the expressions for $\rho_{cW}(t)$, $\rho_{cI}(t)$ are more involved. Moreover, the arrival rates, LoS distributions and transfer probabilities are not available in closed form, which prohibits evaluation of the expectations in (3), (4). Therefore, we do not provide an explicit expression for these means in the general case. In Section 3 we provide an algorithm that predicts the arrival rates and estimates the LoSs and transfer probabilities from the hospital's data warehouse. Subsequently, we use these system parameters to sample the patient trajectories of the Poisson Arrival Location Model resulting in a forecast of $\rho_{cW}(t)$, $\rho_{cI}(t)$ given the initial Poisson distribution of the number of patients as reflected by the initial condition (2) and the residual LoS of these patients.

The distribution of the total number of patients at the ward and ICU is now readily obtained. To this end, let $N_W(t)$ and $N_I(t)$ record the total number of patients in the ward and ICU, respectively, at time t . These random variables have a time-dependent Poisson distribution, for $n_W, n_I = 0, 1, 2, \dots$:

$$\mathbb{P}[N_W(t) = n_W, N_I(t) = n_I] = \frac{\rho_W(t)^{n_W}}{n_W!} \frac{\rho_I(t)^{n_I}}{n_I!} e^{-(\rho_W(t) + \rho_I(t))}, \quad (5)$$

with

$$\rho_W(t) = \sum_{c \in C} \rho_{cW}(t), \quad \rho_I(t) = \sum_{c \in C} \rho_{cI}(t).$$

Observe from (5) that at each time t the random variables $N_W(t)$, $N_I(t)$ are independent. However, for different time points, say t_1 and t_2 , the random variables $N_W(t_1)$, $N_I(t_2)$ are correlated, see [18, Theorem 2.2].

The Poisson distributions for the number of hospitalised patients (1), (5) allow us to evaluate various performance measures. Let $\mathbf{L}_W(s)$, $\mathbf{L}_I(s)$ be tuples of realised LoSs (up to time s) of patients residing in the ward and ICU, respectively. These tuples contain information on the number of patients in the ward and ICU, their patient characteristics c , and how long they have been in the ward and ICU at time s . We aim to forecast the occupancy at the ward and ICU at time $s + t$ given the LoSs of the residing patients at time s :

$$\begin{aligned} \mathbb{E}[N_W(s+t) \mid \mathbf{L}_W(s) = \boldsymbol{\ell}_W, \mathbf{L}_I(s) = \boldsymbol{\ell}_I], \\ \mathbb{E}[N_I(s+t) \mid \mathbf{L}_W(s) = \boldsymbol{\ell}_W, \mathbf{L}_I(s) = \boldsymbol{\ell}_I]. \end{aligned} \quad (6)$$

Furthermore, for a confidence level $\alpha \in [0, 1]$, we are interested in the quantiles

$$l_{\alpha W}(s+t), r_{\alpha W}(s+t), l_{\alpha I}(s+t), r_{\alpha I}(s+t) \quad (7)$$

such that

$$\begin{aligned} \mathbb{P}\left[N_W(s+t) \in [l_{\alpha W}(s+t), r_{\alpha W}(s+t)] \mid \mathbf{L}_W(s) = \boldsymbol{\ell}_W, \mathbf{L}_I(s) = \boldsymbol{\ell}_I\right] \geq \alpha, \\ \mathbb{P}\left[N_I(s+t) \in [l_{\alpha I}(s+t), r_{\alpha I}(s+t)] \mid \mathbf{L}_W(s) = \boldsymbol{\ell}_W, \mathbf{L}_I(s) = \boldsymbol{\ell}_I\right] \geq \alpha. \end{aligned} \quad (8)$$

These quantiles give conditional level α prediction intervals for the occupancy of the ward and ICU at time $s + t$, conditional on all LoS information of the currently

residing patients. A forecast of the expected maximum occupancy at the ward and ICU from time s up to time $s + t$ is:

$$\begin{aligned} & \mathbb{E} \left[\max_{u \in [s, s+t]} N_W(u) \mid \mathbf{L}_W(s) = \ell_W, \mathbf{L}_I(s) = \ell_I \right], \\ & \mathbb{E} \left[\max_{u \in [s, s+t]} N_I(u) \mid \mathbf{L}_W(s) = \ell_W, \mathbf{L}_I(s) = \ell_I \right]. \end{aligned} \quad (9)$$

For some confidence level $\alpha \in [0, 1]$, we are interested in the quantiles

$$l_{\alpha m W}(s+t), r_{\alpha m W}(s+t), l_{\alpha m I}(s+t), r_{\alpha m I}(s+t), \quad (10)$$

such that

$$\mathbb{P} \left[\max_{u \in [s, s+t]} N_W(u) \in [l_{\alpha m W}(s+t), r_{\alpha m W}(s+t)] \mid \mathbf{L}_W(s) = \ell_W, \mathbf{L}_I(s) = \ell_I \right] \geq \alpha, \quad (11)$$

$$\mathbb{P} \left[\max_{u \in [s, s+t]} N_I(u) \in [l_{\alpha m I}(s+t), r_{\alpha m I}(s+t)] \mid \mathbf{L}_W(s) = \ell_W, \mathbf{L}_I(s) = \ell_I \right] \geq \alpha. \quad (12)$$

These quantiles give conditional level α prediction intervals for the maximum occupancy of the ward and ICU during the interval $[s, s+t]$, that is of particular interest for the decision to accept new COVID-19 patients. Other performance measures, including the mixture of patients in the ward and ICU at each time t , may be obtained from (1).

3 On-line forecasting method

This section provides a procedure to predict the arrival rates (Section 3.1) and estimate the LoS distribution and transfer probabilities (Section 3.2) from the hospital's data warehouse. In Section 3.3 we use these system parameters to sample the patient trajectories of the Poisson Arrival Location Model resulting in forecasts of the daily occupancy and the maximum occupancy, including their prediction intervals.

3.1 Richards' curve to predict the arrival rate

Our forecasting method requires $\lambda(t)$, the expected number of arrivals of patients to the hospital at time t . In accordance with literature [15] the cumulative rate follows a 5-parameter Richards' curve:

$$\Lambda(t) = \int_{-\infty}^t \lambda(s) ds = \frac{R - L}{[1 + \delta \exp(-k(t - t_0))]^{1/\delta}} + L, \quad (13)$$

where all parameters R, δ, k, t_0 are positive, resulting in an S-shaped growth curve for the expected number of arrivals up to time t . The parameters have the following interpretations: R represents the total number of arrivals (indeed $\lim_{t \rightarrow \infty} \Lambda(t) = R$), k is a scale parameter related to the growth rate, t_0 determines the offset

along the t -axis, δ is a shape parameter that introduces asymmetry and L is a left asymptote of $\Lambda(t)$. If $\delta = 1$, the resulting growth curve is the logistic growth curve that describes the fraction of infected people in a Susceptible-Infected-Susceptible (SIS) compartmental model [9].

The Richards' curve is fitted on the cumulative number of arrivals using the R package `FlexParamCurve` version 1.5-5 [22] via the following procedure. Our goal is to predict the occupancy of the ward and ICU on a daily basis. To this end, the number of arrivals of COVID-19 patients with characteristics c is determined for each day from the hospital's data warehouse. The parameters of the Richards' curve are estimated minimising the sum of squared errors between the cumulative number of arrivals on each day and the expression in the right-hand side of (13) at that day using the Gauss-Newton algorithm, similar to [30] (that uses a different algorithm to minimise the sum of squared errors). For some sets of arrival data, the nonlinear least squares method fails to converge due to over-parameterisation. In this case, the procedure first fixes the parameter L to 0 (leading to a 4-parameter Richards' curve). If this also does not lead to convergence, the parameter δ is fixed to 1, in which case the procedure fits a logistic growth curve on the arrival data. This procedure results in an estimate $\hat{\Lambda}(d)$ for the expected cumulative number of arrived COVID-19 patients on any given day d . This daily estimate is then linearly interpolated to generate a cumulative arrival intensity $\hat{\Lambda}(t)$ which can be evaluated at each time point t . Let \hat{p} denote the empirically estimated probability $p_c(t)$ that a patient with characteristics c is admitted to the ward at time t , which is assumed stationary and equal for all patient types. Let f_c be the empirically estimated fraction of patients with characteristics c that arrive directly to the hospital. The cumulative arrival rate $A_c(t)$ of patients with characteristics c arriving directly to the hospital is then estimated as $A_c(t) = f_c \Lambda(t)$ for all t , and the cumulative arrival rates to the ward and ICU are estimated as $f_c \hat{p} \Lambda(t)$ and $f_c (1 - \hat{p}) \Lambda(t)$ for all t .

