

BRAIN. Broad Research in Artificial Intelligence and Neuroscience

e-ISSN: 2067-3957 | p-ISSN: 2068-0473

Covered in: Web of Science (ESCI); EBSCO; JERIH PLUS (hkdir.no); IndexCopernicus; Google Scholar; SHERPA/RoMEO; ArticleReach Direct; WorldCat; CrossRef; Peeref; Bridge of Knowledge (mostwiedzy.pl); abcdindex.com; Editage; Ingenta Connect Publication; OALib; scite.ai; Scholar9; Scientific and Technical Information Portal; FID Move; ADVANCED SCIENCES INDEX (European Science

Evaluation Center, neredataltics.org); ivySCI; exaly.com; Journal Selector Tool (letpub.com); Citefactor.org; fatcat!; ZDB catalogue; Catalogue SUDOC (abes.fr); OpenAlex; Wikidata; The ISSN Portal; Socolar; KVK-Volltitel (kit.edu) 2025, Volume 16, Special Issue 1 (April 2025), pages: 87-98.

Special Issue 1: *Neuroscience, Artificial Intelligence, and Innovation in Education*

Submitted: January 23rd, 2025 | Accepted for publication: March 17th, 2025

Performance Assessment of Neural Networks in Medical Treatment Cost Estimation

Kristijan Cincar

West University of Timișoara, Faculty of
Mathematics and Computer Science,
Romania.

Tibiscus University of Timișoara,
Romania.

SC Nikola Tesla, Vrsac, Serbia.

kristijan.cincar80@e-uvt.ro

<https://orcid.org/0000-0002-5740-6680>

Andrea Amalia Minda*

Babeș-Bolyai University of Cluj Napoca,
Engineering Faculty, Romania.

andrea.minda@ubbcluj.ro

<https://orcid.org/0000-0001-7219-827X>

Marius Susa

Petroleum – Gas University of Ploiești,
Systems Engineering, Romania.

marius.susa@upg-ploiesti.ro

<https://orcid.org/0009-0005-1500-0542>

*Corresponding author

Abstract: *This study explores the use of simulations to analyse the performance of predictive neural network models applied in healthcare, specifically in estimating medical treatment costs. The simulation uses synthetic data generated for 30 days, with variables such as treatment cost, disease severity, and patient waiting time. Three neural network models — Feedforward NN, LSTM NN, and 1D CNN — are employed for cost prediction, and their performance is evaluated using the Mean Absolute Error (MAE), Mean Squared Error (MSE) and Root Mean Squared Error (RMSE) metrics. In addition to numerical evaluation, a series of detailed visualizations are created to analyze the relationships between variables and to compare the accuracy of each model. The results indicate that the 1D CNN achieves the best performance, demonstrating its potential for real-life clinical applications. The study also highlights the utility of simulations for testing predictive methodologies in fields where real-world data may be limited.*

Keywords: *hospital scheduling; deep neural learning; patient management; medical treatment cost.*

How to cite: Cincar, K., Minda, A.A., & Susa, M. (2025). Performance assessment of neural networks in medical treatment cost estimation. *BRAIN. Broad Research in Artificial Intelligence and Neuroscience*, 16(Sup1), 87-98. <https://doi.org/10.70594/brain/16.S1/8>



1. Introduction

Estimating the costs of treatments in hospitals is a major challenge in the management of healthcare resources. Ideally, predictive algorithms can help anticipate these costs, improving resource allocation and reducing inefficiencies. However, the development and testing of these algorithms require complex and well-structured data, which in practice can be difficult to obtain due to confidentiality and inter-institutional variations.

