

Intelligence-based Recommendation System for Critical Stroke Management in Intensive Care Units

Luis García Terriza¹, José L. Risco-Martín¹, José L. Ayala¹ and Gemma Reig Roselló²

¹ *Department of Computer Architecture and Automation, Universidad Complutense de Madrid, 28040 Madrid, Spain*

² *Stroke Care Unit, Hospital Universitario de La Princesa, Spain*
{luigar05, jlrisco, jayala}@ucm.es, gemmareig@hotmail.com

Keywords: Machine Learning, Genetic Algorithms, Health Recommendation System, Death Risk Prediction, Decision Support System

Abstract: This work presents an integrated recommendation system capable of providing support in healthcare critical environments such as Intensive Care Units or Stroke Care Units using Machine Learning techniques. The system can manage several patients by reading monitoring hemodynamic data in real-time, presenting current death risk probability, and showing recommendations that would reduce such probability and, in some cases, avoid death. This system introduces a novel method to produce recommendations based on genetic models and supervised machine learning. The interface is built upon a web application where clinicians can evaluate recommendations and straightforwardly provide feedback.

1 INTRODUCTION

Stroke is one of the leading death causes in the world. It is a cerebrovascular disease that can produce death and long-term severe disabilities (Wang et al., 2016). There are two stroke sub-types: hemorrhagic and ischaemic. Hemorrhagic strokes are caused by a vessel rupture, while ischaemic strokes are provoked by blood clots occluding brain arteries (Alexopoulos et al., 1999). According to the World Health Organization (WHO), 6.2 million people pass away every day by strokes, the second main cause of death worldwide and the third prominent cause of disability (World Health Organization, 2018).

On the other hand, Artificial Intelligence (AI) based applications are exponentially growing nowadays, used in many areas such as marketing, transportation, agriculture, education, medicine, etc. Regarding medicine and healthcare, it is stated in (Jiang et al., 2017) that clinicians might be replaced by AI systems at some point, although this will not happen in the short term. However, AI-based recommendation systems are gaining trust as assistance methods to clinicians to help them make better decisions and, therefore, improve the possible outcome of the patient (Asan et al., 2020).

In this work, we propose a real-time recommendation system that can be used in an Intensive Care Unit (ICU) or a stroke care unit, designed to reduce

the death risk probability and, in some cases, avoid death. Our framework has been conceived for patients admitted to the hospital within the first 48 hours after suffering a stroke attack. The AI behind our system uses hemodynamic data to produce a death prediction model along with a real-time recommendation system that can predict current exitus (death) risk and recommend new possible patients' states that would not lead to exitus. These "states" are presented as recommendations to the clinicians in the form of possible hemodynamic values for the patient that should potentially reduce the exitus risk. Obviously, the clinic staff must carefully examine and validate this new state.

To build the predictive model, the AI module is fed with data collected from a Philips PIIC iX monitoring center installed in the stroke care unit of the Hospital Universitario La Princesa since 2017, where around 800 patients have been successfully monitored. Patients' data have been stored in a semi-structured SQL Database so it can be pre-processed, modeled, and read in real-time by this recommendation system. This paper focuses on describing the recommendation system as a high-level application deployed at the stroke care unit mentioned above. Intrinsic technical details about the predictive model definition are detailed in (García-Terriza et al., 2021).

The paper is organized as follows. Section 2 discusses the related work. Section 3 presents the high-level architecture and design of the recommendation

system. Next, Section 4 illustrates the simulations performed to test our hypotheses and shows the results obtained for different patients. Finally, Section 5 draws some conclusions and introduces future work.

2 RELATED WORK

There is an increasing amount of research work where diverse artificial intelligence techniques and models are considered in critical health environments, assisting clinicians in evaluating the current health status of the patient or recommending possible treatments, taking into account the patient's history. In the following, we mention some relevant contributions.

