



Proceeding Paper Evaluating a Recurrent Neural Network Model for Predicting Readmission to Cardiovascular ICUs Based on Clinical Time Series Data⁺

Sobhan Moazemi ^{*,‡,§}, Sebastian Kalkhoff [§], Steven Kessler, Zeynep Boztoprak, Vincent Hettlich, Artur Liebrecht, Roman Bibo, Bastian Dewitz, Artur Lichtenberg, Hug Aubin ^{||} and Falko Schmid ^{||}

Digital Health Lab Düsseldorf, Department of Cardiac Surgery, Medical Faculty and University Hospital Düsseldorf, Heinrich-Heine-University Düsseldorf, 40225 Düsseldorf, Germany; sebastian.kalkhoff@med.uni-duesseldorf.de (S.K.); steven.kessler@med.uni-duesseldorf.de (S.K.); zeynep.boztoprak@hhu.de (Z.B.); vincenthendrik.hettlich@med.uni-duesseldorf.de (V.H.); artur.liebrecht@med.uni-duesseldorf.de (A.L.); roman.bibo@med.uni-duesseldorf.de (R.B.); bastian.dewitz@med.uni-duesseldorf.de (B.D.); artur.lichtenberg@med.uni-duesseldorf.de (A.L.); hug.aubin@med.uni-duesseldorf.de (F.S.) *

* Correspondence: sobhan.moazemi@med.uni-duesseldorf.de

+ Presented at the 8th International Conference on Time Series and Forecasting, Gran Canaria, Spain, 27–30 June 2022.

- ‡ Current address: Moorenstr. 5, 40225 Düsseldorf, Germany.
- § These authors contributed equally to this work.
- \parallel These authors contributed equally to this work as senior authors.

Abstract: Unexpected readmission to intensive care units (ICUs) endangers patients' lives due to premature patient transfers or prolonged stays at the care units. This can be mitigated by stratification of the readmission risk at discharge times using state-of-the-art machine learning (ML) methods. We fitted two alternative recurrent neural network (RNN) models based on long short-term memory (LSTM) on the Medical Information Mart for Intensive Care (MIMIC-III) dataset and evaluated them with an independent cohort from our hospital's ICU (UKD). The first model processed all the available time series data from each patient's ICU stay, whereas the second model focused on the data from the last 48 hours of the ICU stay prior to transfer. Our readmission prediction on MIMIC data reached an area under the curve of receiver operating characteristic (AUC-ROC) of 0.82. Furthermore, the model with the 48 h time frame outperformed the other model, as both models were applied to the independent test cohort. The results suggest that the RNN model for time series forecasting holds promise for future use as a clinical decision support tool, although follow-up studies with larger cohorts as well as user studies should be conducted to assess the generalizability and usability of the methods, respectively.

Keywords: readmission prediction; intensive care unit (ICU); recurrent neural network (RNN); long short-term memory (LSTM); machine learning (ML); time series analysis; health forecasting

1. Introduction

According to different retrospective and review studies [1–3], the rates of readmission to ICUs in hospitals in developed countries are quantified inconsistently in the range from 0.14 to 14.5 percent. Regardless of the inconsistent readmission rates, patients who are readmitted to ICUs due to inappropriate discharge time are subject to life-threatening risks, with respiratory and cardiac complications as the most common causes of readmission [1,4]. At the same time, prolonged stays at ICUs can also lead to increased mortality or poor long-term prognosis for patients [5]. In particular, patients who undergo cardiac surgery are prone to longer ICU stays due to the invasive nature of heart surgery and the resulting increased risk of postoperative complications [6]. Therefore, it is important to find the



Citation: Moazemi, S.; Kalkhoff, S.; Kessler, S.; Boztoprak, Z.; Hettlich, V.; Liebrecht, A.; Bibo, R.; Dewitz, B.; Lichtenberg, A.; Aubin, H.; et al. Evaluating a Recurrent Neural Network Model for Predicting Readmission to Cardiovascular ICUs Based on Clinical Time Series Data. *Eng. Proc.* **2022**, *18*, 1. https:// doi.org/10.3390/engproc2022018001

