Predicting care acuity: A LSTM approach for days-to-day prediction – a case study

Jorg W. R. Bekelaar[1 [0000−0001−7728−1953]], Jolanda J. Luime[2 [0000−0002−1071−7932]], and Renata M. de Carvalho[1 [0000−0001−6129−9278]]

[1] Department of Mathematics and Computer Science
Eindhoven University of Technology, Eindhoven, The Netherlands
j.w.r.bekelaar@student.tue.nl, r.carvalho@ue.nl
[2] Maxima Medical Center, Veldhoven, The Netherlands
jolanda.luime@mmc.nl

Abstract. In recent years, hospitals and other care providers in the Netherlands are coping with a widespread nursing shortage and a directly related increase in nursing workload. This nursing shortage combined with the high nursing workload is associated with higher levels of burnout and reduced job satisfaction among nurses. However, not only the nurses, but also the patients are affected as an increasing nursing workload adversely affects patient safety and satisfaction. Therefore, the aim of this research is to predict the care acuity corresponding to an individual patient for the next admission day, by using the available structured hospital data of the previous admission days. For this purpose, we make use of an LSTM model that is able to predict the care acuity of the next day, based on the hospital data of all previous days of an admission. In this paper, we elaborate on the architecture of the LSTM model and we show that the prediction accuracy of the LSTM model increases with the increase of the available amount of historical event data. We also show that the model is able to identify care acuity differences in terms of the amount of support needed by the patient. Moreover, we discuss how the predictions can be used to identify which patient care related characteristics and different types of nursing activities potentially contribute to the care acuity of a patient.

Keywords: Nurse workload · LSTM prediction model · Event data · Healthcare.

1 Introduction

One hundred years ago, in 1922, the first paper on determination of appropriate nurse staffing levels and bedside nursing time was published [7]. As of today, both still are important topics that have only partially been solved over the last 100 years. In the Netherlands, a number of big steps were made towards the management of nurse staffing levels including the international acceptance of nurse-to-patient ratios, the introduction of policies and regulations regarding shift duration and in general the labour conditions as prescribed in the collective labour agreement. Since this first publication, the effect and impact of bedside nursing time on the quality and the actual level of nursing care provided to the patient continued to become increasingly important for all healthcare actors.

In recent years, hospitals and other care providers are coping with a widespread nursing shortage and a directly related increase in nursing workload in the Netherlands. The nursing shortage results in a higher nursing workload, which is associated with higher levels of burnout and reduced job satisfaction among the nurses. Both could be predecessors for voluntarily stopping clinical nursing work by reschooling to a specialised nurse, nurse practitioner or even by leaving the nursing profession. However, not only the nurses, but also the patients are affected by the nursing shortage and the high workloads: an increasing nursing workload adversely affects patient safety and satisfaction. This emerges by the influence of the nurses on the care process including continuity of care, effective communication at discharge for the continuity of care at home or another care facility, patient centeredness and surveillance.

In order to solve such a complex problem, insights about care acuity - patient characteristics and the nursing care activities that can be expected for a patient - should be gathered. As a first step, we aim at predicting the care acuity expected for an individual patient. Such a prediction already provides a good overview of what nurses should expect for the next admission day and allows for improved decision making in terms of number of nursing staff per shift, which eventually leads to a more equal distribution of the care acuity among the nurses.

In this paper we make use of a Long Short-Term Memory (LSTM) neural network that is able to predict the care acuity per patient for the next admission day. The predictions are based on hospital data collected during the previous days of the corresponding admission, including the amount of care acuity on all previous admission days. Because the conditions of a patient might change due to deterioration or recovery, the LSTM model also considers the conditions of the patient at that particular day.

The remainder of the paper is organized as follows. Section 2 provides the concepts that are relevant for this research. Related research is discussed in Section 3. Section 4 describes the data used in this research, while the approach to data preprocessing and the prediction model are explained in Section 5. Then, Section 6 discusses the obtained results. Section 7 presents the final conclusions.