3.2 Kaplan-Meier estimation of the LoS distribution and transfer probabilities

Our forecasting method requires the LoS distribution F_{cW} at the ward and F_{cI} at the ICU for all patient characteristics c . We use the Kaplan-Meier estimator [13] for the survival function that takes right-censored observations into account, which occur when a patient is still at the respective department or when a patient is transferred to another hospital. Our goal is to predict the occupancy at the ward and ICU on a daily basis. The estimated LoS distribution gives the probability that a patient is at the department at most a certain number of days. Let $e_{cW}(u)$, resp. $e_{cI}(u)$, denote the number of patients with characteristics c at the ward, resp. ICU, with a realised LoS equal to u . Let $n_{cW}(u)$, $n_{cI}(u)$ denote the number of patients with characteristics c that are still being treated at the ward, ICU after $u - 1$ days. The Kaplan-Meier estimates for the LoS distribution at the ward and ICU are then given by [13]:

$$\hat{F}_{cW}(\ell) = 1 - \prod_{u=1}^{\ell} \left(1 - \frac{e_{cW}(u)}{n_{cW}(u)} \right), \quad \hat{F}_{cI}(\ell) = 1 - \prod_{u=1}^{\ell} \left(1 - \frac{e_{cI}(u)}{n_{cI}(u)} \right).$$

Our method requires the probabilities $q_{cW}(\ell)$ ($q_{cI}(\ell)$) for a patient with characteristics c transferring from the ward (ICU) to the ICU (ward) after completing a LoS of ℓ days at the ward (ICU) that are estimated as follows. Let $\hat{F}_{cIW}(\ell)$ denote the empirical probability that a patient with characteristics c is transferred from the ICU to the ward after a LoS of ℓ days at the ICU, and let $\hat{F}_{cWI}(\ell)$ be defined similarly for a patient transferring from ward to ICU. The probabilities $q_{cW}(\ell)$ and $q_{cI}(\ell)$ are estimated as

$$\hat{q}_{cW}(\ell) = \frac{\hat{F}_{cWI}(\ell) - \hat{F}_{cWI}(\ell - 1)}{\hat{F}_{cW}(\ell) - \hat{F}_{cW}(\ell - 1)}, \quad \hat{q}_{cI}(\ell) = \frac{\hat{F}_{cIW}(\ell) - \hat{F}_{cIW}(\ell - 1)}{\hat{F}_{cI}(\ell) - \hat{F}_{cI}(\ell - 1)}.$$

3.3 Generation of the PALM and forecasting ward and ICU bed occupancy

This section presents our method to sample the patient trajectories of the Poisson Arrival Location Model (PALM) resulting in (for instance) forecasts of the conditional means shown in (6). Our method simulates the PALM using Monte Carlo sampling of arrivals, transfers and departures of patients over the forecasting period $[s, s + t]$. We restrict the trajectories to:

1. patients admitted to the ward that leave the hospital from the ward,
2. patients admitted to the ICU that leave the hospital from the ICU,
3. patients admitted to the ICU, then transferred to the ward and leave the hospital from the ward,
4. patients admitted to the ward, then transferred to the ICU and leave the hospital from the ICU.

As a consequence, in our simulation method a patient may visit at most two departments. This restriction is introduced to reduce computational complexity as it avoids a large number of possible patient trajectories. The restriction has a minor effect on our results, as data shows that multiple transfers are very rare during our forecasting horizon of one week. In the description of our method below, we identify the patient characteristics c with the trajectories. Note that randomly assigning patients to these trajectories is equivalent to random selection of transfer or discharge/death upon completion of the LoS at a department, see, e.g., [14, p. 64].

All parameters required for our sampling method are obtained from the hospital's data warehouse as presented in Sections 3.1 and 3.2.

For each time u in the forecasting period $[s, s + t]$, arrivals of new patients with characteristics c are sampled according to a non-homogeneous Poisson process with cumulative rate equal to $f_c \hat{p} \hat{\Lambda}(u)$ for the ward and $f_c (1 - \hat{p}) \hat{\Lambda}(u)$ for the ICU, where the cumulative total arrival rate $\hat{\Lambda}(u)$ is extrapolated from the Richards' curve based on the hospital's data up to the start of the forecasting period at time s . The sampling procedure is executed until the next arrival time exceeds $s + t$. Sampling of inter-arrival times is executed by inverse transformation sampling based on the cumulative arrival rates (see, e.g., [23, p. 312]).

Departure times are generated upon the arrival/transfer of patients and is done by inverse sampling with replacement under the estimated empirical LoS distributions for that combination of patient characteristics and department. These LoS distributions are estimated on all LoS data obtained before the start of the forecasting period at time s . Note that for this specification of patient characteristics, the probability \hat{q}_{cI} (resp. \hat{q}_{cW}) of transfer from the ICU to the ward (resp. ward to ICU) is either always equal to zero or equal to one, depending on whether the patient is of type 3 that is admitted to the ICU and then transfers to the ward (resp. type 4). We do not distinguish the LoS distributions for patients of type 2 and 4 at the ICU, i.e., we assume that $F_{2I} = F_{4I}$. Similarly, we assume that $F_{1W} = F_{3W}$.

For patients residing at the hospital at time s , the patient type is first sampled based on a patient's current realised LoS ℓ . Conditional sampling is based on the ratio between Kaplan-Meier Survival function estimates of general patients residing in the current department and the Kaplan-Meier Survival function estimate of patients being transferred to the other department. For instance, for a patient currently residing at the ICU with current LoS equal to ℓ , the probability that this patient is assigned type 3 is equal to

$$\frac{1 - \hat{F}_{IW}(\ell)}{1 - \hat{F}_I(\ell)},$$

where \hat{F}_I is a Kaplan-Meier estimate of the LoS distribution at the ICU for a general patient, i.e., estimated on all LoS data at the ICU.

The LoS for these currently residing patients is then sampled from the empirical conditional LoS distribution for the sampled patient's type, where the conditioning is done on the already realised LoS. For example, if the patient is assigned type 2 and has a current LoS of ℓ_2 days at the ICU, the total LoS of this patient at the ICU is sampled from the cumulative distribution:

$$\hat{F}_{2I}(\ell_1|\ell_2) = \frac{\hat{F}_{2I}(\ell_1) - \hat{F}_{2I}(\ell_2 - 1)}{1 - \hat{F}_{2I}(\ell_2 - 1)} \quad \forall \ell_1.$$

In order to keep track of the PALM, $N_I(s)$ and $N_W(s)$ are first determined. Then, at each simulated arrival, transfer or departure before time $s + t$, the counters are updated according to the nature of the event. This results in a registered occupancy for both departments at each time point in the forecasting period $[s, s + t]$. The occupancy at a given day at a department is now calculated as the number of patients at the department at the start of that day.

The simulation procedure is repeated 1,000 times in order to accurately estimate the statistics mentioned at the end of Section 2. In order to estimate the expected values at day $s + t$ in (6), the average of the number of patients on day $s + t$ at both departments is taken over all simulations runs. Next, in order to estimate the boundaries in (7) for a given day $s + t$, the respective empirical quantiles are taken over the simulated occupancy on that day for both departments. Furthermore, in order to determine the expressions in (9) and (10), the maximum occupancy is determined for both departments from the forecast date (time s) until the end of the forecasting horizon (time $s + t$). Then, to determine (9) and (10), the empirical mean and quantiles are determined.

4 COVID-19 bed occupancy: case studies and evaluation

This section presents the results of our forecasting method detailed in Section 3 for four Dutch hospitals. Section 4.1 presents results for the Leiden University Medical Centre (LUMC), a medium-sized academic hospital. Subsequently, we present results for the larger general hospitals HagaZiekenhuis (Section 4.2), Rijnstate (Section 4.3), and Elisabeth-TweeSteden Ziekenhuis (Section 4.4). In Section 4.5, we compare these forecasts, elaborate on the quality of our forecasts and provide a statistical evaluation of our forecasting method.