The Emergency Department (ED) is a crucial section in hospitals, as it addresses the needs of patients with critical and unique requirements. These patients often face more complex conditions compared to those in other hospital units. Consequently, researchers have explored various approaches to tackle scheduling challenges in EDs. Azadeh et al. (2014) utilised Genetic Algorithms to optimise scheduling in emergency departments aiming to minimise delays and ensure high-quality patient care. Other researchers implemented methods like ‘Physician-In-Triage’ and ‘the floating patient,’ employing simulation models and Karp-based algorithms to enhance scheduling efficiency (Elalouf & Wachtel, 2015). Kiris et al. (2010) proposed a scheduling system prioritising critical patients to prevent complications or fatalities while managing the waiting times for less urgent cases. Similarly, Carter and Lapierre (2001) investigated EDs in six major Canadian hospitals, identifying key issues through interviews. They applied three precise scheduling techniques—cyclic scheduling, non-rotational cycle scheduling, and rotational cycle scheduling—and compared their results with computerised scheduling systems used in two hospitals.

Efficient hospital planning requires better coordination between doctors, hospital resources, care staff, and patients. The study of Cincar (2020) presents a multi-agent system designed to optimise hospital scheduling by considering factors such as uncertain service duration, patient arrival time, waiting time, and overall treatment time. The system analyses various hospital departments, including hypertension, diabetes, and alcoholism, where each appointment contributes to prioritisation, aiming to achieve cost-effective scheduling and resource allocation.

Ali, Lamsali, and Othman (2018) examined critical components of hospital operations, including nursing, operating rooms, emergency departments, and outpatient services. The research highlighted existing gaps in the field and suggested areas for future investigation and aims to enhance the understanding of hospital scheduling practices and provide valuable insights for healthcare administrators to improve operational efficiency. The research by Yanamala and Suryadevara (2022) employs deep learning methods to forecast hospital readmissions, tackling the significant challenge of imbalanced prediction errors in healthcare analytics. By combining machine learning approaches with healthcare data, the study enhances predictive precision and supports more efficient resource management in clinical environments.

The paper of Sharafat and Bayati (2021) introduces *PatientFlowNet*, a deep learning-based approach that improves patient flow prediction in emergency departments by leveraging both historical and real-time data. Unlike traditional methods, it captures complex temporal patterns and dependencies, enhancing accuracy and enabling better resource allocation and decision-making. A hybrid model combining Long Short-Term Memory (LSTM) networks and decision trees is presented in the paper of Abatal et al. (2025) to optimise patient volume predictions in emergency departments. By integrating deep learning with interpretable machine learning techniques, the approach enhances forecasting accuracy and supports better resource management and decision-making in healthcare settings.

The study of Cincar, Minda, and Varga (2024) focused on integrating discrete event simulation with machine learning models—Decision Trees, Random Forests, Support Vector Machines, and Gradient Boosting—to predict treatment costs and appointment availability. It assessed model performance based on patient throughput, cumulative costs, and cost-effectiveness, demonstrating how advanced analytics can enhance hospital resource management. In this paper, we aim to shift the focus towards predictive neural network models, specifically Feedforward NN, LSTM NN, and 1D CNN, for estimating medical treatment costs. While still using synthetic data

over 30 days, this research introduces a more advanced modeling approach, leveraging deep learning techniques. Unlike the previous study, which compared various traditional machine learning models, this study evaluates deep learning architectures using MAE, MSE, and RMSE metrics. Moreover, it incorporates detailed visualisations to examine variable relationships and model accuracy, offering a more granular analysis of prediction performance.

2. Models and Results

Simulations based on synthetic data offer a powerful alternative. This study uses a simulation methodology to generate data that reflects real scenarios in a hospital. The simulation includes key variables such as the cost of treatment, severity of the disease, and waiting time. The cost of treatment reflects the diversity of medical interventions, shaped by a normal distribution. The severity of the disease is rated on a scale of 1 to 10, simulating the severity of the conditions. The waiting time is a variable that captures the time it takes for patients to receive care, which is important for analysing the efficiency of the system. Schematically, the patient flow can be presented as in Figure 1.




