Artificial Intelligence and Social Computing, Vol. 122, 2024, 96–102
https://doi.org/10.54941/ahfe1004646

Development of an Explainable Pre-Hospital Emergency Prediction Model for Acute Hospital Care

Minjung Lee, Hyunho Park, and Eun Jung Kwon

Electronics and Telecommunications Research Institute, Daejeon, 34129, Republic of Korea

ABSTRACT

This study introduces an eXplainable Artificial Intelligence (XAI) designed to predict which emergency patients require acute hospital care in pre-hospital phase and provide explanations for its reasoning. Emergency medical care is broadly divided into two stages: pre-hospital and in-hospital stages. Various information gathered during the emergency activities performed by paramedics in the pre-hospital stage and while transporting patients is crucial in describing the emergency patient's condition. However, key pre-hospital information, important for the in-hospital medical care of emergency patients, is filtered based on the ambiguous memory of the paramedics, and is verbally shared in a condensed form via phone or radio when transmitted to the hospital. To address this issue, we have developed a model that predicts emergency patients based on pre-hospital information integrating an ensemble model and advanced XAI techniques. This proposed model not only predicts emergency situations requiring acute hospital care but also ensures the model's predictive processes remain transparent and interpretable for medical professionals, addressing the critical need for an information linkage system between the pre-hospital and in-hospital phases.

Keywords: Explainable artificial intelligence (XAI), Pre-hospital emergency prediction, Acute hospital care

INTRODUCTION

Effective emergency medical care is crucial in saving lives and improving the recovery prospects of patients in critical health conditions or those suffering severe injuries. The emergency medical service process is broadly divided into two phases: the pre-hospital and the in-hospital phase, as illustrated in Figure 1. The pre-hospital phase encompasses medical services provided before a patient arrives at a hospital, typically by paramedics or emergency medical technicians. This phase includes the preliminary assessment, stabilization, and transport of patients. The in-hospital care consists of medical services rendered upon the patient's arrival at the hospital, ranging from further assessments to treatment and recovery efforts (Wilson et al., 2015).

A significant challenge in the seamless integration of these two stages is the effective utilization of pre-hospital information (Zhang et al., 2021).

However, much of the information about the patient's condition during transport and the care performed by paramedics at the pre-hospital stage has not been systematically delivered to the hospital. In South Korea, typically, only very brief information, filtered through the paramedics' subjective judgement, is relayed to the hospital orally via radio. Consequently, valuable information that could influence patient treatment plans and outcomes may not be fully utilized in hospital care. Currently, this is the case in many countries due to the absence of a real-time information linkage system between the pre-hospital and in-hospital phases.

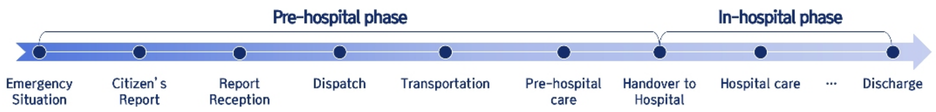

Figure 1: Process of emergency medical service.

To tackle this issue, we propose a novel approach employing an explainable artificial intelligence (XAI) model. This model is designed to identify emergency patients requiring acute hospital care in the pre-hospital stage and elucidate the reasoning behind its predictions to the medical professionals, thus enabling medical professionals to comprehend the predictive elements behind its decisions and perform treatments considering this information. Such transparency aims to foster trust and make the model's predictions more interpretable. This approach can be used for bridging the information gap between pre-hospital and in-hospital care, ensuring that decisions of the in-hospital stage are made based on a comprehensive understanding of the patient's condition and the pre-hospital measures undertaken.

PROBLEM DEFINITION AND DATASET DESCRIPTION

The aim of this study is to develop and validate an XAI model capable of accurately predicting the need for urgent hospital care based on information obtained during the pre-hospital phase. The primary data source used in this research is paramedic activity logs. These logs are comprehensive documents that record the patient rescue efforts and emergency medical treatment activities administered by paramedics in response to emergency situations during transport. They include a wide range of information, including the patient's condition, emergency treatments provided, and details of the transportation process.