Patient	Origin	Destination	datetime start	datetime end	stay at ICU
1	Home	Care facility (other)	5-3-2020 22:06	3-4-2020 11:13	no
2	Home	Care facility (rehabilitation)	15-3-2020 16:38	6-4-2020 11:42	no
3	Home	Home	15-3-2020 16:21	16-3-2020 18:39	no
4	Home	ICU	15-3-2020 20:19	18-3-2020 10:43	no
4	Ward	Deceased (without autopsy)	18-3-2020 10:43	6-4-2020 23:03	yes
5	Home	Ward	15-3-2020 22:15	16-3-2020 13:51	yes
5	ICU	Home	16-3-2020 13:51	18-3-2020 14:49	no
6	Home	Care facility (other)	16-3-2020 18:58	27-3-2020 11:11	no
7	Home	ICU	17-3-2020 02:00	24-3-2020 10:23	no
7	Ward	Other hospital	24-3-2020 10:23	26-3-2020 03:27	yes
8	Home	Home	18-3-2020 01:15	22-3-2020 14:00	no
9	Home	Deceased (without autopsy)	18-3-2020 03:10	5-4-2020 03:00	yes
10	Home	Home	18-3-2020 21:52	19-3-2020 21:08	no
11	Home	ICU	19-3-2020 00:16	30-3-2020 00:08	no
11	Ward	Deceased (without autopsy)	30-3-2020 00:08	7-4-2020 12:25	yes
12	Home	Home	19-3-2020 11:54	27-3-2020 21:43	no

Fig. 2: Extract of the input data for our model, where the name of the hospital, patient identification number, as well as patient characteristics such as age, BMI and gender have been removed for privacy reasons. Care facility denotes a long term care facility outside the hospital.

Our forecasting method requires a complete record of time stamps for patient admissions, transfers and discharges that we obtained from all hospitals included in this paper. Figure 2 shows an extract of the input table, where the name of the hospital, patient identification number, as well as patient characteristics have been removed for privacy reasons. The rows of this table describe the trajectory of the patient at either the ward or the ICU. The patient ID is replaced by a number in the first column. The next two columns describe the Origin and Destination of the patient before and after his/her stay at the current department, while the last column indicates whether the current department is ICU (yes) or ward (no). The fourth and fifth columns give information about the start and end time of the patient's stay at the current department. In this table, the trajectory of the patient is explicitly characterised. For example, patient nr. 5 is admitted to the ICU on March 15 at 22:15 from his/her own home environment, transferred to the ward on March 16 at 13:51 and discharged from the ward to return to his/her home environment on March 18 at 14:49.

Using the method described in Section 3.3, the PALM is simulated 1,000 times. In the figures in this section, performance measures of the true occupancy (at the start of each day, depicted in red) are compared with these measures forecast by our

method. The graphs in this section report estimates of the following performance measures:

- Expanding window forecasts of the bed occupancy at the ward and ICU 1, 2, 3 and 5 days ahead:

$$\mathbb{E}[N_W(s+t) \mid \mathbf{L}_W(s) = \ell_W, \mathbf{L}_I(s) = \ell_I], \quad (14)$$

$$\mathbb{E}[N_I(s+t) \mid \mathbf{L}_W(s) = \ell_W, \mathbf{L}_I(s) = \ell_I], \quad (15)$$

where s is the current day and $t = 1, 2, 3$, and 5 the forecasting horizon for which the bed occupancy is forecast. These expected values are estimated by taking the average of the occupancy at both departments for all 1,000 PALM simulations at $t = 1, 2, 3$ and 5 days. The forecasts are obtained using an expanding window procedure. This means that for each day, the forecasts are generated using only the data (e.g. arrivals, transfers and departures) up to that day. After predicting the daily occupancy at a given date, the date is incremented by one day, i.e., $s \rightarrow s + 1$ and the forecasts are generated using all data up to that new date.

The graphs show, for a certain period, the forecasts of the occupancy at 1, 2, 3 and 5 days ahead, as well as the realised daily occupancy.

- The corresponding 95% prediction interval 3 days ahead based on (8). These quantiles are estimated by taking the empirical quantiles of level $1/2 - \alpha/2$ and $1/2 + \alpha/2$ of the occupancy at both departments for all 1,000 PALM simulations at $t = 3$ days ahead for each day s in the forecasting period.
- Expanding window forecasts of the maximum expected occupancy at the ward and ICU for forecasting horizon t

$$\mathbb{E} \left[\max_{u \in [s, s+t]} N_W(u) \mid \mathbf{L}_W(s) = \ell_W, \mathbf{L}_I(s) = \ell_I \right], \quad (16)$$

$$\mathbb{E} \left[\max_{u \in [s, s+t]} N_I(u) \mid \mathbf{L}_W(s) = \ell_W, \mathbf{L}_I(s) = \ell_I \right]. \quad (17)$$

These expected values are estimated by taking the average of the maximum occupancy realised for all 1,000 PALM simulations in the interval $[s, s+t]$.

- The corresponding 95% prediction interval 3 days ahead based on (11), (12). These quantiles are estimated by taking the empirical quantiles at levels $1/2 - \alpha/2$ and $1/2 + \alpha/2$ of the maximum occupancy, at both departments for all 1,000 PALM simulations in the interval $[s, s+t]$ for $t = 3$ for each day s in the forecasting period.

The expected daily bed occupancy, along with its quantiles, can be used to visualise the expected evolution of the occupancy at both COVID-19 ward and ICU.

The maximum expected occupancy in a certain forecasting period of $[s, s+t]$ days expresses the risk of overcrowding of the ward and ICU in the next t days. These forecasts can be used to evaluate the expected minimally required number of available beds in the next days. The maximum expected occupancy is an important performance measure to control admittance of COVID-19 patients to a hospital at both ward and ICU.

In the following sections, six figures are shown for each hospital. The left-hand figures consider occupancy at the ICU and the right-hand figures consider occupancy at the ward. For all figures, the patients transferred from other hospitals are excluded from the calculation of the occupancy. This is because the predicted arrival rate only pertains to autonomous arrivals to the hospital. In the figures in the top row, expanding window forecasts of the daily occupancy are shown for each day s . Furthermore, for each day s , the forecast at day $s - t$ of the occupancy at day s is shown for forecast horizons t . The realised occupancy at those dates is shown in red. In the top row, the forecasts are shown without prediction intervals. For the LUMC, forecasts for horizons $t = 1, 2, 3$ and 5 days are shown. For the other hospitals, forecasts for the horizons $t = 1$ and 3 are shown only, to enhance the clarity of the figures. The colours for forecast horizon $t = 1, 2, 3$ and 5 are orange, cyan, blue and purple, respectively. The figures in the middle row display the expanding window forecasts for daily occupancy at forecast horizon $t = 3$ including the 95% prediction interval. The figures in the bottom row show the expanding window forecasts of the maximum expected occupancy for a forecast horizon of $t = 3$ days.

4.1 Leiden University Medical Centre

Leiden University Medical Centre (LUMC) is an academic hospital in Leiden. Together with the other general hospitals in the region, it serves a community of around two million people in an urban area in the south-west of the Netherlands. The main focus of the LUMC is top clinical and highly specialised care. It is the smallest and oldest of the eight academic hospitals in the Netherlands.

Figure 3 presents our forecasts for the LUMC for the period April 15, 2020 until July 10, 2020, the second part of the first COVID-19 peak in this region. The first COVID-19 patient arrived at the LUMC on March 3, 2020. Hence at the start of the forecast interval, 1.5 months of data is available on the arrival rates and LoSs of COVID-19 patients, including 99 (30) COVID-19 patients that had left the ward (ICU) before then.

4.2 HagaZiekenhuis

HagaZiekenhuis (Haga) is a top clinical hospital in The Hague with 600 beds and approximately 29,000 inpatient admissions per year. Next to secondary care, top clinical hospitals in the Netherlands also provide tertiary care for particular patient groups that the hospital specialises in. Moreover, top clinical hospitals play an important role in the education of medical professionals, and perform clinical research. Haga is located in the same urban area in the south-west of the Netherlands as the LUMC.