Academic Editors: Ignacio Rojas, Hector Pomares, Olga Valenzuela, Fernando Rojas and Luis Javier Herrera

Published: 16 June 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https:// creativecommons.org/licenses/by/ 4.0/). optimal point in time for transfer, for this particular vulnerable group of patients. Intensivists need to quantify patients' readmission risks in order to determine this point in time and avoid unplanned readmissions. This implies an increasing need for clinical decision support tools which complement current discharge routines through the use of predictive models taking advantage of artificial intelligence (AI).

AI in general and ML in particular have been widely used in many diagnostic and prognostic domains, such as oncology [7,8] and computational neuroscience [9,10], to provide clinical decision support tools, which either replace or complement established diagnostic and prognostic routines. More specifically, time series analysis has been successfully applied in many medical domains, including management of type 2 diabetes [11], hospital admission prediction [12], age-related death prediction [13] and ICU readmission prediction [3,14].

From a technical point of view, as related works suggest, a variety of techniques ranging from simple probabilistic predictors such as logistic regression (LR) and ML-based classifiers such as support vector machines (SVMs) to more sophisticated models such as convolutional neural networks (CNNs) have shown their potential for the prediction of readmission to ICUs [3,15,16].

In a previous study [17], currently under review, we trained a model with a recurrent architecture on a subset of the MIMIC-III dataset, an open access dataset of time series data from more than 50,000 care unit stays with different health conditions at a single center [18], without any limitation on the time frames of the input time series data, and showed its superiority to logistic regression and a feed forward network in predicting readmission. In the current study, we analyze the performance of a recurrent neural network (RNN) model processing health-related times series data, taking advantage of the long short-term memory (LSTM) method with slight architectural enhancements compared to the previous model, considering time series data from the last 48 h of patients' stays at ICUs, sampled in 60-minute intervals.

The main contribution of our work is to propose and evaluate two alternative MLbased models, which are trained and fit using a publicly available dataset of health-related time series data, with and without concrete time windows (in compliance with state-of-theart metrics in the domain) prior to patients' discharge from ICUs. Moreover, as a further contribution, the proposed models were evaluated using a real-world hospital cohort with data from our own cardiovascular intensive care unit at UKD to assess how they would generalize against new sets of unseen data.

In the next sections, first, a detailed overview of the patient data and methods is given, and then the results of model training and evaluations are presented and discussed. Finally, the relevance of the findings and possible future follow-ups will be discussed.

2. Materials and Methods

This section explains the methodology of the study. First, the details of the dataset curation for training, validation and test steps will be elaborated, including the selection criteria and cohort statistics, which will be followed by the description of the integrated times series data. Then, the details of the training and validation step using 5-fold cross-validation on the MIMIC-III dataset are discussed. Finally, the details of the test step on an unseen set of data from a held-out MIMIC cohort as well as the UKD database are presented. The entire in-house-developed analysis pipeline is developed in the Python (V3.6.) programming language. The LSTM models for time series data analysis are implemented using the PyTorch library [19].

The inclusion criteria of subjects from both MIMIC and UKD datasets have been defined as follows:

 Patients who were transferred from the ICU to a normal station and then returned to the ICU within 48 h. Due to possible logistical reasons, patients whose first ICU stay was less than 24 h are excluded. These patients are labeled as 'readmitted'.

- 2. **Patients who died during the hospital stay.** The dead subjects are also labeled as 'readmitted'.
- 3. Patients who were transferred from the ICU to a normal station and then were not returned to the ICU within 30 days. Those who did not return within 30 days after transfer from ICU are labeled as 'non-readmitted'.

2.1. Patient Cohorts

For the training, we focused on the cases of patients from the MIMIC-III dataset who visited cardiovascular care units. As a result, the training subject cohort consisted of 11,513 patients, out of which 966 patients were labeled as readmitted and 10,547 were labeled as non-readmitted. From the MIMIC dataset, the train subjects' ages ranged from 17 to 89 years (mean(M) = 66, standard deviation(SD) = 13.51). For anonymization purposes, the real admission times are not attainable from the MIMIC dataset.