Figure 1. Patient flow

To test the accuracy of the predictions, three types of neural networks were implemented: Feedforward NN, long short-term memory LSTM NN, and Convolutional Neural Networks CNN 1D.

The Feedforward Neural Network model, being a densely connected neural network without temporal structure, is suitable for tabular data without sequential dependencies. In this case, because it does not use the temporal relationships in the disease severity history, the model may have a poorer performance (higher root mean square error, RMSE) than LSTM and 1D CNN, which are optimised for sequential data. The LSTM network excels at capturing temporal relationships and long-term dependencies in historical data series. This allows the model to “learn” from the evolution of the patient’s health status over the past few days and make more accurate predictions of

treatment costs. Therefore, the LSTM tends to have a lower RMSE compared to the Feedforward network, indicating a superior performance in terms of accuracy. 1D CNN uses a different approach to process time series of disease severity. Instead of analysing the data in series like LSTM, 1D CNN applies convolutions to “filter” and identify local features within the time history. Although 1D CNN does not capture long-term dependencies as well as LSTM, it can still achieve good prediction performance (relatively low RMSE) by extracting relevant features from the data series.

Each model is evaluated using the MAE, MSE, and RMSE metrics. They provide a measure of the deviation of the predictions from the real values, the MFA being more intuitive and the MSE penalising large errors more severely. Mean Absolute Error (MAE) and Mean Squared Error (MSE) are commonly used metrics for evaluating the performance of regression models. MAE measures the average absolute differences between predicted and actual values, providing a straightforward interpretation of error magnitude. MSE calculates the average of the squared differences between predicted and actual values, giving more weight to larger errors due to squaring. Root Mean Squared Error (RMSE) plays a crucial role in providing a more comprehensive evaluation of prediction accuracy in medical cost estimation. RMSE helps measure the magnitude of prediction errors while giving more weight to larger errors. Since cost prediction errors can lead to substantial financial misestimations, RMSE ensures that models minimising high-cost deviations are prioritised.

The study aims to generate a realistic set of synthetic data to simulate hospital operations, implement and compare three neural network architectures for predicting treatment costs, and explore relationships between variables through graphical visualisations.

Data were generated for a 30-day period between 7:00 am and 2:00 pm to simulate patient flow in a limited-hours hospital. The main parameters are: the number of patients seen per hour, cost of treatment, severity of illness, and waiting time.

- The number of patients per hour is modeled using Poisson distribution with a mean of 5 patients/hour. This distribution reflects variability in patient flow.
- Treatment cost is modeled using a normal distribution with a mean of 2000 RON and a standard deviation of 500 RON. The values reflect the diversity of costs associated with medical interventions.
- Disease severity was uniformly generated between 1 and 10, representing a spectrum of severity.
- The waiting time shows a uniform distribution between 5 and 30 minutes to simulate common delays in the healthcare system.

This data was stored in a Pandas DataFrame, allowing further analysis to be easily managed and visualised.

The Python code developed simulates patient management in a hospital setting using three deep learning models: Feedforward NN, LSTM NN, and CNN 1D to predict treatment costs and optimise scheduling. The code simulates an analysis of medical data and tests the performance of prediction models in a clinical context. Results are visualised to highlight relationships between parameters and the performance of neural network models.

For the generation of synthetic patient data in the main loop, the number of patients per hour is iterated by days and hours and is generated using the Poisson distribution. Treatment cost, severity of illness, waiting time, and arrival time are generated using different normal and uniform distributions. This is the essential part of simulation, which creates a realistic dataset based on statistical distributions.

Probabilistic distributions, such as the Poisson distribution, the Normal distribution, or the Uniform distribution, allow the modelling of typical behaviours in real-world settings. Poisson distribution is used to model rare events in a given time interval (e.g., number of patients per hour). Normal Distribution models continuous variables centered around a mean value (e.g., cost of treatment). Uniform Distribution assigns the same probability to any value within a specific interval (e.g. time of arrival within an hour).