Figure 4 presents our forecasts for Haga for the period March 30, 2020 until July 10, 2020, the second part of the first COVID-19 peak in this region. The first COVID-19 patient arrived at Haga on March 5, 2020. Hence at the start of the forecast interval, 25 days of data is available on the arrival rates and LoSs of

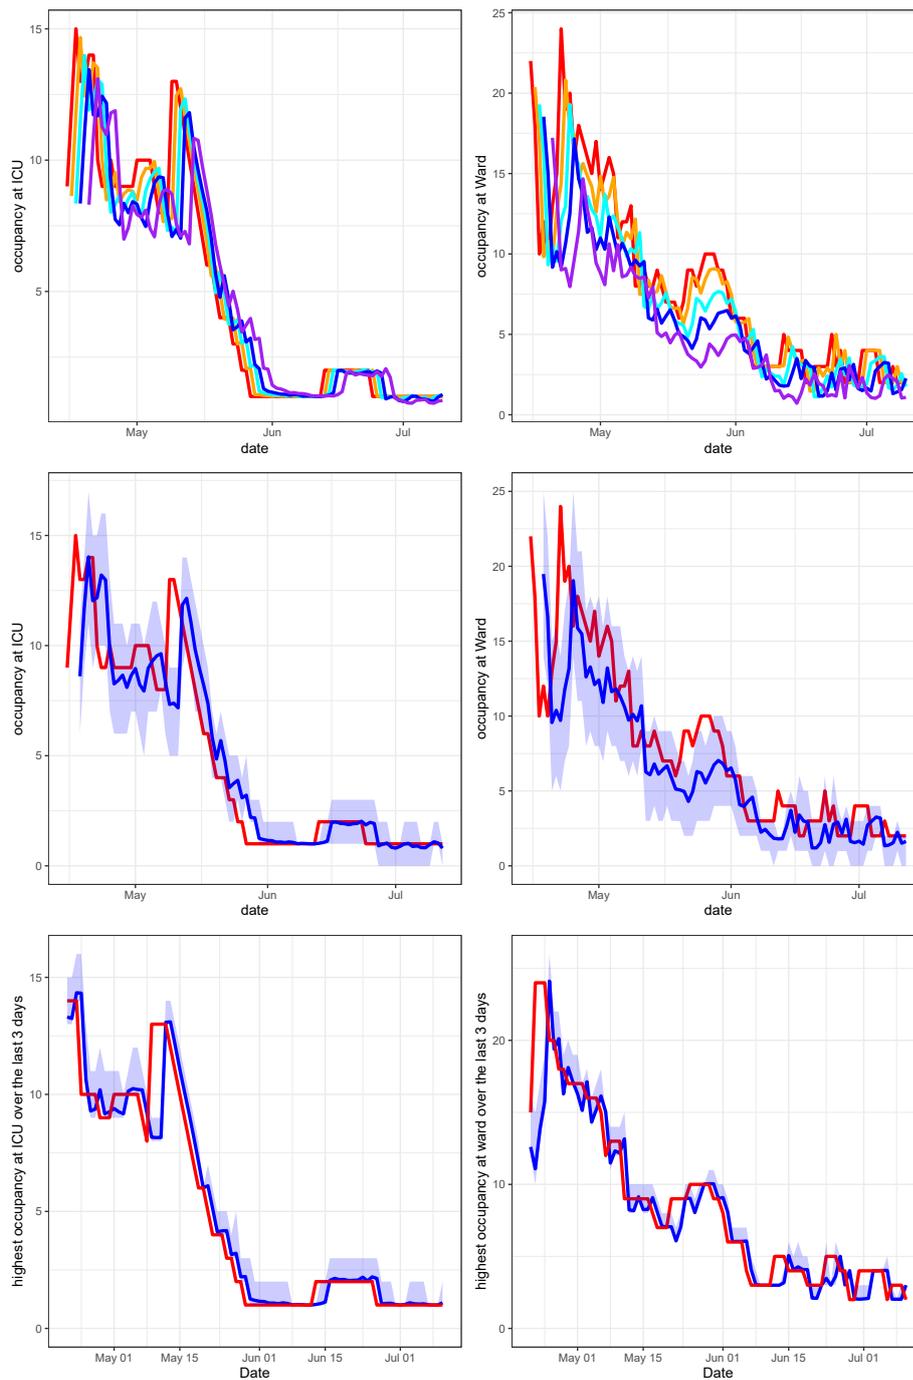


Fig. 3: Leiden University Medical Centre April 15, 2020 until July 10, 2020.

Top row: expanding window forecasts 1, 2, 3 and 5 days ahead at the COVID-19 ICU (left) and ward (right).

Middle row: expanding window forecasts 3 days ahead at the COVID-19 ICU (left) and ward (right), along with a 95% prediction interval.

Bottom row: expanding window forecasts of the maximum occupancy, including the 95% prediction interval and realised maximum occupancy of patients at the COVID-19 ICU (left) and ward (right) over the last 3 days.

The realised occupancy is shown in red, while the forecasts for 1, 2, 3 and 5 days ahead are shown in orange, cyan, blue and purple respectively.

COVID-19 patients, including 40 (5) COVID-19 patients that had left the ward (ICU) before then.

4.3 Rijnstate

Rijnstate is a top clinical hospital in Arnhem with 766 beds and approximately 33,000 inpatient admissions per year. Rijnstate serves a community of 450 thousand people in and around Arnhem, a city in the east of the Netherlands.

Figure 5 presents our forecasts for Rijnstate for the period March 30, 2020 until July 10, 2020, the second part of the first COVID-19 peak in this region. The first COVID-19 patient arrived at Rijnstate on March 3, 2020. Hence at the start of the forecast interval, approximately one month of data is available on the arrival rates and LoSs of COVID-19 patients, including 157 (7) COVID-19 patients that had left the ward (ICU) before then.

4.4 Elisabeth-TweeSteden Ziekenhuis

Elisabeth-TweeSteden Ziekenhuis (ETZ) is a top clinical hospital in Tilburg with 782 beds and approximately 37,000 inpatient admissions per year. Tilburg is a city in North-Brabant, the province that experienced the initial outbreak of COVID-19 in the Netherlands at the end of February, 2020. As a consequence, ETZ was the first Dutch hospital to admit a COVID-19 patient. Just like HagaZiekenhuis and Rijnstate, ETZ is a top clinical hospital.

Figure 6 presents our forecasts for ETZ for the period March 23, 2020 until July 10, 2020, the second part of the first COVID-19 peak in this region. The first COVID-19 patient arrived at ETZ on February 28, 2020. Hence at the start of the forecast interval, almost one month of data is available on the arrival rates and LoSs of COVID-19 patients, including 104 (24) COVID-19 patients that had left the ward (ICU) before then.

4.5 Evaluation of our method

In this section, we discuss the results presented in Sections 4.1 through 4.4. Moreover, we compare the performance of our forecasting method to the performance of a moving average forecaster for all four hospitals.

When investigating trends in the COVID-19 occupancy during the first peak, we see a very similar trend for the LUMC (Figure 3) and Haga (Figure 4), which is natural as these two hospital are located in the same region. ETZ (Figure 6) has the earliest and highest peak and also admitted the highest total number of COVID-19 patients, because ETZ is located in North-Brabant, the region that was hit first and hardest during the first COVID-19 peak in the Netherlands. In Rijnstate (Figure 5) the first COVID-19 peak ended the latest.