(Moghadam et al., 2020) aim to predict hypotension events using a Logistic Regression model in an ICU environment. With a minimum of 5 minutes of physiological data monitoring, the system outputs the eventual risk of a hypo-tension crisis in the short term. In (Nemati et al., 2018), the authors developed a framework based on artificial intelligence to predict a sepsis crisis event in advance from 4 to 12 hours. Using real-time monitoring from several hospitals and a Weibull-Cox proportional hazard model, they obtained an accuracy between 56% and 72%, depending on the temporal window and the experiment. In the same line, (Lukaszewski et al., 2008) attempted to predict sepsis in patients who received surgery and were admitted to ICU afterward. Using data from daily blood analysis and PCR markers from 92 patients, they trained a neural network model, obtaining an average accuracy of 83% in a range from 1 to 4 days before clinical diagnosis.

Closer to our research, other studies try not only to predict critical events but also to provide or offer recommendations to clinicians. To compare and analyze the current state of the art with our approach, we have included Table 1, where similar works regarding recommendation systems in ICU environments are presented. The comparison has been performed taking into account four main aspects: (i) the system is deployed in ICU, (ii) the current health status of the patient is an input to the system, (iii) the system has real-time capabilities, and finally, (iv) the technique or algorithm used in the study.

As Table 1 shows, (Utomo et al., 2019) train a Bayesian model using Reinforcement Learning to predict the best treatment for the patient in the ICU. The resulting model is used in real-time in ICU using patient monitoring data as input for the system. (Neloy et al., 2019) perform classifications of critical patients in ICU by using association rule mining and K-Nearest Neighbors to reduce patient mortal-

ity risk. As the system is deployed on a Cloud environment, the system is capable of being retrained while the system is running and replacing the previous model with the new one. In (Chen et al., 2016), the authors propose an approach capable of providing recommendations based on the Electronic Medical Record (EMR) upon the patient's admission. Recommendations are not based on the ongoing status of the patient nor their current health status but their historical record. (Masud et al.,) recommend medications for critical care patients by using the first 24 hours of monitoring data and comparing them with the historical patient database. This work produces, therefore, a one-shot recommendation. (Varatharajah et al., 2020) propose a recommendation system based on reinforcement learning to support the clinical management of COVID-19 patients. The system predicts the severity of the COVID-19-infected patients and then helps with their treatment. Finally, (Thong et al., 2015) aim to diagnose patients' disease from given symptoms. They have used intuitionistic fuzzy sets and a recommendation system to achieve such a purpose, improving standalone methods of state of the art.

Our work proposes a real-time recommendation system capable of predicting current death risk and providing recommendations that would reduce such risk. We achieve that using (i) real-time data monitored by the patient, (ii) a trained machine learning model, and (iii) a genetic algorithm. In terms of capability to track and advance the health status of the patient, our study, along with (Utomo et al., 2019), and (Neloy et al., 2019), are the most complete systems as they are capable not only of operating in an ICU environment but also to read patient's monitoring data in real time and provide recommendations about the control variables, each one on its corresponding area. The rest of the works presented lack this real-time component which is fundamental to assist the clinicians at every moment while the patient is admitted to the hospital. Another relevant advantage of our method is the minimum required time to produce recommendations or diagnoses. While (Masud et al.,) require patient monitoring for the first 24 hours, our work, on the contrary, needs a minimum monitoring time of 2 minutes and a half. Finally, previous works have not addressed the complex process of dealing with real-time continuous hemodynamic signals that present artifacts and produce vast data. At the same time, our recommendation system can manage hemodynamic data from hospitalized patients.

Table 1: Definition of hemodynamic variables used. The acronyms used are: RL (Reinforcement Learning), KNN (K-Nearest Neighbors), DM (Data Mining), NB (Naive Bayes), FL (Fuzzy Logic), RF (Random Forest), and GA (Genetic Algorithm)

Reference	System in ICU	Patient's data ingestion	Real-time	Algorithm
Utomo <i>et al.</i>	Yes	Yes	Yes	RL
Neloy <i>et al.</i>	Yes	Yes	Yes	KNN
Chen <i>et al.</i>	No	No	No (one-shot)	DM + NB
Masud <i>et al.</i>	Yes	Yes (24h)	No	KNN
Varatharajah <i>et al.</i>	No	No	No (one-shot)	RL
Thong <i>et al.</i>	No	No	No	FL
Our Work	Yes	Yes	Yes	RF + GA

3 SYSTEM ARCHITECTURE AND DESIGN

The recommendation system presented in this work uses a death risk model based on machine learning techniques that have been proved to predict patients' exitus probability with an F-Score performance above 98%, tested within the first 48 hours after patient's admission (García-Terriza *et al.*, 2021). It is capable of predicting current death risks in real time for every patient who is being monitored. Our method uses a genetic algorithm to find the pseudo-optimal state of the patient's hemodynamic system, reducing the exitus risk. The genetic algorithm evaluates different states through a patient in-silico model defined with machine learning techniques described in the following subsections. And the last part of the recommendation system is the presentation to the clinicians, which is done through a web interface where clinicians may review the current patient's status, the death risk, and the recommendations, so the death risk is reduced.