2 Background

2.1 Process Mining

Predictive process monitoring [1] is a segment of process mining interested in predicting the future of an ongoing process execution. For that, it relies on the *event log*, which is a structured dataset containing information about different executions of a process and can also be seen as a collection of traces. A *trace* is a non-empty sequence of events related to the same process execution, ordered by time. An *event* is then an atomic part of the process execution and is characterized by various properties, e.g., an event has a timestamp and it corresponds to an activity. Many approaches in predictive process monitoring leverage machine learning techniques, such as neural networks (cf. Section 3). For such approaches, the event log information should be encoded in terms of features. Usually, the event and its data *payload* are part of this feature set.

The predictive process is split into two phases: training and prediction. The *training* phase counts on the information of previously completed process executions to learn relations in the data. The *prediction* phase considers an ongoing process execution. This means that such a process execution is not completed yet. The known part of the trace is defined as the *prefix* and is used as the input for the prediction model. The future sequence of events that is supposed to take place after the prefix, is defined as the *suffix* and represents the prediction made by the model. Predictive process monitoring is also applied to predict the outcome of a process execution, or its completion time. For some organizations, it can be highly valuable to be able to predict in advance what is going to happen to a process execution. In this context, the organization can focus on *preventing* issues from happening, rather than *reacting* to them after their occurrence [10].

2.2 Long Short-Term Memory Neural Networks

The Long Short-Term Memory (LSTM) model is an advanced form of a Recurrent Neural Network (RNN) that allows information to persist [4]. LSTM models are explicitly designed to solve the problem of long-term dependencies by changing the structure of hidden neurons in a traditional RNN. A LSTM model can be used for predicting on the basis of time series data due to the characteristic of retaining the information for a long period of time. Hence, it is considered effective and general at capturing long term temporal dependencies [2].

In practice, the LSTM architecture consists of a set of recurrently connected sub-networks called memory blocks. An individual memory block contains a functional module that is known as the memory cell and a number of different gates. The memory cell is responsible for remembering the temporal state of the neural network over arbitrary time intervals, while the gates formed by multiplicative units regulate the flow of information associated with the memory cell. Together, the memory blocks form the key part of the LSTM that enhances its capability to model long term dependencies. A memory block contains both a hidden state and a cell state known as short term memory and long term memory respectively. The cell state encodes an aggregation of the data from all previous time steps that have been processed by the LSTM, while the hidden state is used to encode a characterisation of the input data of solely the previous time step.

A memory block contains three gates that together regulate the information flow: the *forget* gate, which decides what information should be removed from the previous cell state, the *input* gate, which quantifies the importance of the new information carried by the input and the *output* gate, which extracts the useful information from the current LSTM block by computing the new hidden state. LSTM models are appropriate to handle sequential data of different sizes and hence, process executions of different lengths. They are also able to consider additional information about the events, resources and any other data payload.

3 Related Work

Over the years, several attempts have been made to quantify and predict care acuity [3]. Some of the earliest methods that have been developed to quantify care acuity are the Therapeutic Intervention Scoring System (TISS), the TISS-28 method, the Nine Equivalents of Nursing Manpower (NEMS) and the Nursing Activities Score (NAS). All methods are based on an identical principle that distinguishes a number of activities that are scored on a $1-4$ basis according to the intensity of involvement. Subsequently, the assigned scores are used to group patients into separate classes. The Project Research of Nursing (PRN) assigns a score to each nursing activity based on, among other factors, the corresponding duration, frequency and the number of nurses required to execute the activity, while the Time Oriented Scoring System (TOSS) is a time based system for quantifying care acuity that exactly times a number of preselected nursing tasks. Alternatively, the Rafaela method relates each activity to a domain with varying nursing intensities. The points assigned to the different domains are added up per patient and department to compute the actual workload per nurse.

As is mentioned in section 2.2, a LSTM can be used for predicting on the basis of time series data. Today, applications and research of LSTM for time series prediction include usage in the healthcare sector to predict the day of discharge [9], hospital performance metrics [5] or to make clinical predictions [8].