Naturally, the accuracy of forecasts increases when forecasts are made at a point in time that is closer to the actual realisation. This can clearly be seen in the top

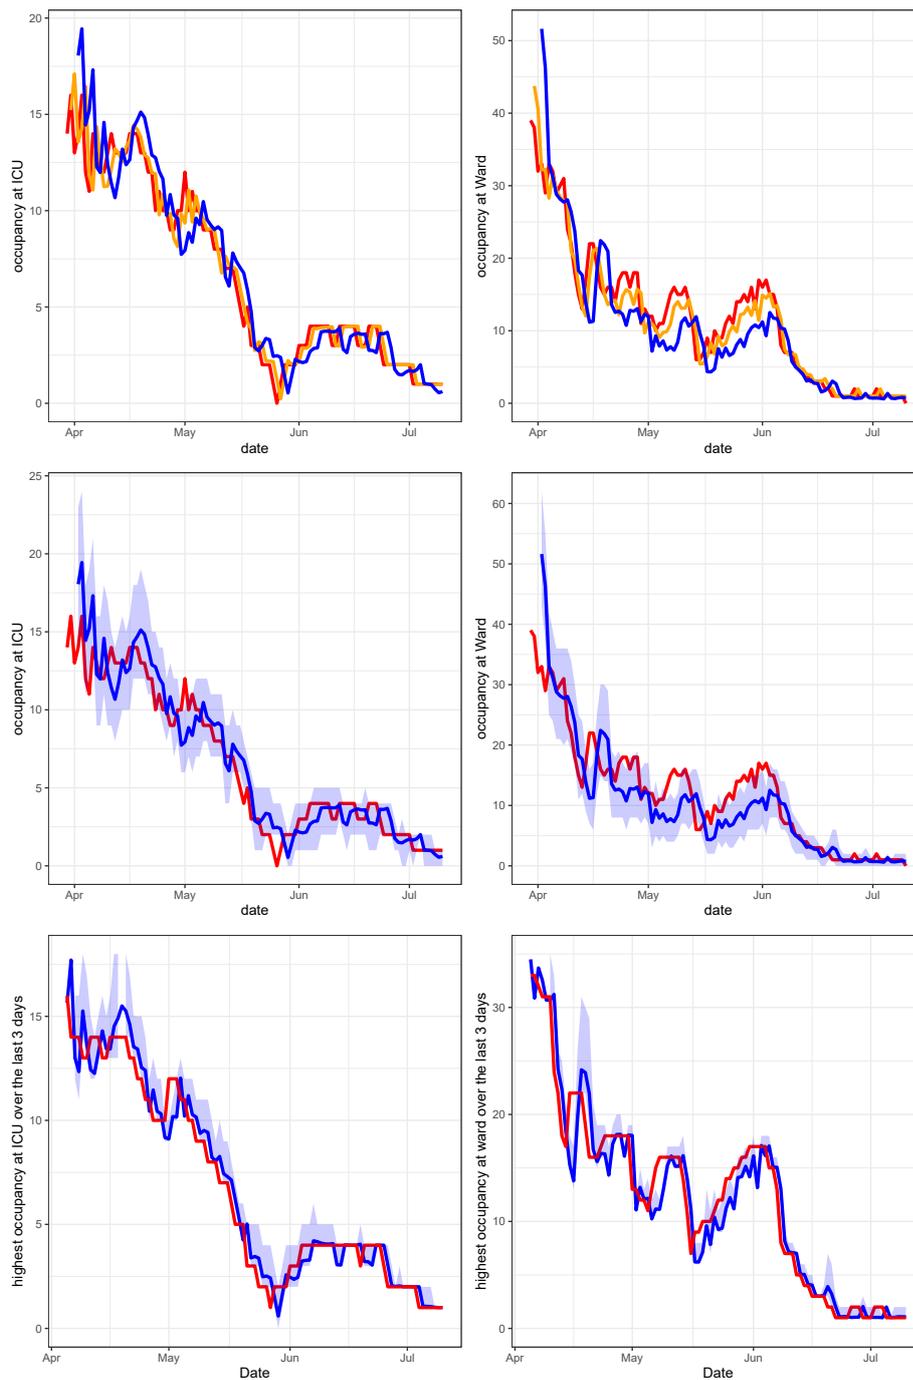


Fig. 4: HagaZiekenhuis March 30, 2020 until July 10, 2020.

Top row: expanding window forecasts 1 and 3 days ahead at the COVID-19 ICU (left) and ward (right).

Middle row: expanding window forecasts 3 days ahead at the COVID-19 ICU (left) and ward (right), along with a 95% prediction interval.

Bottom row: expanding window forecasts of the maximum occupancy, including the 95% prediction interval and realised maximum occupancy of patients at the COVID-19 ICU (left) and ward (right) over the last 3 days.

The realised occupancy is shown in red, while the forecasts for 1 and 3 days ahead are shown in orange and blue respectively.

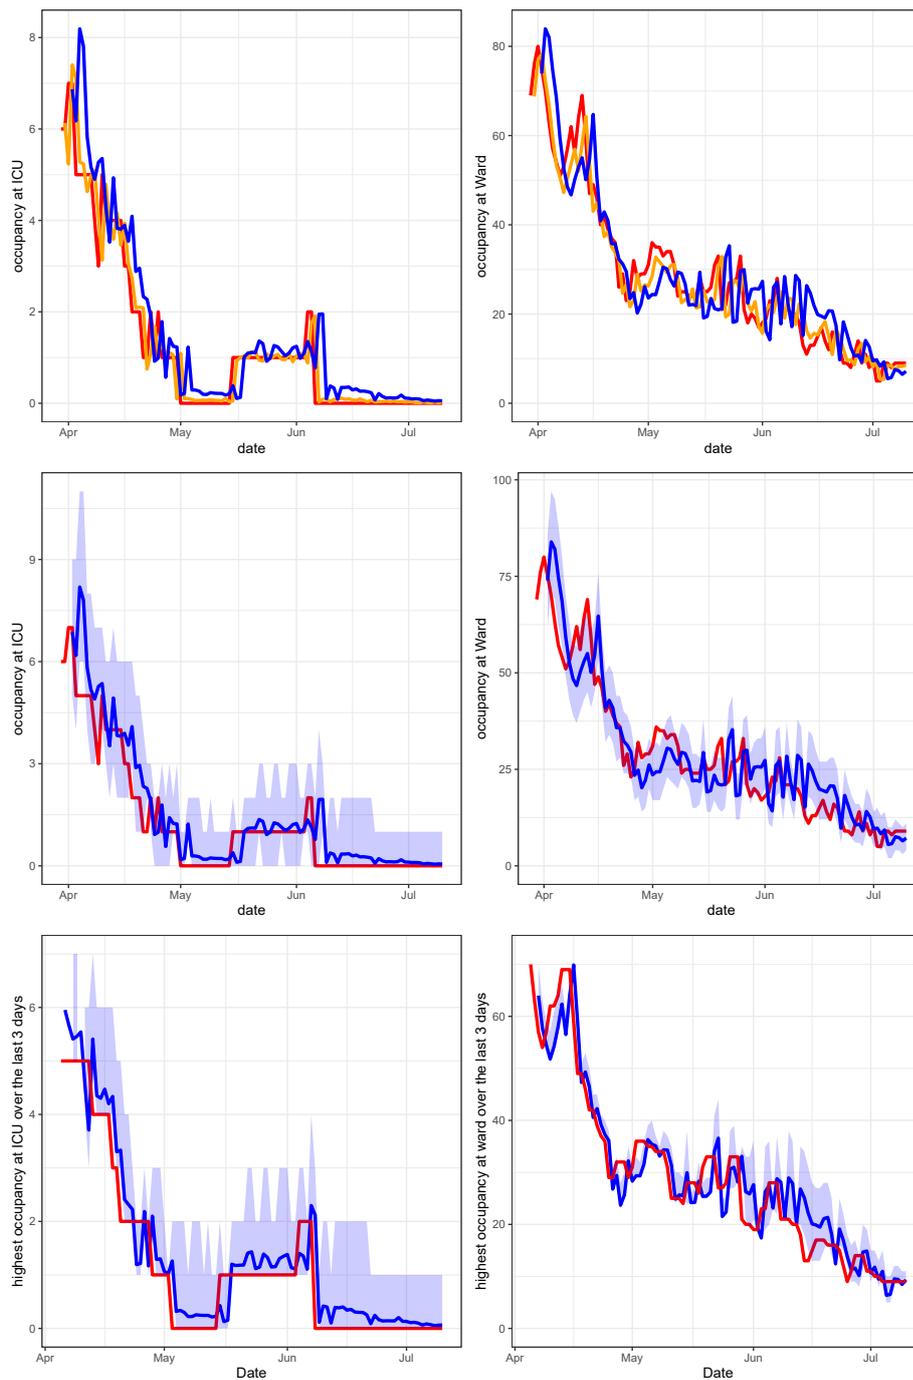


Fig. 5: Rijnstate March 30, 2020 until July 10, 2020.