Figure 1 depicts our system architecture. Firstly, hemodynamic data is consumed by the recommendation system, which is connected to a SQL Database where the monitoring center Philips PIIC iX stores patients' monitored data in real time. Secondly, the recommendation system publishes a web interface with two main purposes. On the one hand, it provides clinicians with an interface to add static variables of the patient. On the other hand, a visual interface shows current risks and recommendations for each patient admitted to the stroke care unit. These components are detailed in the following subsections.

3.1 Model

As mentioned before, the recommendation system uses a machine learning model to predict the patient's current death risk based on the patient's hemodynamic variables at that moment. The machine learning model has been trained from the patient's hemodynamic data and static features. The set of Hemo-

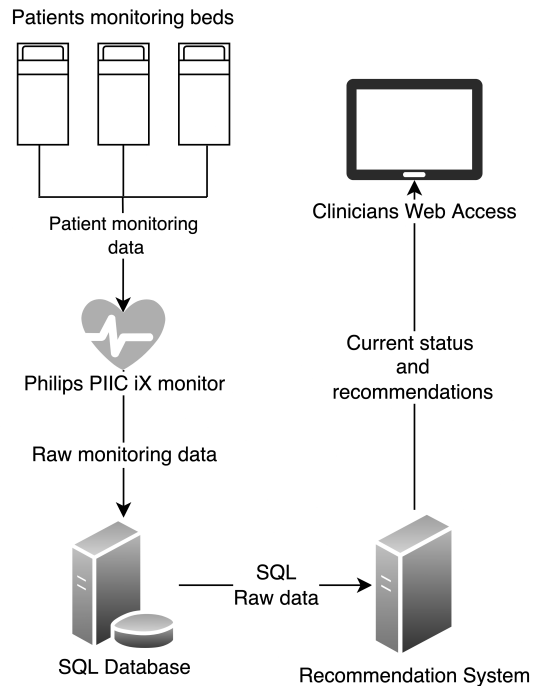


Figure 1: High-Level system architecture

dynamic variables are described in Table 2, whereas the static features are shown in Table 3. The categorical input variables were coded as numerical variables, so the model input is suitable for machine learning and deep learning algorithms. Hemodynamic data may present inaccurate measurements as it is gathered from sensors. Hence data is preprocessed before using it for training/testing or production purposes.

Preprocessing is compounded of three stages: (1) missing values filling, (2) outlier values removal, and (3) standardization. The first stage substitutes missing values (null values) with the mean of that same patient and variable. The second stage replaces values by the mean when those values are greater or equal to four times the standard deviation. Finally, the third preprocessing stage consists of applying the Z-Score standardization to the input values of the model. This last stage is especially relevant as variables compre-

hend different orders of magnitude, which might lead to sub-optimal training.

Algorithm selection is one of the key points of the study. One of our objectives was to find the best type of machine learning or deep-learning algorithm for this problem. We aim to predict the patient’s exitus based on the pair of hemodynamic data and static features. These were the algorithms tested:

- Logistic Regression
- Naive Bayes
- Support Vector Machines (SVM-SVC)
- Tree-based algorithms: Decision Trees (CART), Random Forests, and Gradient Boosting Trees
- Distance-based algorithms: K-Nearest Neighbors, Dynamic Time Warping + 1-NN.
- Multi-Layer Perceptron Neural Network
- Deep Neural networks: 1D Convolutional Neural Network (CNN 1-D), Long-Short Term Memory (LSTM)

Out of all these algorithms, the ones that performed the best were the tree-based algorithms, followed by the Deep Neural Networks. The final algorithm chosen, and the one used in this recommendation system was the **Gradient Boosting Trees**, which performed the highest F1-Score, a 98.5%. Every training and testing procedure was done using the K-Fold cross-validation method with K=5 to avoid overfitting and independent results from training and testing datasets.