Also in the context of process mining, the usage of LSTM is not new. LSTM models are notably suitable to deal with problems that involve sequences, such as event traces. Mostly, LSTMs are used in attempting to predict the next activity in a trace. Tax *et al.* [12] and Tello-Leal *et al.* [13] employ LSTM to predict the next event of a running case. Tax goes beyond, predicting its timestamp and showing how the method can be used to predict the full continuation of a case and its remaining time. Pham *et al.* [11] also uses LSTM models to predict the next activity in a trace and who would perform such an activity.

Building on the aforementioned works, we believe that simply representing the trace events with its data attributes is not enough for the problem of predicting care acuity. Firstly, we think it necessary to use all the historic hospital data of an admission, as it can show how fast the patient recovery process is. Therefore, we cannot use predefined prefix lengths. Secondly, there might be data to learn about this recovery process that are crucial for the model, but are not associated to any event directly. Such data mostly come from monitoring the patient's vital parameters.

To help in the decision making of the distribution of patients among nurses, we do not need to know the explicit sequence of events that will happen for a patient. However, it is necessary to know how much support is expected to be required from a nurse and if any critical task will take place. So, this work is different from previous research mainly on how we group and represent the workload of tasks. Besides this, we are interested in predicting a numeric value representing the care acuity on the next admission day. This objective is a regression task, rather than a classification task such as predicting the next event label.

4 Data Description

The data for this research have been obtained from the clinical departments of a hospital located in the Netherlands throughout 2018. The clinical departments consist of 8 different departments that are responsible for providing different levels of adult patient care services: cardiology, gastrointestinal surgery, general surgery, gynaecology, internal medicine, neurology, oncology, orthopaedics and pulmonary. The data from the intensive care unit (ICU) and the short stay unit (SSU) were not considered in this research. The resulting dataset distinguishes 62 features per record, including:

– five *patient features*: age, BMI, pre-hospitalisation physical mobility, sex, social economic status and the unique patient identifier;
– eight *admission features*, such as the admission department, inter-department transfers, reason and type of the admission and the specialism of the doctor;
– seven *time features*, such as the current date, day of the week, month, season, time and the current length of stay;
– eleven *medication features*, such as the daily number of inhaled, injected, intravenous, oncology, oral and pain medications and the daily number of either newly started or discontinued medications;
– four *examination features*: the daily number of bloodcount, imaging, laboratory and microbiology tests;
– eleven *vital parameter features*, including vital functions such as the daily maximum and minimum body temperature, early warning score, oxygen saturation level, pain score, respiratory rate and systolic blood pressure;
– six *nursing activity features*, such as the description, explanation and the type of the nursing activity and the maximum number of daily occurrences;
– three *nursing notes features*, such as the length (in terms of lines and words) and the daily amount of nursing notes;
– three *DBC features*, such as the amount of diagnosis treatment combinations in the three previous years per specialism;
– two *operation room (OR) features*: whether the patient underwent a procedure, or more than two procedures in the OR.
– two *discharge features*: whether the patient is going to be discharged in the next 24 or 48 hours.

5 Methodology

Data was extracted using a software package on a copy of the electronic patient file system. This software (CTcue) allows for immediate pseudonymisation of the data using NLP and pseudo-IDs to de-identify all doctor and patient names in unstructured text. Data from the nurse activity plan was extracted by a dedicated SQL query. A common data model was created by representing each nursing care activity as a single record with the activity, the data source and the timestamp. Each of the records was assigned to a department, room and bed and an unique admission number. If multiple clinical departments were visited by the patient during a day, we assigned the final department to this day.

5.1 Workload Assignment

Based on the current daily nursing practice, 11 core activity categories were identified: *Activities of Daily Living (ADL)*, *Bed rest*, *Communication*, *Drains*, *Excretion*, *Feeding tube*, *Infusions & Lines*, *Measurements & Observations*, *Reporting*, *Respiration* and *Woundcare*. Initially, a tree was built containing these categories and the individual nursing care activities, classified into each category. Points were assigned to each individual activity based on existing work by Jonker [6]. The research conducted by Jonker implemented a dedicated scoring system based on the Rafaela method [3] and assigns $0-3$ points to each individual activity. This research implemented a similar scoring system with a similar scoring scale, but with a number of exceptions that received either 4 or 5 points, based on suggestions by the board of nurses. For each patient, all the activities in each category were summed to compute their daily care acuity. Care acuity is a latent variable that has no golden or reference standard, but of which the validity could be constructed via nurse opinion. Figure 1 provides an overview of the care acuity distribution in the training, validation and test datasets.