Top row: expanding window forecasts 1 and 3 days ahead at the COVID-19 ICU (left) and ward (right).

Middle row: expanding window forecasts 3 days ahead at the COVID-19 ICU (left) and ward (right), along with a 95% prediction interval.

Bottom row: expanding window forecasts of the maximum occupancy, including the 95% prediction interval and realised maximum occupancy of patients at the COVID-19 ICU (left) and ward (right) over the last 3 days.

The realised occupancy is shown in red, while the forecasts for 1 and 3 days ahead are shown in orange and blue respectively.

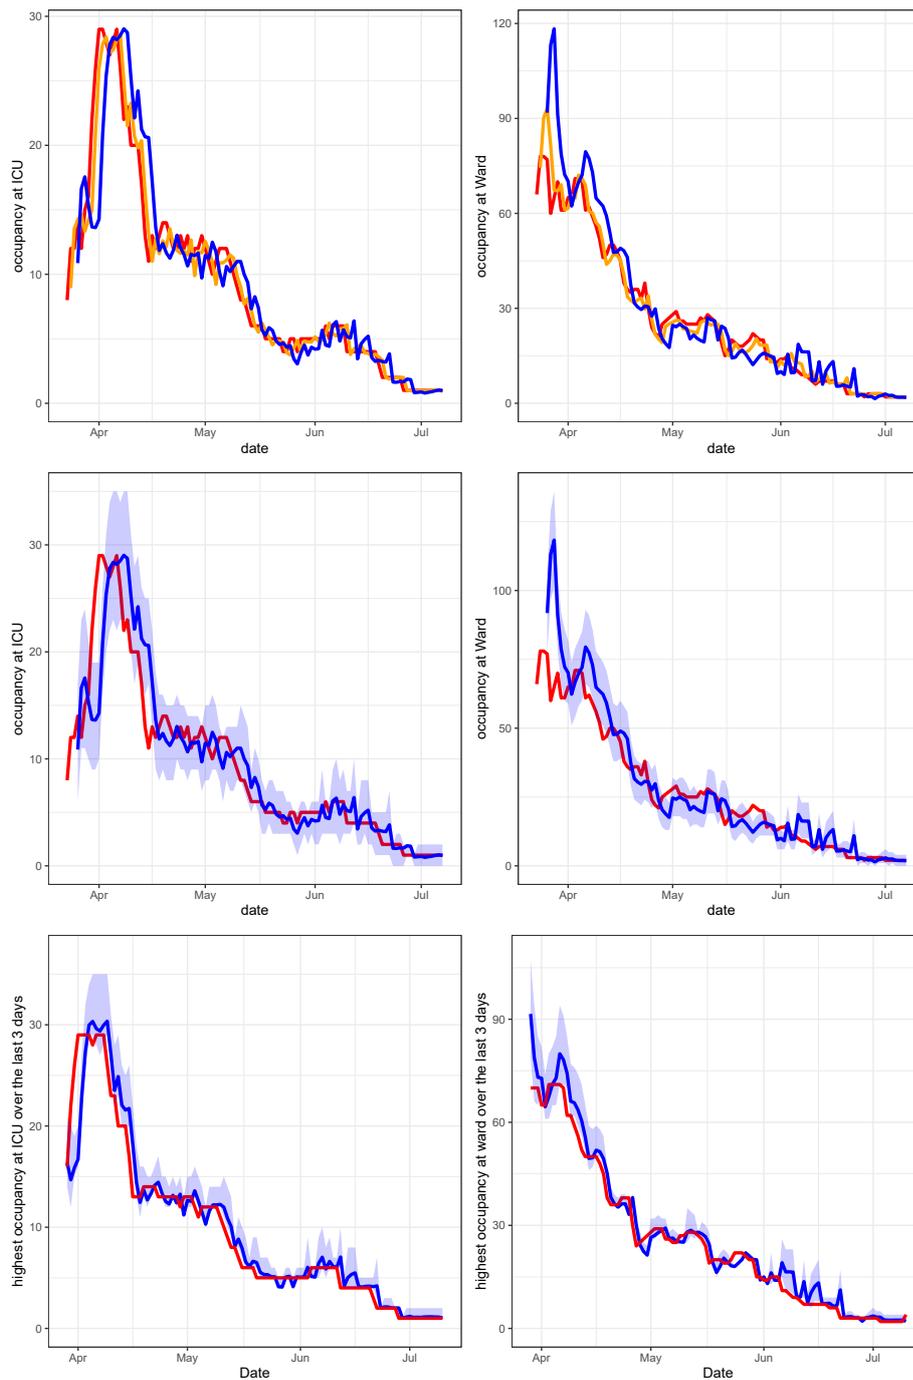


Fig. 6: Elisabeth-TweeSteden Ziekenhuis March 23, 2020 until July 10, 2020.

Top row: expanding window forecasts 1 and 3 days ahead at the COVID-19 ICU (left) and ward (right).

Middle row: expanding window forecasts 3 days ahead at the COVID-19 ICU (left) and ward (right), along with a 95% prediction interval.

Bottom row: expanding window forecasts of the maximum occupancy, including the 95% prediction interval and realised maximum occupancy of patients at the COVID-19 ICU (left) and ward (right) over the last 3 days.

The realised occupancy is shown in red, while the forecasts for 1 and 3 days ahead are shown in orange and blue respectively.

rows of Figures 3 through 6. In particular, the top row of Figure 3 shows that our expanding window forecasts 1 day ahead are very accurate, while the accuracy reduces for forecasts 5 days ahead. One of the reasons is that the 1-day forecast is able to pick up sudden changes in the trend the next day, whereas this obviously takes five days for the 5-day forecast. This is, for example, visible in the top-right graph of Figure 3 when, half May, the downward trend changes to a sudden peak. Similar effects are visible in the same graph at the end of May and also in Figures 4 through 6 with a sudden decline in the number of hospitalised patients. This delay in picking up sudden changes in the trend also results in larger over- or undershoots for the forecasts further into the future.

The top and middle rows of the figures also show clearly that the accuracy of forecasts increases for larger population sizes. ETZ has seen the highest number of COVID-19 patients on its ward and ICU during the first peak, and indeed the middle row of Figure 6 displays narrow confidence intervals that often contain the realisation. LUMC (Figure 3) and Rijnstate (Figure 5) saw the smallest number of COVID-19 ward and ICU patients, respectively, resulting in broader confidence intervals that contain the realisation less often.

In contrast to the forecasts of the daily occupancy that we just discussed, the forecast of the maximum occupancy 3 days ahead of time is spot-on, as displayed in the bottom rows of Figures 3 through 6. This is exactly the forecast that is most valuable for hospitals, as it provides quantitative support for several decisions, for example on the admittance of additional COVID-19 patients, the necessity of COVID-19 patient transfers to other hospitals, and the (im)possibility of providing care for non-COVID-19 patients.

Table 1 displays error measures for our forecasting method and also for a moving average forecaster, which enables us to compare the results. For the moving average forecaster, the average occupancy of the last week is used as forecast of the daily occupancy for several horizons t . Error measures for the expected maximum bed occupancy are not displayed, as it is not straightforward to forecast the maximum occupancy using a moving average. The stochastic nature of our own forecasting method results in some noise in the performance measures, which we have reduced by averaging the performance measures over five independent replications for each of the hospitals. The columns ‘CR’ (coverage rate) indicate how often the realised bed occupancy was covered by the 95% prediction interval. For the ICU, this happened in 74-97% of cases. The lower coverage rate at the ward, resulting from narrow prediction intervals, is most likely due to the fact that we do not incorporate the uncertainty about the estimated LoS distributions and predicted arrival rates in our prediction intervals. Most of the occupancy in the ward is due to direct arrivals to the hospital, while the uncertainty was seen to be highest for the arrival rate predictions. Observe from Figures 3 through 6 that when the prediction interval does not cover the realisation, it is mostly very close to the realisation. The bias is calculated as the average of ‘forecast minus realised’. Hence, the closer to zero the better. For the ICU, the bias of our forecast is always close to zero and lower than that of the moving average forecaster. For the ward, the negative bias indicates that our forecast is slightly too low on average. For the wards at ETZ and Rijnstate, that admitted the highest numbers of patients, the bias of our forecast is much lower than that of the moving average forecaster. The

Table 1: Results for our forecasting method (average of 5 simulation runs of 1,000 replications per hospital) compared to a moving average forecaster for the occupancy at the departments of the hospitals.