As input data used in this work involves a particularly relevant time component, time windows while training models became a bottom-line condition. After training a large variety of models within different time windows, we found the approach of training one model for each hour since the patient’s admission as the most feasible on both computational and performance aspects. Thus, every model would be most accurate in its specific time window.

For further details of the data preprocessing, training/testing process, and results, see (García-Terriza et al., 2021).

3.2 Recommendation Algorithm

The most challenging part of the system was finding the patient’s clinical state to minimize the death risk probability. As described in Table 2, seven hemodynamic variables are continuously monitored and are candidates to be adjusted, so the patient’s exitus odds are reduced. Hence, our system must search for valid combinations, from a clinical point of view, of those seven physiologically relevant variables that would

not lead to a patient’s exitus in a short-term period. The number of combinations of seven numerical variables is not computationally feasible in near-real-time situations, so we figured out an algorithm to search the solutions space and find possible values for each patient. To this end, we have designed a **genetic algorithm-based** method that offers possible combinations that meet the clinical requirements.

A genetic individual is comprised of the seven hemodynamic variables monitored from the patient and loaded in the exitus risk model. The initial population is created by producing random values for each variable in every individual. Random values generation is controlled within their respective physiological limits. For instance, it makes no sense to generate an individual composed of cardiac frequency above 220 bpm (beats per minute) or oxygen saturation level less than 80% or greater than 100%. These conditions are also fulfilled at the gene mutation phase. The most relevant part of a genetic algorithm is often the fitness function, which must evaluate individuals objectively and quantify them. In our case, the fitness function is defined as the combination of probabilities from several exitus models since we have computed different models for different temporal monitoring windows. Altogether provide a probability that must be minimized and ensure that, according to all the models, the individual has a p probability of becoming an exitus defined as:

$$p = \prod_{i=1}^n (1 - m_i) \quad (1)$$

where m_i is the death probability given by the i -th death risk model between 0.0 and 1.0 and n is the number of models (or temporal monitoring windows) included in the system. As the genetic model tries to maximize the fitness function, the death probabilities are subtracted from 1.

Our genetic algorithm uses a one-point crossover function, where a point between the seven variables is randomly chosen for pairs of individuals so that the left part of the first individual is merged with the right part of the second individual. The right part of the first individual is merged with the left part of the second one, forming two children. Elitism is applied to ensure that the best individuals are not lost between generations by preserving a reduced group of the best individuals after each generation.

All the individuals that have survived at the end provide several states of the patient that, in turn, define a set of acceptable solutions to our problem, i.e., a combination of seven hemodynamic values that minimize the death probability of a patient in near-real-time situations. These solutions generated by the sys-

Table 2: Definition of hemodynamic variables used

Variable	Abbreviation	Type	Role	Description
Rhythm Estimation	RE	Categorical	Input	A Rhythm Indicator
VE	VE	Numerical	Input	Ventricular extra systole
CF	CF	Numerical	Input	Cardiac Frequency
Breathing Frequency	BF	Numerical	Input	Respiratory Rate
Perf	Perf	Numerical	Input	Pulmonary Perfusion
SpO2	SpO2	Numerical	Input	Oxygen Saturation
ST-II	ST	Numerical	Input	Syst. Time Interval Index

Table 3: Definition of patient’s static variables used

Variable	Abbreviation	Type	Role	Description
Age	Age	Categorical	Input	Ischemic/Hemorrhagic
Gender	Gender	Categorical	Input	Male/Female
Monitoring time	MT	Numerical	Input	Monitoring observation number
Type of Stroke	TS	Categorical	Input	Ischemic/Hemorrhagic
Risk Prediction	RP	Numerical	Output	Exitus probability

Patient 9858406

Figure 2: Patient features input form

tem are the recommendations offered to the clinicians through the web interface.

3.3 Visual Interface

To end this Section, we will describe the web interface and the defined action protocol for clinicians.