For test purposes, a cohort of 502 patients who visited the cardiovascular ICU at the UKD were retrospectively analyzed. Among these, 100 patients were returned to the ICU within the 48 h after discharge (labeled as 'readmitted') and 402 patients visited the ICU just once, without being readmitted within 30 days after their discharge (labeled as 'non-readmitted'). The ages of the subjects from the UKD database were quantified in the range from 19 to 95 years (M = 67, SD = 12.25). The cohort corresponds to cardiovascular ICU stays in the time period from January 2017 to December 2020. Due to the retrospective character of the study, patient consent was waived.

2.2. Time Series and Patient Data

For each subject in the train or test cohort, 10 vital or laboratory variables (selected by highly qualified cardiovascular surgeons) were captured during the patient's stay at the corresponding intensive care unit. To form the time series data for all the subjects, 60 min time intervals have been taken into account. The missing values were extrapolated using the mean value of the corresponding variable throughout the entire corresponding cohort. Moreover, from the patients' files, the age and weight values were extracted. Finally, the length of stay (LoS) of each subject was quantified from the admission time entries from the corresponding database (MIMIC or UKD) and added to the rest of the features to end up with feature vectors of size 13. Table 1 provides the list of the 10 variables used for the time series data analysis and the 3 extra patient variables.

Table 1. The input parameters including time series data from laboratory or vital values and information from patient files (ABP: ambulatory blood pressure, LoS: length of stay). The table is adapted from [17].

Lab Values	Vital Signs	Patient Information
Creatinine	Body Temperature	Age
Blood PH	ABP	Weight
Sodium	Heart Rate	LoS
Potassium	Oxygenation	
Hematocrit		
Bilirubin		

2.3. Preprocessing

Preprocessing of the input data is applied, similarly to our previous work [17], consisting of the following steps: First, the time series features and patient variables are obtained from the corresponding dataset (MIMIC or UKD). Then, all values are standardized and missing data are replaced with the corresponding distribution's mean value. In the last step, the data are both re-sampled and oversampled before being processed by the corresponding model. Figure 1 gives an overview of the preprocessing pipeline.



Figure 1. Preprocessing pipeline. First, the data are extracted and standardized, and then the missing data are filled in. Finally, re-sampling and oversampling are applied before the data are fed to the models. This figure is adapted from previous work [17].

In the feature extraction step, the case IDs of the subjects who met our selection criteria have been processed to acquire the feature values from the corresponding database using appropriate queries. Then, the normalization is applied at two different levels. First, the features which might have been stored in different units (such as weight) have been unified and converted accordingly and cases with extreme and outlier values have been removed from the cohorts. In the second step, we applied standardization around the mean value to end up with normalized feature vectors. The standardization of each variable is applied as defined in Equation (1):

$$\hat{s}_i = \frac{s_i - \mu_x}{\sigma_x},\tag{1}$$

where \hat{s}_i is the standardized value, s_i is the original value, μ_x is the mean of the variable over the whole cohort *x* and σ_x is the standard deviation over *x*.

As required by the LSTM method, all the input feature vectors must be of equal size before being analyzed. Thus, for all the timestamps without an actual measurement, the missing value was filled in with zero (which is the mean value after standardization). Moreover, as the vital and lab values of patients are usually measured in fragmented timestamps, the feature vectors extracted from the databases were non-uniformly distributed along subject cohorts. Thus, to unify the entire cohorts, the time series parameters are resampled to a frequency of one entry per hour. In case there exist more than one measurement per hour, the mean of each hour is used as the value. Alternatively, if there is no record for a time step, the mean over the whole cohort is used. Furthermore, the imbalance in the label distribution (there are more non-readmitted cases than readmitted ones in both cohorts) might result in the training of a deep learning model that is not able to predict the minority class labels accurately [20]. Therefore, the data for training the LSTM models were oversampled in order to have evenly distributed classes.