CR: coverage rate of the occupancy by the 95% prediction interval, bias: the bias estimated by averaging errors, MAE: mean absolute error.

	Forecast						Mov. av. forecaster			
	ICU			Ward			ICU		Ward	
	CR	bias	MAE	CR	bias	MAE	bias	MAE	bias	MAE
LUMC										
1 day ah.	0.94	0.04	0.45	0.47	-0.52	1.16	0.62	1.06	0.55	1.62
2 days ah.	0.91	0.01	0.77	0.54	-1.31	1.81	0.75	1.24	0.83	1.79
3 days ah.	0.91	0.00	0.92	0.54	-1.89	2.28	0.86	1.40	1.05	1.99
5 days ah.	0.86	0.03	1.37	0.55	-2.77	2.99	1.10	1.67	1.47	2.27
Max. 3d. ah.	0.74	0.07	0.62	0.17	-0.55	1.43	-	-	-	-
Haga										
1 day ah.	0.89	0.06	0.60	0.59	-0.83	1.60	0.46	1.04	1.32	2.59
2 days ah.	0.90	0.02	0.90	0.65	-1.71	2.57	0.59	1.20	1.65	2.98
3 days ah.	0.92	0.00	1.14	0.60	-2.42	3.32	0.72	1.33	1.99	3.35
5 days ah.	0.89	0.25	1.58	0.55	-2.86	4.45	1.00	1.58	2.64	4.19
Max. 3d. ah.	0.76	0.16	0.74	0.19	-0.19	1.69	-	-	-	-
Rijnstate										
1 day ah.	0.96	0.04	0.33	0.76	-0.61	3.54	0.23	0.37	2.44	4.65
2 days ah.	0.95	0.14	0.50	0.67	-1.19	5.18	0.29	0.44	2.95	5.19
3 days ah.	0.97	0.25	0.58	0.66	-1.72	6.21	0.34	0.52	3.44	5.60
5 days ah.	0.96	0.48	0.79	0.61	-1.62	7.61	0.43	0.65	4.54	6.31
Max. 3d. ah.	0.93	0.24	0.41	0.29	0.72	4.22	-	-	-	-
ETZ										
1 day ah.	0.97	-0.11	0.92	0.85	-0.40	2.00	0.52	1.89	2.75	3.57
2 days ah.	0.95	-0.34	1.42	0.78	-0.68	3.45	0.74	2.21	3.38	4.19
3 days ah.	0.92	-0.57	1.91	0.69	-0.71	4.92	1.01	2.50	4.02	4.75
5 days ah.	0.89	-0.59	2.60	0.59	0.30	7.00	1.63	2.93	5.43	5.76
Max. 3d. ah.	0.76	0.21	1.21	0.22	1.71	2.80	-	-	-	-

mean absolute error (MAE) of our forecast is again close to zero for the ICU, and lower than that of the moving average forecaster for all hospitals except Rijnstate (the hospital with the smallest number of COVID-19 ICU admissions). For the ward, the MAE of our forecast is lower than that of the moving average forecaster up to a horizon of 2 days. In conclusion, our forecasting method often outperforms the moving average forecaster. Moreover, our forecasting method is richer as it also produces prediction intervals and generates forecasts for the maximum occupancy that are spot-on.

To summarise, our forecasting method shows to be very accurate, which has convinced hospitals to embrace our forecasting method and incorporate it in their COVID-19 control or capacity dashboard (e.g., Figure 7).

5 Discussion and conclusion

In this paper, we have presented a data-driven approach that forecasts the number of hospitalised COVID-19 patients in the ward and the ICU based on predicted arrival rates and estimated LoS, justified by an underlying network of infinite server queues driven by a Poisson Arrival Location Model. As demonstrated in Section 4, that reports the results of our method for the first COVID-19 peak, the forecasts produced by our method are very accurate. In particular, the forecasts of the maximum occupancy in the ward and the ICU three days ahead are spot-on. This enables hospitals to make informed decisions about whether or not to admit additional COVID-19 patients at their ward or ICU. Indeed, our forecasts are currently being used in four Dutch hospitals during the second COVID-19 peak the Netherlands is facing. For example, the LUMC has incorporated our forecasting graphs in their capacity dashboard (see Figure 7), which is now being reviewed on a daily basis by their physicians.

Now that we have developed a forecasting method that enables informed decision-making for individual hospitals, in future research we aim to build on this method to develop a regional model. Our regional model will not only forecast the COVID-19 occupancy in several hospitals, but also use these forecasts to provide decision support for proactively transferring COVID-19 patients from one hospital in the region to another when the first faces a risk of overcrowding. In our regional model, we will apply a Richards' curve to predict the daily regional number of COVID-19 patients that require hospitalisation, instead of autonomous direct arrivals to each of the hospitals, as this provides an accurate approximation at the regional level.

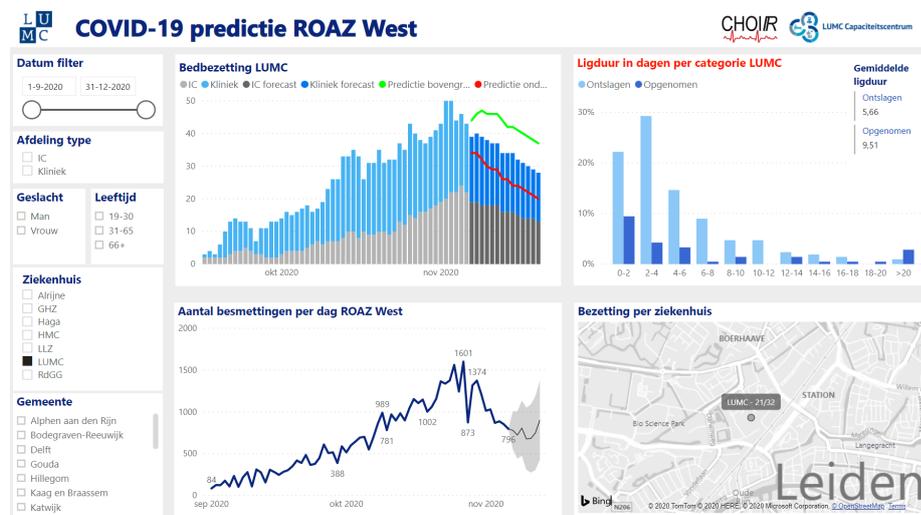


Fig. 7: The LUMC capacity dashboard, with our forecasts of the bed occupancy incorporated in the darker (right) part of the top-left graph provided with the prediction interval (green and red), and our prediction of the daily number of regional infections incorporated in the right (grey) part of the bottom-left graph.

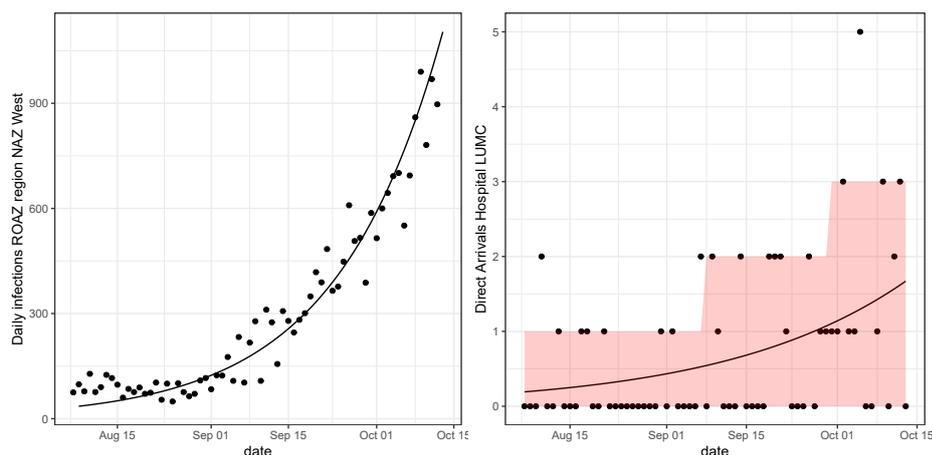


Fig. 8: Left: the number of daily infections in ROAZ region NAZ West and the trend given by our prediction method using the Richards' curve fitted on the (cumulative) daily number of infections. Right: the daily number of autonomous direct arrivals to the LUMC. The trend and 80% prediction interval given by our prediction method using the Richards' curve fitted on this arrival data is also shown for LUMC. Both plots are shown for the period starting from 8-07-2020 until 14-10-2020.