As mentioned in previous subsections, a set of static variables is not available automatically. Thus, they must be added to the system manually by clinicians every time a patient is admitted to the stroke care unit. Figure 2 presents the web interface where the patient’s age, gender, and stroke-subtype features can be inserted. Age is limited to integer numbers between 20 and 100. Gender is delimited to biological genders, male and female. The stroke type is categorized into hemorrhagic, ischaemic, and stroke mimics (this last one refers to events initially diagnosed as ischaemic stroke with non-cerebrovascular causes). In Figure 2, the number “9858406” identifies the patient’s MRN (Medical Record Number), which is automatically provided by the registration system.

Lastly, we have defined an action protocol for clin-

icians comprised of a set of repeating steps for every patient. After an acute stroke, the patient is admitted to the stroke care unit, where they are signed up in the infirmary and begin to be monitored. The patient will start showing up on the recommendation system web interface, where the patient features must be inserted. At this point, the clinicians must intervene and add the patient’s features. Since this process is manual, the clinicians must check the patient’s dashboard at least once per day. Once features have been inserted, the system will start providing the current risk and recommendations. Finally, when recommendations are available, the clinician can analyze the suitability of these recommendations from a medical point of view and provide positive feedback. This feedback will be used in the future to fine-tune the recommendation systems and improve the clinician’s satisfaction.

4 EXPERIMENTS

Our experimental work has processed patients with acute stroke monitored at the stroke unit of the Hospital Universitario de La Princesa (Madrid, Spain). Inclusion criteria were: all patients with acute stroke admitted to the stroke unit of the hospital and who are susceptible to noninvasive multiparameter monitoring, according to the clinical protocols in force in the aforementioned unit. Exclusion criteria were: patients admitted to the stroke unit for scheduled procedures or other processes other than the acute phase of stroke; no possibility of monitoring (technical problems, damaged monitoring equipment, patient intolerance/agitation); loss of data due to error in patient coding; poor quality of data recording.

In this section, we will first introduce experiments from a technical and computational point of view, and finally, we will present the system's web interface.

Our experimental work has covered a broad spectrum of patient features (age, stroke type, and gender) to validate the accuracy of the recommendation system. However, due to space limits, we will present cases of study aggregating by age ranges of 10 years, combining with the two-stroke subtypes, ischaemic and hemorrhagic. Another combination would be the patient's gender. Still, to keep the table size down, we have decided to aggregate by mean, as there are no relevant variations between genders.

All the experiments were run using the same genetic model parameters and the same set of death risk prediction models for every case. The initial population is 20 individuals, the individual mutation probability is 20%, and the gene mutation probability, evaluated for each gene compounding the individual, is 10%. The cross probability between two individuals is 50% using the one-point cross method, which is selected randomly between the existing possible cross-points. Finally, the genetic model is executed with a maximum of 20 generations, defined after several tests as sufficient to reach a stable solution. There is elitism where the three best individuals are preserved between generations.

Table 4 shows aggregated results from our experiments for each age range and stroke sub-type (hemorrhagic and ischaemic). As it may be seen, the optimal solution found by the algorithm (i.e., the best individual found) presents no exitus (death) risk likelihood in almost every case, according to the models used in the genetic algorithm. As age rises, the genetic algorithm takes more generations to find optimal solutions, which seems reasonable as the exitus risk grows proportionally with age. Therefore solutions are harder to be found if they exist. The results also show that the stroke sub-type is a relevant feature of the study case, as exitus risk is higher in hemorrhagic cases than ischaemic, making it harder for the algorithm to find optimal solutions. Typically, the optimal solution is found within the first five generations when the patient is younger than 70 years or suffers an ischaemic stroke.

In terms of computation, as individuals are not compounded of a significant number of genes, the population size is relatively reduced (20 individuals). Usually, optimal solutions are found within the first five generations. Simulations are computed fast enough to provide results in a short-term period that offers the recommendation system near real-time characteristics. As stated before, the solutions found by the genetic algorithm are, at last, the recommenda-

tions provided to clinicians.