The obtained data is converted into a Pandas DataFrame that includes information on day, arrival time of each patient, treatment cost, severity of illness, and waiting time. DataFrame Pandas is used here as a very useful library for manipulating and analysing tabular data. Structuring data in tabular form allows for storing data in a clear and accessible format, calculating statistics and aggregating data, easy visualisations, and comparability with other methods. Each model utilised in this study (Feedforward, LSTM, CNN) generates synthetic predictions based on the actual treatment cost. Predictions are calculated as variations of the true value, using random coefficients between 0.8 and 1.2 for Feedforward, 0.85 and 1.15 for LSTM, and 0.9 and 1.1 for CNN. These variations simulate how models can approximate the true cost with some error.

Table 1. Performance metrics of each model

Models	Average cost per patient (RON)	Average cost per day (RON)	Cumulative costs (RON)
Feedforward	9232,902286	323151,58	9694547,4
LSTM	9325,237143	326383,3	9791499
CNN	9690,178286	339156,24	10174687,2

Table 1 presents the costs associated with different models: Feedforward model, LSTM, and CNN. It includes three cost categories: average cost per patient, average cost per day, and cumulative costs. The lowest average cost per patient is for the Feedforward model (9232,90 RON), while the highest is for CNN (9690,18 RON). The average cost per day follows a similar trend, increasing from 323151,58 RON to 339156,24 RON. The highest cumulative costs are associated with the CNN model (10174687,2 RON), while the lowest total costs are for the Feedforward model (9694547,4 RON).

Figure 2 illustrates the distribution of treatment costs (in RON) based on two key variables: the severity level of cases and the waiting time before treatment. The X-axis (Waiting Time) is divided into five intervals: 5-10 minutes, 10-15 minutes, 15-20 minutes, 20-25 minutes, and 25-30 minutes, while the Y-axis (Severity Level) is categorised into five severity ranges: 1,005-2,811 (very low severity), 2,811-4,608 (low severity), 4,608-6,406 (medium severity), 6,406-8,203 (high severity), and 8,203-10 (very high severity).




Figure 2. Cost distribution by severity and waiting time

The values inside the cells represent the average treatment cost (RON) for each combination of severity level and waiting time. The general trend shows that costs increase with severity, as more severe cases (lower on the Y-axis) tend to have higher treatment expenses. Additionally, shorter waiting times are associated with higher costs since urgent cases require immediate and often more expensive resources. A notable pattern appears in the (2,811-4,608, 5-10 minutes) cell, which has the highest recorded cost (9638 RON), indicating that certain combinations of severity and short waiting times substantially increase treatment costs. Conversely, lower costs are observed in cases with lower severity and longer waiting times. Potential anomalies may exist in isolated high-value cells, which could indicate special cases or unjustified costs. This heatmap is useful for resource allocation, helping hospitals identify where more funding is needed, waiting time management, where reducing waiting times for severe cases could optimise costs, and improving medical process efficiency by analysing the relationship between severity, waiting time, and treatment cost to enhance patient flow management.

Figure 3 illustrates the daily cost distribution estimates made by three prediction models: Feedforward, LSTM, and CNN. On the Y-axis, we have the estimated daily costs, ranging from approximately 250000 RON to 500000 RON. The median costs for each model are 323151,58 RON for Feedforward, 326383,30 RON for LSTM, and 339156,24 RON for CNN, with CNN having the highest median value, indicating that it estimates higher costs on average. The box in each plot represents the interquartile range (IQR), containing 50% of the data, while the whiskers extend to the minimum and maximum values within a reasonable range, excluding outliers, which would appear as separate points. Feedforward Prediction has a more concentrated distribution, indicating lower variation, whereas CNN Prediction has the widest distribution, suggesting that its predictions vary more significantly across days. In terms of model performance, a model with less variation in daily costs is preferable for consistency, meaning Feedforward and LSTM are more stable, while CNN, despite its higher median cost, might be better at capturing more severe cases.