Figure 8 shows the Richards' curve for the LUMC and for the region NAZ West that contains the LUMC. As we aim for a model that provides decision support for patient transfers between hospitals in the region, interaction between the hospitals clearly plays a major role. Therefore, we plan to combine our current research with earlier work on managing the overflow of ICU patients within a region [16], where we will invoke the Modified Offered Load approximation [1, 18] to take into account the capacity constraints on the number of available beds in the hospitals.

Given the quality of our forecasts and the swift implementation of our forecasting method in four Dutch hospitals, we are confident that hospitals will also embrace our regional model. As such, the outlook is that we can provide decision support for one of the major COVID-19 challenges in the Netherlands – transferring COVID-19 patients between hospitals.

References

1. Abdalla N, Boucherie RJ (2002) Blocking probabilities in mobile communications networks with time-varying rates and redialing subscribers. *Annals of Operations Research* 112:15–34
2. Alban A, Chick SE, Dongelmans DA, van der Sluijs AF, Wiersinga WJ, Vlaar AP, Set D (2020) ICU capacity management during the COVID-19 pandemic using a stochastic process simulation. *INSEAD Working Paper No 2020/20/TOM/HMI* Available at <http://dx.doi.org/10.2139/ssrn.3570406>

3. Aleva F, van Mourik L, Broeders M, Paling A, de Jager C (2020) COVID-19 in critically ill patients in North Brabant, the Netherlands: Patient characteristics and outcomes. *Journal of Critical Care* 60:111–115
4. Boucherie RJ, Taylor PG (1993) Transient product form distributions in queueing networks. *Discrete Event Dynamic Systems* 3(4):375–396
5. de Bruin AM, Bekker R, van Zanten L, Koole G (2010) Dimensioning hospital wards using the Erlang loss model. *Annals of Operations Research* 178:23–43
6. van Dijk N, Kortbeek N (2009) Erlang loss bounds for OT-ICU systems. *Queueing systems* 63(1-4):253–280
7. Farcomeni A, Maruotti A, Divino F, Lasinio GJ, Lovison G (2020) An ensemble approach to short-term forecast of COVID-19 intensive care occupancy in Italian regions. arXiv preprint Available at <http://arxiv.org/abs/2005.11975>
8. Goic M, Bozanic-Leal M, Badal M, Basso LJ (2020) COVID-19: Short-term forecast of ICU beds in times of crisis. Available at <http://dx.doi.org/10.2139/ssrn.3693447>
9. Hethcote HW (1989) Three basic epidemiological models. In: Levin SA, Hallam TG, Gross LJ (eds) *Applied Mathematical Ecology*, Springer, Berlin, Heidelberg, pp 119–144
10. van den Heuvel E, Regis M, Zhan Z (2020) Statistical approach for making predictions of confirmed infection and deaths on Corona virus. Tech. rep., Eindhoven University of Technology, available at http://assets.tue.nl/fileadmin/content/pers/2020/03%20March/TUe%20-%20Technical_Report_Prediction_Corona_Virus.pdf
11. Huang C, Wang Y, Li X, et al (2020) Clinical features of patients infected with 2019 novel coronavirus in Wuhan, China. *The Lancet* 395(10223):497–506
12. Ioannidis JPA, Cripps S, Tanner MA (2020) Forecasting for COVID-19 has failed. *International Journal of Forecasting* Available at: <https://doi.org/10.1016/j.ijforecast.2020.08.004>
13. Kaplan EL, Meier P (1958) Nonparametric estimation from incomplete observations. *Journal of the American statistical association* 53(282):457–481
14. Kelly FP (1978) *Reversibility and Stochastic Networks*. Wiley, Chichester
15. Lee SY, Lei B, Mallick BK (2020) Estimation of COVID-19 spread curves integrating global data and borrowing information. arXiv preprint Available at arXiv:2005.00662
16. Litvak N, van Rijsbergen M, Boucherie RJ, van Houdenhoven M (2008) Managing the overflow of intensive care patients. *European Journal of Operational Research* 185(3):998–1010
17. Manca D, Caldiroli D, Storti E (2020) A simplified math approach to predict ICU beds and mortality rate for hospital emergency planning under COVID-19 pandemic. *Computers & Chemical Engineering* 140:106945
18. Massey WA, Whitt W (1993) Networks of infinite-server queues with nonstationary Poisson input. *Queueing Systems* 13(1-3):183–250
19. Massonnaud C, Roux J, P C (2020) COVID-19: Forecasting short term hospital needs in France. medRxiv preprint Available at <https://doi.org/10.1101/2020.03.16.20036939>
20. Murk J, van de Biggelaar R, Stohr J, Verweij J, Buiting A, Wittens S, van Hooft M, Diederer B, Kluiters-de Hingh Y, Ranschaer E, Brouwer A, Retera J, Verheijen M, Ramnarain D, van Ek I, van Oers J (2020) The first 100 COVID-19 patients admitted to the Elisabeth-Tweesteden Hospi-

- tal, Tilburg, The Netherlands [De eerste honderd opgenomen COVID-19-patiënten in het Elisabeth-Tweesteden Ziekenhuis]. *Nederlands tijdschrift voor geneeskunde* Available at <http://www.ntvg.nl/system/files/publications/d5002.pdf> (In Dutch)
21. Nationale Intensive Care Evaluatie (NICE) (2020) COVID-19 in Dutch Intensive Care Units; Patient characteristics and outcomes, version 2020-06-29. Available at https://www.stichting-nice.nl/COVID_report.pdf (In Dutch)
 22. Oswald SA, Nisbet IC, Chiaradia A, Arnold JM (2012) FlexParamCurve: R package for flexible fitting of nonlinear parametric curves. *Methods in Ecology and Evolution* 3(6):1073–1077
 23. Resnick SI (1992) *Adventures in stochastic processes*. Birkhäuser, Boston
 24. Richards FJ (1959) A flexible growth function for empirical use. *Journal of Experimental Botany* 10(2):290–301
 25. RIVM (2020) Kort-cyclische rapportage indirecte effecten COVID-19 op zorg en gezondheid - vierde rapportage. Available at <http://www.rivm.nl/documenten/kort-cyclische-rapportage-indirecte-effecten-covid-19-op-zorg-en-gezondheid-14-september> (In Dutch)
 26. Ule A, Boucherie RJ (2003) On the distribution of calls in a wireless network driven by fluid traffic. *European Journal of Operational Research* 147(1):146–155
 27. Vekaria B, Overton C, Wisniowski A, et al (2020) Hospital length of stay for COVID-19 patients: Data-driven methods for forward planning. *BMC Infectious Diseases*, Preprint (Version 1) Available at <http://www.researchsquare.com/article/rs-56855/v1>
 28. van de Vrugt NM, Schneider AJ, Zonderland ME, Stanford DA, Boucherie RJ (2018) Operations research for occupancy modeling at hospital wards and its integration into practice. In: Kahraman C, Topcu Y (eds) *Operations Research Applications in Health Care Management*. International Series in Operations Research & Management Science, vol 262, Springer, Cham, pp 101–137
 29. Whitt W, Zhang X (2017) A data-driven model of an emergency department. *Operations Research for Health Care* 12:1–15
 30. Wu K, Darcet D, Wang Q, Sornette D (2020) Generalized logistic growth modeling of the COVID-19 outbreak: comparing the dynamics in the 29 provinces in China and in the rest of the world. *arXiv preprint* Available at [arXiv:2003.05681](https://arxiv.org/abs/2003.05681)