A web interface displays the formerly presented results. The interface must provide information about the patient's current status, clinical recommendations to improve such status, and a mechanism to send feedback about the recommendations. Figure 3 depicts the all-patients dashboard where previous requirements are satisfied. First, the static feature insertion process required for every patient on admission may be spotted in (1), stating "Insert patient information". Regarding the patient risk information, in (2) and (3), the real-time exitus likelihood and the historical accumulated exitus likelihood of the patient are displayed, respectively. The main goal of showing the historical risk is to provide a statistical marker that can be used to compare with the current risk and evaluate the evolution of the patient. Next, to see the recommendations generated for a patient by the system, the clinician should click on (4). A drop-down appears depicting a table where the last observation (5) and recommendations (6) are displayed. Aiming to reduce the number of recommendations, we decided to limit it to 3 unique recommendations at maximum. When the clinician considers the recommendation suitable for medical criteria, the clinician must click on (7), changing the icon's color as shown in (8). Finally, if the patient has been monitored for less time than the predefined temporal window or the input monitoring data are incomplete, the system shows a message in the interface stating, "Wait a few more seconds for results".

5 CONCLUSION AND FUTURE WORK

This study has demonstrated that our recommendation system can diagnose death risk probability in real-time for hospitalized patients and provide recommendations in near real-time, which may be of assistance to clinicians. Using genetic algorithms, we have developed a model that can find suitable solutions for each patient within a short period of time. The system's recommendations may help clinicians decrease the outcomes of the stroke in ICUs or stroke care units, and, in some cases, those recommendations may help save the patient's life. In terms of architecture design, this system is quite flexible as its core is based on a bundled machine learning model, which may be changed by a different model built of almost any other underlying technology such as, for instance, deep learning or reinforcement learning, without changing the overall architecture. This could lead to enhancements in the model section in

Table 4: Experiments results

Age Range	Stroke sub-type	Generations until best individual	Average risk	Best individual risk
[40,50)	Ischaemic	2	0.29 %	0.0 %
[40,50)	Hemorrhagic	3	0.41 %	0.0005 %
[50,60)	Ischaemic	1	0.5 %	0.0 %
[50,60)	Hemorrhagic	3	0.55 %	0.0001 %
[60,70)	Ischaemic	2	0.748 %	0.0 %
[60,70)	Hemorrhagic	3	0.0085 %	0.0001 %
[70,80)	Ischaemic	4	0.5 %	0.04 %
[70,80)	Hemorrhagic	20	36.3 %	34.1 %
[80,90)	Ischaemic	3	0.32 %	0.075 %
[80,90)	Hemorrhagic	20	33.8 %	26.5 %

Exitus probability

#	NHC	Exitus probability (%)	Mean Accumulated exitus probability (%)
1	647298 <input type="button" value="Edit"/>	Wait a few more seconds for results	N/A
2	587229 <input type="button" value="Edit"/>	0.00%	0.00%
3	9858406 <input type="button" value="Edit"/>	Insert patient information ⁽¹⁾	N/A
4	143127 <input type="button" value="Edit"/>	0.05%	0.04%
5 ⁽⁴⁾	9971526 <input type="button" value="Edit"/>	99.49% ⁽²⁾	86.84% ⁽³⁾

#	EV	FR	FC	SPO2	Rhythm Estimation	ST-II	Perfusion	Feedback
Last observation ⁽⁵⁾	0	23	66	92	Ritmo SV	0.2	1.8	
Recommendation #1 - (keep for 2.5 mins) ⁽⁶⁾	0	17	67	96	Ritmo sinusal	-0.57	4.34	<input type="button" value="Feedback"/> ⁽⁸⁾
Recommendation #2 - (keep for 2.5 mins)	0	17	61	96	Ritmo sinusal	-0.57	4.31	<input type="button" value="Feedback"/> ⁽⁷⁾
Recommendation #3 - (keep for 2.5 mins)	0	17	67	96	Ritmo sinusal	-0.57	4.31	<input type="button" value="Feedback"/>

Figure 3: Web patient dashboard view

the future, which is crucial to reduce outcomes and mortality. Further improvements in architecture and computation capabilities would imply implementing the architecture on the (Public) Cloud, adding new parts to the architecture, and decoupling some components. A Cloud-based architecture would reduce the maintenance load and make more accessible the automated deployment of new retrained models based on the feedback and real-time results.