An exception was formed by the *ADL* and *Communication* categories that do not contain standalone activities. For the *ADL* category, we assigned a fixed baseline workload based on whether a patient was independent, partially independent or fully dependent on the nurse during *ADL* activities. These scenarios are used in daily nursing practice and were regarded as relevant by the nurses. Additional points were assigned for auxiliaries used for patient movement and transfer, equipment and medications, medical devices and the patient's mobility status, as they complicate the execution of the nursing care activities contained in the *ADL* category. The communication baseline workload consists of one scenario with four components that contribute to the care acuity: bedside rounds, communication with family members, medical handovers across shifts and time spent by nurses on registering medical notes and additional reporting. Special attention was paid to patients with a delirium. Delirium is a sudden change in

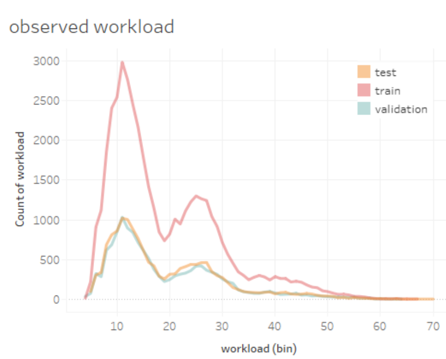

Fig. 1. The distribution of the care acuity.

the mental state of a patient. If patients are delusional or get a tendency to walk away, nurses need to attend the patient more often. We used the results of the Delirium Observation Questionnaire for these patients to assign a workload to the delirium on a daily basis.

5.2 Prediction Preparation and Target

Finally, the time series data are compressed using a compression interval of exactly 1 admission day. As a results, each admission day is represented by a single record and the input for the LSTM model consists of a large input vector. Because of the different magnitude and unit of the different features, each feature is individually scaled and translated such that its values are in a range between 0 and 1. The scaler is fit on the training set and used to transform the data contained in both the validation and the test set. After the LSTM model has produced the predictions for the care acuity, it is necessary to reverse the scaling to retrieve the actual predictions of the care acuity.

The output of the LSTM model consists of the total care acuity of an individual patient for the next admission day. For the final day of an admission, there are no registered nursing activities on the next day. As a consequence, the value of target variable for each final day of an admission will be equal to 0. In order to predict the care acuity of the patient, the input variables are formed by the features that were contained in the original dataset and specified in Section 4, together with the additional features that were generated in the feature engineering steps during the data preprocessing. The LSTM model uses these features to predict the care acuity corresponding to an individual patient for the next admission day, based on the available date and time stamped hospital data of the current day. If the current length of stay is longer than 1 day, the prediction is made based on the available date and time stamped hospital data of the current day and the previous admission days.

5.3 LSTM Prediction Model

A LSTM neural network consists of different types of layers, including at least one LSTM layer. Figure 2 depicts the architecture of the LSTM neural network that was used in this research. The initial layer of the sequential model that represents the LSTM model is made up by a LSTM input layer. In order to pass the input data to the LSTM layer of the sequential model, the *input shape* parameter of the LSTM layer must be specified. The input to the LSTM layer must be three-dimensional and consists of samples, time steps and features. A sample is one sequence and represented by a unique admission in the dataset. A time step is one point of observation in the sample and represented by a single admission day, while each feature is one observation at a time step. Furthermore, the *units* parameter of the LSTM layer indicates the dimension of the hidden state and the number of parameters in the LSTM layer. Lastly, the *return_sequences* parameter of the LSTM layer ensures that the full output sequence is returned. By enabling the *return_sequences* parameter, one is allowed to access the output of the hidden

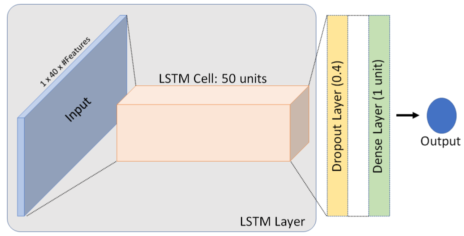

Fig. 2. The architecture of LSTM neural network

state for each time step, leading to a prediction of the care acuity on the next day for each admission day. This way, the LSTM layer eventually facilitates the prediction and subsequently the evaluation of the care acuity for each day contained in an admission, based on the previous admission days.