Figure 3. Average cost per day

This graph is useful for model comparison, helping identify which models generate more realistic predictions, and for performance analysis, showing that models with tighter distributions are more stable, whereas those with higher medians may better represent actual costs for severe cases.




Figure 4. Distribution of treatment costs by severity level

The distribution of treatment costs (RON) across different severity levels is shown in in Figure 4, with severity categories (Very Low, Low, Medium, High, and Very High) on the X-axis and the Y-axis displaying treatment costs ranging from approximately 2000 RON to 16000 RON. The box represents the interquartile range (IQR), with the bottom (Q1) and top (Q3) edges marking the 25th and 75th percentiles, while the line inside the box indicates the median cost for each severity level. The whiskers extend to the minimum and maximum values within 1,5 times the IQR, and outliers, shown as dots outside the whiskers, represent extreme costs. The graph reveals that the median cost increases with severity, indicating that more severe cases incur higher treatment expenses. Low-severity cases show lower cost variation, whereas high-severity cases exhibit greater fluctuations, reflecting the unpredictability of medical expenses for severe conditions. The general trend indicates that as severity rises, treatment costs increase and their variability becomes more pronounced. This visualization is valuable for resource allocation, helping determine which severity levels require more funding, cost control, by assessing whether extreme costs are justified, and policy development, by aiding in the design of treatment strategies based on severity levels.

The predictions for treatment costs (in RON) generated by three machine learning models, Feedforward, LSTM, and CNN is presented in Figure 5. The Y-axis shows the estimated treatment costs, which range from 0 to approximately 18000 RON. Each dot in the plot represents an individual predicted treatment cost. The shape of the swarm reflects the distribution and concentration of the predicted values, where the Feedforward Prediction model shows a tighter distribution with lower variability, the LSTM Prediction model exhibits a wider spread, indicating more diverse predictions, and the CNN Prediction model demonstrates the most dispersed predictions, especially in the higher cost range. Outliers are present in all models, with Feedforward containing a few extreme values near 0 RON and 17500 RON, while LSTM and CNN also show high-cost outliers, with CNN displaying the widest range. Most predictions for all models are clustered between 5000 RON and 10000 RON, which aligns with typical treatment costs. In comparing the models, the Feedforward Prediction model is more consistent but may not capture extreme cases well, the LSTM model provides greater variation, potentially capturing more complex cases, while the CNN model, having the widest spread, may overestimate certain costs.




Figure 5. Predictions for treatment costs

The graph is useful for model evaluation, helping to identify the most accurate and consistent model, as well as for error analysis, where outliers may indicate areas of poor performance. Additionally, it assists in decision-making, allowing stakeholders to select the model that best aligns with real-world treatment costs.

Figure 6 shows the Mean Absolute Error (MAE) of three machine learning models: Feedforward Neural Network (NN), LSTM NN, and CNN 1D. Lower MAE values indicate better model performance, and the chart highlights the differences among the three models. Feedforward NN has the highest MAE, approximately 190, making it the least accurate, while LSTM NN, with an MAE of around 150, performs better but is still less accurate than CNN 1D NN, which achieves the lowest MAE of 96,18. The chart emphasises CNN 1D NN as the best-performing model, marked in orange, due to its superior accuracy. MAE provides a straightforward way to understand prediction errors in the same units as the target variable (RON in this case), and a lower MAE signifies predictions closer to the actual treatment costs.




Figure 6. Mean Absolute Error for each model

This visualisation underscores the effectiveness of CNN 1D NN in accurately estimating treatment costs, making it the preferred model for cost predictions while also pointing out areas where other models could be improved.