2.4. Model Training

In a previous study [17], we have shown the superiority of an RNN model based on LSTM architecture to logistic regression and a feed forward network using the MIMIC-III dataset. Therefore, in the current study, two alternative models with LSTM architecture have been analyzed, compared and further evaluated on unseen sets of data from open access and available in-house data. Both models are trained and validated using a subset of the MIMIC-III dataset, which was selected randomly. The first model, which will be denoted as the 'cropped' model from now on, applies a fixed window of the last 48 h to each ICU stay. The second model, which will be denoted as 'uncropped' in the next parts of the text, applies no limits on the available time series data from each subject.

As a special kind of recurrent neural network (RNN), long short-term memory (LSTM) is able to learn long-term dependencies in time series data. The input to an LSTM is a batch of arrays containing time series data. Although the LSTM model can handle inputs with varying lengths, the input length has to be equal inside each batch. Thus, to pad each time series with zeros until they have equal length, a PyTorch object is used.

To be able to capture non-linear relations between input variables and target labels, a hidden layer with the size of 50 neurons is applied after the LSTM pipeline. The input to the models are the time series data. For the cropped model, a fixed window of 48 h is used, while, for the uncropped model, no constraint on the length of time series data is applied. After processing the input time series batches, the three patient variables (age, weight and LoS) are concatenated to the feature vectors. Then, a rectified linear unit (ReLU) is applied as the activation function, while a dropout layer with a value of 0.3 is used for regularization purposes. The final fully connected layer consists of 2 neurons, which are then passed through to the softmax function, which quantifies the final class probabilities (one for each label).

Figure 2 illustrates the architecture of the LSTM-based models. The x_i input refers to the i-th time step of all time series. Table 2 summarizes the hyperparameters for model fitting.



Figure 2. LSTM models' architecture: First, the time series data are passed to the LSTM part. Then, the patient information (age, weight and length of stay) is concatenated to the feature vectors from LSTM output. Next, the ReLU and dropout are applied before the feature vectors are passed through the hidden layers. Finally, the softmax function is used to assign probabilities for each binary label.

Table 2. The hyperparameters of the LSTM-based model. The table is adapted from [17].

Hyperparameters	Values
optimization algorithm	Adam
learning rate	0.003
loss function	cross-entropy
batch size	32
epochs	early stopping

2.5. Model Validation

The training cohort from the MIMIC dataset is first split into separate train and validation cohorts with similar ratios of target labels. Then, to validate the model internally, 5-fold cross-validation is applied with random folds. As a result, the hyperparameters of the best performing models are stored to be used for the corresponding test steps. The training of each model continued until the area under the curve of precision recall (AUC-PR) stopped increasing after several epochs. Subsequently, the model that achieved the highest AUC-PR in the cross-validation was applied to the separate validation set. This procedure was repeated 10 times and the mean and standard deviation of the performance metrics were computed accordingly.

2.6. Model Evaluation

Both of the trained models are evaluated against unseen cohorts from the MIMIC-III and UKD datasets. To this end, two independent test cohorts have acted as representatives of unseen data. To quantify the model performance on each of the held-out test cohorts, each of the cropped and uncropped models, which were trained and fit in the validation step, are applied to the independent test cohorts. Then, the performance of the models is quantified and compared in terms of balanced accuracy, AUC-PR, precision, recall and AUC-ROC.

3. Results

In this section, the results of the model validation and evaluation are presented. First, the results obtained by different models in the training and validation phase will be elaborated. Then, the models' performance on the test cohorts will be quantified and compared.

3.1. Model Training and Validation

The results of the model training with five-fold cross-validation on the MIMIC dataset are summarized in Table 3. Both cropped and uncropped models performed well on held-out validation subsets within the training cohort, with AUC-ROCs of 0.877 and 0.859, respectively.

Table 3. Evaluation metrics for the validation set from MIMIC database.