The LSTM layer is followed by a Dropout layer, which is a regularisation technique that randomly selects nodes to be dropped during each weight update cycle. To drop the inputs, the layer randomly sets input units to 0 with a frequency that is equal to the *dropout rate* parameter. To ensure that no values are dropped during inference, the Dropout layer only applies during training and not when the performance of the model is evaluated. It ensures that neurons do not end up relying too much on other neurons or on specific inputs, but that the model learns the meaningful interactions and patterns in the data. It produces a robust LSTM model, has the effect of reducing overfitting and eventually improving model performance, by ensuring that the weights are optimised for the general problem instead of for noise in the data.

Lastly, a dense output layer is added to the model. The dense layer is a regular densely-connected neural network layer that is used to consolidate output from the LSTM layer to the predicted values. Because the *return_sequences* parameter of the LSTM layer is enabled, the dense layer receives the hidden state output of the LSTM layer for each input time step. In order to ensure that the output of the LSTM model has the dimensionality of the desired target, so that that the output of the dense layer consists of a prediction of the care acuity for each admission day contained in the test dataset, the value of the *unit* parameter for the dense layer is set to 1.

In order to determine the optimal values of the different hyperparameters, a random search followed by a grid search are executed. The random search randomly samples from a wide range of hyperparameter values to narrow down the search range for each hyperparameter, by performing k-fold cross validation. Subsequently, the grid search further refines the optimal values for the hyperparameters by evaluating the best hyperparameter values returned by the random search. The results of the grid search together with the specific characteristics of the hospital dataset indicate that the optimal model performance is achieved by training the model for 10 epochs, with a batch size of 4, using the *nadam* optimizer and the *mean absolute error* as the loss function.

6 Results

6.1 Data and Descriptive Statistics

In total, 16755 unique admissions of 12224 unique clinical patients of all ages in 2018 were available for analysis. We kept 15477 adult admissions corresponding to patients above 18 years of age, as there are specific medical procedures in place for patients under the age of 18 that come with specialized nursing care activities outside the scope of this research. Of those, we excluded 2782 admissions of women that were in labor, because the associated patient care at the Gynecology department differs significantly from the remaining clinical departments. In order to avoid incomplete admissions, we only included those that solely have admission days in 2018. This resulted in 12492 unique admissions corresponding to 9931 unique patients. The average length of stay was equal to 7 days (SD 10 days) with a median of 4 days (IQR 2-9 days). We subsequently split the dataset randomly in a 60% training set that included 7495 admissions, a 20% validation set ($n = 2498$) and a 20% test set ($n = 2499$). This resulted in an equal distribution of care acuity as is shown in Figure 1, with a median equal to 16 (IQR 9-26) for the training dataset. The fact that the care acuity increases with the length of stay can be explained by the fact that sicker patients that require additional nursing care remain admitted to the hospital, while the relatively fitter patients that require less nursing care are discharged from the hospital.

6.2 Evaluation Metrics

To evaluate the performance of the model, a selection of different performance metrics are used. The three most well-known metrics that are used for evaluating and reporting the performance of a regression model are the Mean Absolute Error (MAE) – calculated as the average of the absolute error values –, the Mean Squared Error (MSE) – calculated as the mean of the squared differences between the predicted care acuity and the actual care acuity values – and the Root Mean Squared Error (RMSE) – calculated as the square root of the Mean Squared Error. Besides this, the R-squared score (R^2) indicates how well the model is able to predict the value of the target variable and is the percentage of the target variable variation that can be explained by the model. It is calculated by dividing the variance explained by the model by the total variance. Lastly, the symmetric Mean Absolute Percentage Error (sMAPE) returns the error of the model as a percentage, making it easy to compare and understand the model accuracy across different configurations, datasets and use cases.