Figure 7. Mean Squared Error for each model

The Mean Squared Error (MSE) of the three models is presented in Figure 7. MSE is a metric that measures the average squared differences between predicted and actual values, penalising larger errors more heavily. Among the models, Feedforward NN has the highest MSE, exceeding 50000, indicating that it is the least accurate model and produces the most significant errors. LSTM NN performs better than Feedforward NN, with an MSE of approximately 35000, but it is still outperformed by CNN 1D NN, which achieves the lowest MSE, marked in orange, at less than 15000. The significantly lower MSE of CNN 1D NN demonstrates its ability to produce predictions that are closer to actual values while minimizing the impact of large prediction errors. This comparison highlights CNN 1D NN as the most reliable model for predicting treatment costs, while the higher MSE of the other models suggests a need for further optimization to improve their predictive accuracy.

By incorporating the RMSE metric alongside MAE and MSE, the evaluation process becomes more robust, offering deeper insights into error distribution and overall model performance as can be seen in Figure 8. This is especially important in healthcare, where even small cost miscalculations can impact budgeting, insurance claims, and hospital resource allocation. The model with the smallest RMSE performs best in minimising prediction errors. If a model has a higher RMSE, it means its predictions deviate significantly from actual values.




Figure 8. Root Mean Squared Error for each model




Figure 9. Predicted treatment costs vs the actual treatment costs

Next, we compared the predicted treatment costs to the actual treatment costs for the three machine learning models in Figure 9. Each plot includes a diagonal dashed line, representing the ideal case where predicted costs perfectly match actual costs. The scatter points show the relationship between actual and predicted values for each model.

The Feedforward NN plotted with blue has more points widely scattered around the diagonal line, indicating a larger deviation between predictions and actual costs and, hence, lower prediction accuracy. The LSTM NN plotted with orange shows improved alignment, with most points closer to the diagonal, reflecting better prediction accuracy compared to the Feedforward NN. Finally, the CNN 1D NN plotted with green exhibits the highest concentration of points around the diagonal line, signifying the most accurate predictions and the best model performance among the three.

The feedforward neural network model can be used for applications where response time and cost are priorities and the data does not contain significant temporal relationships. It is easy to implement and explain, making it suitable for rapid estimation scenarios in small clinics.

The LSTM Neural Network excels in long-term patient monitoring applications, offering excellent accuracy in capturing complex sequential relationships. However, it is more expensive and complex to implement, being recommended in large hospitals or medical centers that have the infrastructure and resources to support advanced deep learning models.

The CNN 1D Neural Network is a model suitable for capturing local features from time sequences and offers a balance between complexity and performance. This model is recommended for hospitals or clinics that have access to specialised hardware (GPU) and want a computationally efficient model for short-term sequential data.

This study allows us to perform clinical simulations by modeling real hospital scenarios, to evaluate different neural network architectures (Feedforward, LSTM, CNN) for accuracy, and to perform exploratory analysis by discovering relationships between variables such as cost, severity, and waiting time.

A key novelty in this study is the demonstration that 1D CNN outperforms other neural network models, highlighting its potential for real-world clinical applications. Additionally, the study emphasises the role of simulations in testing predictive methodologies in scenarios where real-world data is scarce.

In the future, the authors propose to generate more complex data by introducing additional variables (e.g., age of patients, type of disease), using other evaluation methods, calculating performance metrics to compare model predictions, or automating visualisations by automatically generating graph-based reports.

In the future, we aim to compare simulation-based models with real-world studies on medical cost prediction to ensure their accuracy and practical relevance. By integrating anonymised patient data, using collaborative learning, and refining synthetic datasets with real-world insights, we can bridge the gap between theoretical models and actual healthcare scenarios. This comparison

will help assess the reliability of neural network predictions and enhance their applicability in clinical decision-making while maintaining strict privacy compliance. Incorporating real patient data into future research while maintaining privacy compliance requires a combination of ethical considerations, legal frameworks, and technical safeguards.