	Balanced Accuracy	Recall	Precision	AUC-PR	AUC-ROC
cropped to 48 h	0.799	0.694	0.389	0.666	0.877
uncropped	0.778	0.649	0.389	0.636	0.859

3.2. Model Evaluation

The results of the evaluation steps on held-out cohorts from the MIMIC and UKD databases are illustrated in Table 4 and Figures 3 and 4, comparing the models' performance together and to an unskilled classifier. In general, we observe a remarkable gap in the performance of the models as applied to the two groups of held-out data. Both the cropped and uncropopped models perform reasonably well on the held-out set from MIMIC, with AUC-ROCs of 0.796 and 0.828, respectively. From the test results of the UKD data, the cropped model (AUC-ROC = 0.554) performed better than the uncropped model (AUC-ROC = 0.517).

Table 4. Evaluation metrics for the test sets from UKD and MIMIC databases.

UKD					
	Balanced Accuracy	Recall	Precision	AUC-PR	AUC-ROC
cropped to 48 h	0.514	0.06	0.318	0.223	0.554
uncropped	0.496	0.22	0.196	0.221	0.517
MIMIC					
	Balanced Accuracy	Recall	Precision	AUC-PR	AUC-ROC
cropped to 48 h uncropped	0.74 0.766	0.575 0.628	0.35 0.477	0.538 0.571	0.796 0.828



Figure 3. Results of the two alternative models' predictions on the unseen cohort from MIMIC dataset: (**left**) the precision–recall curve; (**right**) the receiver operating characteristic curve. The dashed lines denote the performance expected by an unskilled classifier.



Figure 4. Results of the two alternative models' predictions on the unseen dataset from UKD: (**left**) the precision–recall curve; (**right**) the receiver operating characteristic curve. The dashed lines denote the performance expected by an unskilled classifier.

4. Discussion

Predicting the risk of readmission to intensive care units at discharge times through the use of artificial intelligence can support healthcare systems to manage their resources in a more efficient way. As unplanned readmissions might result in longer stays as well as health hazards for patients [1,4], it is critical to provide quantification of the readmission risks as an additional asset for decision making for the intensivists. Related works suggest that different vital and laboratory values of patients sampled in the last 48 h prior to the patient's discharge from the intensive care unit contribute the most to the fact of whether the patient will return to the care unit [21]. Thus, state-of-the-art time series forecasting methods on health-related data from ICU stays could be of great importance for the analysis of readmission risk for patients. Therefore, the motivation behind this study has been to propose and evaluate two alternative models, based on time series forecasting, for readmission prediction for patients visiting cardiovascular ICUs after heart surgery.

The first model analyzes all available vital and lab values for all the participating subjects during their stays at the corresponding ICU. Alternatively, the second model takes into account only the values in the last 48 h prior to discharge from the care unit. As elaborated in the results, both of the alternate models performed reasonably well in predicting patients who would be readmitted to the ICU when they were applied to an unseen cohort from the MIMIC dataset. Furthermore, we showed that the model with the 48 h time limit outperformed the other model in predicting readmissions for a cohort of unseen cases from UKD, which conforms with the findings from related work [21]. The superiority of the model with a persistent number of timestamps can be justified, as,

typically, data-driven ML-based methods perform more efficiently on data cohorts with higher consistency [22]. Moreover, it signifies the relevance of the patients' clinical factors measured in the most recent timestamps prior to transfer for predicting readmission risks.

Overall, the findings reveal the predictive potential of the AI-based approach for health forecasting after evaluating its performance on unseen datasets. However, there are some drawbacks to the currently applied methodology, which should be mitigated in future follow-ups. As for any other time series model, the methodology towards health forecasting as described in this paper would also be affected by the issue of missing data. For this study, we replaced missing measurements with the corresponding mean value throughout the whole cohort. For the future follow-ups, we will consider alternative AI-based imputation methods such as bidirectional recurrent imputation for time series (BRITS) [23], which would most likely help to further improve the robustness of our methods. Furthermore, most of the state-of-the-art ML-based models with deep architectures are less interpretable due to their 'black box' nature. Therefore, one of our future works will be to provide an ML-based approach towards ICU readmission prediction in an explainable manner, i.e., by following explainable AI (XAI) best practices. This will further help to give the medical experts higher levels of interpretability and confidence in AI analyses.