6.3 Performance Measures

First of all, Figure 3 shows the predicted care acuity for an exemplary patient contained in the test set. The lower part of the figure represents the actual daily workload for each activity category. The actual care acuity, which is the sum of the different categories, and the predicted care acuity for each admission day are represented by the green and the red line in the top part of the figure respectively.

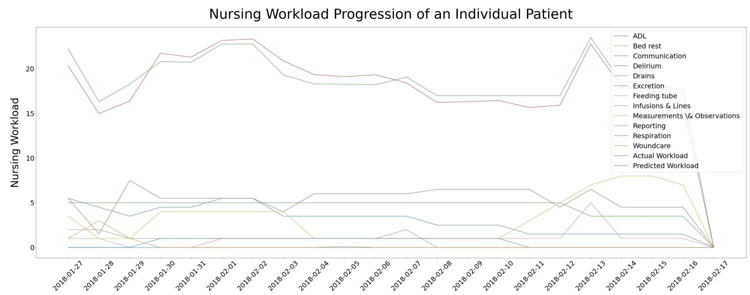

Fig. 3. Predicted Workload vs Actual Workload for an exemplary patient.

On top of this, Figure 4 shows the average prediction error per consecutive admission day. Each bar represents the average prediction error for the admission day indicated by the value on the x-axis of the plot.

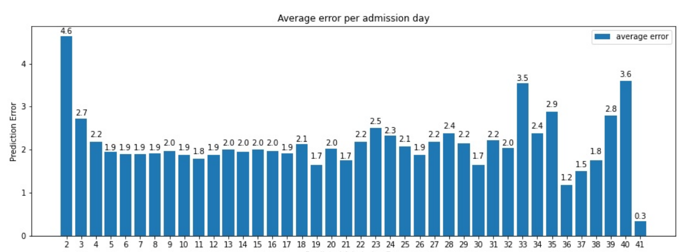

Fig. 4. Average prediction error per admission day.

It can be observed that the error decreases during the first five days of the admissions contained in the test set. After the fifth admission day, the prediction error stabilizes around a value of 2 and after the thirty-second admission day, the figure shows multiple outliers for which the value of the prediction error is 0.5 larger or smaller than the previously stable value of 2.0. This can be explained by the fact that it is harder for the model to learn about longer admissions, as they become more scarce. Longer admissions often consider patients that are exceptional, causing different and unexpected things to happen, such as infections or relapse. The highest errors are caused by an unexpected change of either the *ADL* or the *Communication* baseline scenario.

Furthermore, Table 1 displays the evaluation metrics for the predictions on the test set that indicate the performance of the LSTM model. The MAE indicates that the model performs well in general, as most of the errors are low. However, the high MSE value indicates that when the error is on the high side, it is far above the average. Finally, the R^2 score shows a high level of correlation.

Evaluation Metric	**Value**
MSE	19.376712799072266
RMSE	4.4018988609313965
MAE	2.230191946029663
R^2	0.8627777362127644
sMAPE	47.75%

Table 1. The evaluation metrics of the LSTM model on the test dataset.

7 Conclusions and Discussion

Nurses have a strong influence on the quality-of-care that patients receive in the hospital. To maintain high quality of care under the stress of the nurse staffing shortage, it has become critical to distribute workload evenly and to see what type of work maybe automated or done by others. This requires easy access to insights in the observed and expected care acuity of each patient in daily clinical nursing care. In this paper, we addressed this by digitally identifying and quantifying the care acuity corresponding to individual nursing activities and subsequently, predicting care acuity with a one-day time horizon to allow for an equal assignment of workload using an LSTM model. The architecture of the LSTM model proved itself suitable to facilitate this time series prediction for the hospital dataset. It displays the ability to adapt and make reliable predictions for the consecutive admission days.