3. Conclusions

The results highlight the usefulness of using synthetic data to simulate clinical scenarios. The 1D CNN model provided superior performance compared to Feedforward NN and LSTM NN, with the lowest values for MAE and MSE. This suggests that the CNN architecture is better equipped to capture the complex relationships between the analysed variables. LSTM generally provides the best accuracy for sequential data with obvious temporal dependencies, making it more suitable for monitoring patients over time. 1D CNN is also competitive but may be less accurate for data with long-term dependencies. If computational efficiency is a priority, the Feedforward network is the fastest option. However, the 1D CNN offers a better balance between accuracy and training time, making it a viable alternative when both capturing temporal features and having a shorter training time than LSTM are desired.

The simulation also demonstrated the ability to identify trends between severity, waiting time, and costs, which may be valuable for resource management in hospitals. The use of intuitive visualizations allowed detailed interpretation of the results, highlighting the importance of combining statistical analysis with graphical representation.

For future research, extending the simulation to include variables such as medical history, length of hospitalization, or type of intervention may improve the accuracy of predictive models. Applying the methodology to real data would also validate the findings and highlight the limitations of the models. This approach provides a robust framework for exploring and optimising predictive algorithms in healthcare and beyond.

Acknowledgement

The publication of this article was supported by the 2024 Development Fund of the Babeş-Bolyai University.

References

- Abatal, A., Mzili, M., Benlalia, Z., Khallouki, H., Mzili, T., El Kaim Billah, M., & Abualigah, L. (2025). Hybrid long short-term memory and decision tree model for optimizing patient volume predictions in emergency departments. *International Journal of Electrical and Computer Engineering*, 15(1), 669-676. <https://doi.org/10.11591/ijece.v15i1.pp669-676>
- Ali, H., Lamsali, H., & Othman, S. (2018). Hospital scheduling analysis: A contemporary review and proposed schematic understanding. *Journal of Advanced Research in Dynamical and Control Systems*, 10, 164-173.
- Azadeh, A., Hosseinabadi Farahani, M., Torabzadeh, S., & Baghersad, M. (2014). Scheduling prioritized patients in emergency department laboratories. *Computer Methods and Programs in Biomedicine*, 117(2), 61-70. <https://doi.org/10.1016/j.cmpb.2014.08.006>
- Carter, M. W., & Lapierre, S. D. (2001). Scheduling emergency room physicians. *Health Care Management Science*, 4(4), 347-360. <https://doi.org/10.1023/A:1011802630656>
- Cincar, K. (2020). Hospital management and scheduling with multi-agents approach. *2020 International Conference on e-Health and Bioengineering (EHB)*, Iasi, Romania, 1-4. <https://doi.org/10.1109/EHB50910.2020.9280281>
- Cincar, K., Minda, A. A., & Varga, M. (2024). A simulation-based analysis using machine learning models to optimize patient flow and treatment costs. *BRAIN: Broad Research in Artificial Intelligence and Neuroscience*, 15(3), 180-187. <http://doi.org/10.70594/brain/15.3/13>

- Elalouf, A., & Wachtel, G. (2015). An alternative scheduling approach for improving patient flow in emergency departments. *Operations Research for Health Care*.
<https://doi.org/10.1016/j.orhc.2015.08.002>
- Kiris, S., Yüzügüllü, N., Ergün, N., & Çevik, A. A. (2010). A knowledge-based scheduling system for emergency departments. *Knowledge-Based Systems*, 23(8).
<https://doi.org/10.1016/j.knosys.2010.06.005>
- Sharafat, A., & Bayati, M. (2021). PatientFlowNet: A deep learning approach to patient flow prediction in emergency departments. *IEEE Access*, PP(1), 1-1.
<https://doi.org/10.1109/ACCESS.2021.3066164>
- Yanamala, A. K. Y., & Suryadevara, S. (2022). Cost-sensitive deep learning for predicting hospital readmission: Enhancing patient care and resource allocation. *International Journal of Advanced Engineering Technologies and Innovations*, 1(3), 56-81.