The results of the evaluation step reveal the inconsistency between the time series measurements from the two databases (MIMIC-III and UKD), also known as the dataset shift problem [24]. This might reflect different protocols and standards for patient care and discharge in the corresponding care units and should be diminished in the future using appropriate transfer learning techniques or by training new models on larger cohorts which better fit in protocols. Thus, we are gathering more annotated data to prepare a suitable cohort of ICU stays at UKD for further follow-ups. Nonetheless, even these relatively poor preliminary results on the data from UKD are slightly better than the performance of an unskilled classifier. These findings identify room for improvement in readmission prediction accuracy by the AI model. Although the AI performance might be considered far from perfection, even highly experienced intensivists will not predict ICU readmissions quite perfectly. Therefore, to further quantify how the proposed AI approach would compete with experienced domain experts, we intend to conduct experiments in which the performance of the AI model for ICU readmission prediction is compared to that of a group of experienced intensivists. Nevertheless, it should be noted that the proposed methods are designed as decision support tools for medical experts and therefore should not be used as an autonomous system for decision making.

Another important aspect of any clinical decision support tool is usability. In this regard, as part of the upcoming follow-ups, we intend to conduct proper empirical studies to assess how a provided AI assistant user interface would satisfy user needs in terms of experience and performance. Moreover, to account for better generalizability, the integration of datasets from other hospitals and intensive care units would be a useful next step to take.

5. Conclusions

The aim of the study has been to train and evaluate a deep model with a recurrent architecture to predict the readmission of patients who visited cardiovascular ICUs based on times series data of vital and lab values. To this end, the RNN model based on the LSTM method was trained with open access data and tested using independent hospital data. Our findings revealed the potential of the proposed AI-based methodology to assist domain experts at cardiovascular intensive care units for the better quantification of readmission risks at discharge times. To further address the usability and generalizability aspects of the methods, empirical user studies and experiments with data from other care units should be conducted. Author Contributions: Conceptualization, S.M., S.K. (Sebastian Kalkhoff), S.K. (Steven Kessler), H.A. and F.S.; methodology, S.M., S.K. (Sebastian Kalkhoff), S.K. (Steven Kessler) and Z.B.; software, S.K. (Steven Kessler) and Z.B.; validation, S.K. (Steven Kessler), Z.B. and H.A.; formal analysis, S.M., S.K. (Sebastian Kalkhoff), H.A., F.S. and A.L. (Artur Lichtenberg); resources, A.L. (Artur Liebrecht), R.B., B.D., H.A., F.S., A.L. (Artur Lichtenberg) and A.L. (Artur Lichtenberg); data curation, V.H., S.K. (Steven Kessler), Z.B. and H.A.; writing—original draft preparation, S.M., S.K. (Sebastian Kalkhoff), S.K. (Steven Kessler), Z.B. and H.A.; writing—original draft preparation, S.M., S.K. (Sebastian Kalkhoff), S.K. (Steven Kessler), Z.B., A.L. (Artur Liebrecht), R.B. and B.D.; writing—review and editing, A.L. (Artur Liebrecht), R.B., B.D., A.L. (Artur Lichtenberg), H.A. and F.S.; visualization, S.K. (Steven Kessler) and Z.B.; supervision, S.M., S.K. (Sebastian Kalkhoff), A.L. (Artur Lichtenberg), H.A. and F.S.; project administration, A.L. (Artur Lichtenberg), H.A. and F.S.; funding acquisition, A.L. (Artur Lichtenberg), H.A. and F.S.; funding acquisition of

Funding: The authors disclose receipt of the following financial support for the research, authorship and/or publication of this article: This work was supported by the Federal Ministry of Education and Research of Germany (grant number 16SV8601).