Our next step consists of deploying the recommendation system in the stroke care unit of the Hospital Universitario La Princesa. This installation comprises the connection to the production SQL Database so that the system can read the patient's hemodynamic data in real-time and the web publishing within the hospital networking systems. Furthermore, the clinician feedback through the web must be stored on a persistent device that should be easily accessible, so reports are available. The clinician information will be analyzed to improve the recommendation system procedure by using clustering and segregation of the

different patient clusters.

Acknowledgements

This research has been funded by Instituto de Salud Carlos III (RICORS-RD21/0006/0009) and co-financed with FEDER Funds and/or from the European funds of the Recovery, Transformation and Resilience Plan and by NextGenerationEU. This work is also supported by Spanish Ministry of Science and Innovation under project PID2019-110866RB-I00.

REFERENCES

Alexopoulos, E., Dounias, G. D., and Vemmos, K. (1999). Medical diagnosis of stroke using inductive machine learning. In *Machine learning and applications*, pages 20–23.

- Asan, O., Bayrak, A. E., Choudhury, A., et al. (2020). Artificial intelligence and human trust in healthcare: focus on clinicians. *Journal of medical Internet research*, 22(6):e15154.
- Chen, J. H., Podchiyska, T., and Altman, R. B. (2016). Orderrex: clinical order decision support and outcome predictions by data-mining electronic medical records. *Journal of the American Medical Informatics Association*, 23(2):339–348.
- García-Terriza, L., Risco-Martín, J. L., Roselló, G. R., and Ayala, J. L. (2021). Predictive and diagnosis models of stroke from hemodynamic signal monitoring. *Medical & Biological Engineering & Computing*, 59(6):1325–1337.
- Jiang, F., Jiang, Y., Zhi, H., Dong, Y., Li, H., Ma, S., Wang, Y., Dong, Q., Shen, H., and Wang, Y. (2017). Artificial intelligence in healthcare: past, present and future. *Stroke and Vascular Neurology*, 2(4):230–243.
- Lukaszewski, R. A., Yates, A. M., Jackson, M. C., Swingler, K., Scherer, J. M., Simpson, A., Sadler, P., McQuillan, P., Titball, R. W., Brooks, T. J., et al. (2008). Presymptomatic prediction of sepsis in intensive care unit patients. *Clinical and Vaccine Immunology*, 15(7):1089–1094.
- Masud, M. M., Cheratta, M., and Hayawi, K. Medication recommendation for critical care patients using patient similarity in clinical records.
- Moghadam, M. C., Abad, E. M. K., Bagherzadeh, N., Ram-singh, D., Li, G.-P., and Kain, Z. N. (2020). A machine-learning approach to predicting hypotensive events in icu settings. *Computers in Biology and Medicine*, 118:103626.
- Neloy, A. A., Shafayat Oshman, M., Islam, M., Hossain, M., Zahir, Z. B., et al. (2019). Content-based health recommender system for icu patient. In *International Conference on Multi-disciplinary Trends in Artificial Intelligence*, pages 229–237. Springer.
- Nemati, S., Holder, A., Razmi, F., Stanley, M. D., Clifford, G. D., and Buchman, T. G. (2018). An interpretable machine learning model for accurate prediction of sepsis in the icu. *Critical care medicine*, 46(4):547.
- Thong, N. T. et al. (2015). Intuitionistic fuzzy recommender systems: an effective tool for medical diagnosis. *Knowledge-Based Systems*, 74:133–150.
- Utomo, C. P., Kurniawati, H., Li, X., and Pokharel, S. (2019). Personalised medicine in critical care using bayesian reinforcement learning. In *International Conference on Advanced Data Mining and Applications*, pages 648–657. Springer.
- Varatharajah, Y., Chen, H., Trotter, A., and Iyer, R. K. (2020). A dynamic human-in-the-loop recommender system for evidence-based clinical staging of covid-19. In *HealthRecSys@ RecSys*, pages 21–22.
- Wang, H., Naghavi, M., Allen, C., Barber, R. M., Bhutta, Z. A., Carter, A., Casey, D. C., Charlson, F. J., Chen, A. Z., Coates, M. M., et al. (2016). Global, regional, and national life expectancy, all-cause mortality, and cause-specific mortality for 249 causes of death, 1980–2015: a systematic analysis for the global burden of disease study 2015. *The lancet*, 388(10053):1459–1544.
- World Health Organization (2018). The top 10 causes of death.