The LSTM model was able to learn from the data and on group level resembles the observed data very well. If patients' care acuity fluctuates only slightly, the model is very well equipped to pick up small changes and predict the care acuity on the next day correctly. However, if there is a sudden deterioration of the patients' conditions, the model picks up the changes, but seemingly with one day delay. This suggest that the model drives to much on the previously observed care acuity and less on the change of the patient's condition. Future work needs to be done to weight the patients' characteristics and condition differently to stress the model to learn more from these features. Also, the initial care acuity that the model assigns to the first admission day seems rather arbitrary. This was to be expected as there is no data available to learn from. As a consequence, the model has no other means than to assign the average workload and optimize from there. One solution here could be to train another model to learn the care acuity for first day based on patient characteristics (e.g., reason for admission, vital functions) first and use these input values for the LSTM model to use. A similar approach was applied to predict the day of discharge by using the input of a GPboost model to determine the day of discharge.

More work is needed on the validity of the assignment of care acuity points to individual nursing care activities. The initial approach we took worked rather well. Consecutive rounds of discussion with nurses were performed to agree upon and optimize the current assignment of points. However, constructing validity against the nurses' opinions should be further researched to reflect daily nursing care well. Moreover, we need to put more weight on features related to patient characteristics, so that the ability of the model to make predictions for individual patients improves. In the end, this research contributed in digitally identifying and quantifying the care acuity corresponding to individual nursing activities and show that patients' care acuity can be predicted one-day ahead.

References

1. Di Francescomarino, C., Ghidini, C.: Predictive Process Monitoring, pp. 320–346. Springer International Publishing, Cham (2022)
2. Greff, K., Srivastava, R.K., Koutník, J., Steunebrink, B.R., Schmidhuber, J.: Lstm: A search space odyssey. IEEE Transactions on Neural Networks and Learning Systems **28**(10), 2222–2232 (2017)
3. Griffiths, P., Saville, C., Ball, J., Jones, J., Pattison, N., Monks, T.: Nursing workload, nurse staffing methodologies and tools: A systematic scoping review and discussion. International Journal of Nursing Studies **103**, 103487 (2020)
4. Houdt, G.V., Mosquera, C., Nápoles, G.: A review on the long short-term memory model. Artificial Intelligence Review pp. 1–27 (2020)
5. Jia, Q., Zhu, Y., Xu, R., Zhang, Y., Zhao, Y.: Making the hospital smart: using a deep long short-term memory model to predict hospital performance metrics. Industrial Management & Data Systems (Jan 2022)
6. Jonker, J.: De validiteits- en betrouwbaarheidstest van het verpleegkundig zorgzwaartemodel. (2019)
7. Lewinski-Corwin, E.H.: The hospital nursing situation. The American Journal of Nursing **22**(8), 603–606 (1922), http://www.jstor.org/stable/3406790
8. Lu, W., Ma, L., Chen, H., Jiang, X., Gong, M.: A clinical prediction model in health time series data based on long short-term memory network optimized by fruit fly optimization algorithm. IEEE Access **8**, 136014–136023 (2020)
9. Luo, L., Xu, X., Li, J., Shen, W.: Short-term forecasting of hospital discharge volume based on time series analysis. 2017 IEEE 19th International Conference on e-Health Networking, Applications and Services (Healthcom) pp. 1–6 (2017)
10. Ly, L.T., Maggi, F.M., Montali, M., Rinderle-Ma, S., van der Aalst, W.M.: Compliance monitoring in business processes: Functionalities, application, and tool-support. Information Systems **54**, 209–234 (2015)
11. Pham, D.L., Ahn, H., Kim, K.S., Kim, K.P.: Process-aware enterprise social network prediction and experiment using lstm neural network models. IEEE Access **9**, 57922–57940 (2021)
12. Tax, N., Verenich, I., La Rosa, M., Dumas, M.: Predictive business process monitoring with lstm neural networks. In: Dubois, E., Pohl, K. (eds.) Advanced Information Systems Engineering. pp. 477–492. Springer International Publishing, Cham (2017)
13. Tello-Leal, E., Roa, J., Rubiolo, M., Ramirez-Alcocer, U.M.: Predicting activities in business processes with lstm recurrent neural networks. In: 2018 ITU Kaleidoscope: Machine Learning for a 5G Future (ITU K). pp. 1–7 (2018)