Institutional Review Board Statement: The study was conducted according to the guidelines of the Declaration of Helsinki and approved by the Ethics Committee at the Faculty of Medicine of Heinrich Heine University Düsseldorf (Study No. 2021-1388, 5 May 2021).

Informed Consent Statement: Patient consent was waived due to the retrospective nature of the study.

Data Availability Statement: This study applied an open access dataset (MIMIC-III [18]), which is referenced accordingly, as well as an independent cohort from UKD that, according to German data protection policies, cannot be accessed from outside our facilities.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

UKD	University Hospital Düsseldorf
ICU	Intensive Care Unit
RNN	Recurrent Neural Network
LSTM	Long Short-Term Memory
ML	Machine Learning
AI	Artificial Intelligence
MIMIC	Medical Information Mart for Intensive Care
BRITS	Bidirectional Recurrent Imputation for Time Series
XAI	eXplainable AI

References

- Rosenberg, A.L.; Charles, W. Patients Readmitted to ICUs: A Systematic Review of Risk Factors and Outcomes. *Chest* 2000, 118, 492–502. [CrossRef] [PubMed]
- van Sluisveld, N.; Bakhshi-Raiez, F.; de Keizer, N.; Holman, R.; Wester, G.; Wollersheim, H.; van der Hoven, J.G.; Zegers, M. Variation in rates of ICU readmissions and post-ICU in-hospital mortality and their association with ICU discharge practices. BMC Health Serv. Res. 2017, 17, 281. [CrossRef] [PubMed]
- Desautels, T.; Das, R.; Calvert, J.; Trivedi, M.; Summers, C.; Wales, D.J.; Ercole, A. Prediction of early unplanned intensive care unit readmission in a UK tertiary care hospital: A cross-sectional machine learning approach. *BMJ Open* 2017, 7, e017199. [CrossRef] [PubMed]
- 4. Lin, W.T.; Chen, W.L.; Chao, C.M.; Lai, C.C. The outcomes and prognostic factors of the patients with unplanned intensive care unit readmissions. *Medicine* **2018**, *97*, e11124. [CrossRef] [PubMed]
- Azarfarin, R.; Ashouri, N.; Totonchi, Z.; Bakhshandeh, H.; Yaghoubi, A. Factors influencing prolonged ICU stay after open heart surgery. *Res. Cardiovasc. Med.* 2014, 3, e20159. [CrossRef] [PubMed]
- Almashrafi, A.; Elmontsri, M.; Aylin, P. Systematic review of factors influencing length of stay in ICU after adult cardiac surgery. BMC Health Serv. Res. 2016, 16, 318. [CrossRef] [PubMed]
- Andrearczyk, V.; Fontaine, P.; Oreiller, V.; Castelli, J.; Jreige, M.; Prior, J.O.; Depeursinge, A.: Multi-task Deep Segmentation and Radiomics for Automatic Prognosis in Head and Neck Cancer. In *Predictive Intelligence in Medicine*. *PRIME* 2021. *Lecture Notes in Computer Science*; Rekik, I., Adeli, E., Park, S.H., Schnabel, J., Eds.; Springer: Cham, Switzerland, 2021. [CrossRef]

- Moazemi, S.; Essler, M.; Schultz, T.; Bundschuh, R.A. Predicting Treatment Response in Prostate Cancer Patients Based on Multimodal PET/CT for Clinical Decision Support. In *Multimodal Learning for Clinical Decision Support*. ML-CDS 2021. Lecture Notes in Computer Science; Springer: Cham, Switzerland, 2021. [CrossRef]
- 9. Khatami, M.; Wehler, R.; Schultz, T. Parcellation-Free prediction of task fMRI activations from dMRI tractography. *Med. Image Anal.* **2022**, *76*, 102317. [CrossRef] [PubMed]
- 10. Khan, A.R.; Wang, L.; Beg, M.F. FreeSurfer-Initiated Fully-Automated Subcortical Brain Segmentation in MRI Using Large Deformation Diffeomorphic Metric Mapping. *NeuroImage* **2008**, *41*, 735–746. [CrossRef] [PubMed]
- 11. Hsu, G.C. Using Time-Series and Forecasting to Manage Type 2 Diabetes Conditions (GH-Method: Math-Physical Medicine); Universidad de Granada: Granada, Spain, 2019.
- 12. Spiga, R.; Batton-Hubert, M.; Sarazin, M. Predicting hospital admissions with integer-valued time series. In *ITISE 2019. Proceedings* of *Papers*; ITISE: Grenade, Spain, 2019. pp. 897–998.
- Šimpach, O.; Dotlačilová, P. Age-Specific Death Rates Smoothed by the Gompertz–Makeham Function and Their Application in Projections by Lee–Carter Model. In *Time Series Analysis and Forecasting. Contributions to Statistics*; Rojas, I., Pomares, H., Eds.; Springer: Cham, Switzerland, 2016. [CrossRef]
- 14. Lin, Y.W.; Zhou, Y.; Faghri, F.; Shaw, M.J.; Campbell, R.H. Analysis and prediction of unplanned intensive care unit readmission using recurrent neural networks with long short-term memory. *PLoS ONE* **2019**, *14*, e0218942. [CrossRef] [PubMed]
- Huang, Y.; Talwar, A.; Chatterjee, S.; Rajender, R.A. Application of machine learning in predicting hospital readmissions: A scoping review of the literature. *BMC Med. Res. Methodol.* 2021, 21, 96. [CrossRef] [PubMed]
- Ashfaq, A.; Sant'Anna, A.; Lingman, M.; Nowaczyk, S. Readmission prediction using deep learning on electronic health records. J. Biomed. Inform. 2019, 97, 103256. [CrossRef] [PubMed]
- 17. Kessler, S.; Schroeder, D.; Korlakov, S.; Hettlich, V.; Lichtenberg, A.; Schmid, F.; Aubin, H. Predicting readmission to the cardiovascular ICU using recurrent neural networks. J. Digital. Health 2022, manucript in final decision.
- Johnson, A.; Pollard, T.; Shen, L.; Lehman, L.H.; Feng, M.; Ghassemi, M.; Moody, B.; Szolovits, P.; Celi, L.A.; Mark, R.G. MIMIC-III, a freely accessible critical care database. *Sci. Data* 2016, *3*, 160035. [CrossRef] [PubMed]
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelsshein, N.; Antiga, L.; et al. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In Proceedings of the 33rd Conference on Neural Information Processing Systems (NeurIPS 2019), Vancouver, Canada, 8–14 December 2019; pp. 8024–8035.
- Bauder, R.A.; Khoshgoftaar, T.M.; Hasanin, T. An Empirical Study on Class Rarity in Big Data. In Proceedings of the 2018 17th IEEE International Conference on Machine Learning and Applications (ICMLA), Orlando, FL, USA, 17–20 December 2018; pp. 785–790. [CrossRef]
- Brown, S.E.; Ratcliffe, S.J.; Halpern, S.D. An empirical derivation of the optimal time interval for defining ICU readmissions. *Med. Care* 2013, 51, 706–714. [CrossRef] [PubMed]
- 22. Gudivada, V.; Apon, A.; Ding, J. Data Quality Considerations for Big Data and Machine Learning: Going Beyond Data Cleaning and Transformations. *Int. J. Adv. Softw.* **2017**, *10*, 1–20.
- Cao, W.; Wang, D.; Li, J.; Zhou, H.; Li, L.; Li, Y. BRITS: Bidirectional Recurrent Imputation for Time Series. Adv. Neural Inf. Process. Syst. 2018, 31, 6775–6785.
- 24. Quiñonero-Candela, J.; Sugiyama, M.; Schwaighofer, A.; Lawrence, N.D.; Jordan, M.I.; Dietterich, T.G. Dataset Shift in Machine Learning; MIT Press: Cambridge, MA, USA, 2008; ISBN 978-0-262-17005-5.