For the proposed XAI model training, the paramedic activity logs data is labeled with the disease information contained in the discharge summary records. Our research specifically focuses on 24 critical emergency diseases that necessitate acute care, as outlined in Table 1. Of the 1.6 million data entries we used, only 9.13 percent of the patients belong to critical emergency diseases that require acute care. Furthermore, as shown in Figure 2, there is a clear difference in the proportions of each of the 24 diseases of interest, indicating that the data has significantly imbalanced classes.

Table 1. 24 critical emergency diseases.

No	Disease	No	Disease
1	Lethal Arrhythmia	13	Disseminated Intravascular Coagulation (DIC)
2	Gallbladder and Biliary Disease	14	Poisoning
3	Shock	15	Severe Infection and Sepsis
4	ARDS/Pulmonary Edema	16	Severe Trauma
5	Status Epilepticus	17	Acute Aortic Syndrome
6	Acute Renal Failure, Diabetic Coma	18	Acute Abdominal Conditions
7	Acute Myocardial Infarction STEMI	19	Bronchial Hemorrhage and Foreign Body
8	Acute Myocardial Infarction NSTEMI	20	Obstetric Emergency
9	Acute Cerebrovascular Accident	21	Severe Burns
10	Cerebral Hemorrhage	22	Emergency Vascular Disease
11	Gastrointestinal Bleeding/Foreign Body	23	Ophthalmologic Emergency
12	Cardiac Arrest	24	Urological Emergency

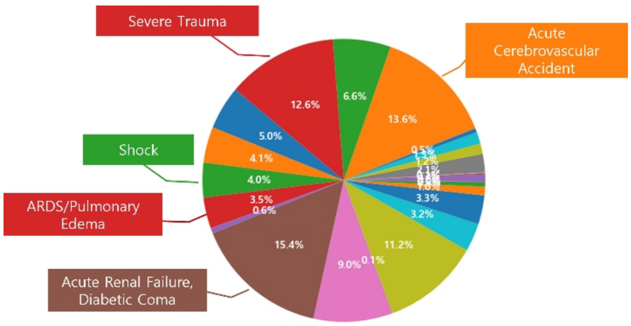

Figure 2: Proportions of each of the 24 diseases.

Data pre-processing involved data resampling and cleaning processes. Resampling techniques, such as under-sampling and over-sampling, were employed to balance the data relating to emergency conditions, addressing the issue of data imbalance. Furthermore, the data cleaning process included the removal of incomplete or erroneous data entries, elimination of duplicate information, handling of missing values to enhance data quality. Lastly, the transformation of features for the XAI model training was conducted. These pre-processing steps are essential for maximizing the efficiency of model training and improving prediction accuracy.

EXPLAINABLE PRE-HOSPITAL EMERGENCY PREDICTION MODEL

In this section of this paper, we delve into the specifics of the explainable severe emergency diseases prediction model, incorporating an ensemble approach with a focus on XGBoost (eXtreme Gradient Boosting) (Chen and Guestrin, 2016), a tree-based ensemble model, and SHAP (SHapley Additive exPlanations) (Lundberg and Lee, 2017) for explainable AI techniques.

Our predictive model harnesses the power of XGBoost, a highly efficient and scalable implementation of gradient boosting, to forecast the need for acute hospital care among emergency patients at the pre-hospital stage. XGBoost stands out for its capability to handle large-scale tabular data over other machine learning approaches while providing an intrinsic overfitting regularization function. As depicted in Figure 3, XGBoost operates on the principle of sequentially building trees, where each new tree corrects errors made by previously constructed trees, thereby continuously improving prediction accuracy. To augment the transparency and interpretability of our model, we integrate SHAP values, a model-agnostic explainable AI technique that clarifies the contribution of each feature to the model's predictions. The model is trained on the pre-hospital information collected from the paramedic activity logs, which include patient demographics, vitals, and treatment actions. To optimize the model's performance, we incorporate Optuna (Akiba et al., 2019) for hyperparameter optimization, tuning the hyperparameters of model to achieve the most favorable outcomes. Following the training of the XGBoost model, we apply SHAP to render its decision-making process transparent and interpretable. SHAP values elucidate the impact of each feature on the prediction, furnishing a comprehensive and easily understandable breakdown for medical professionals.

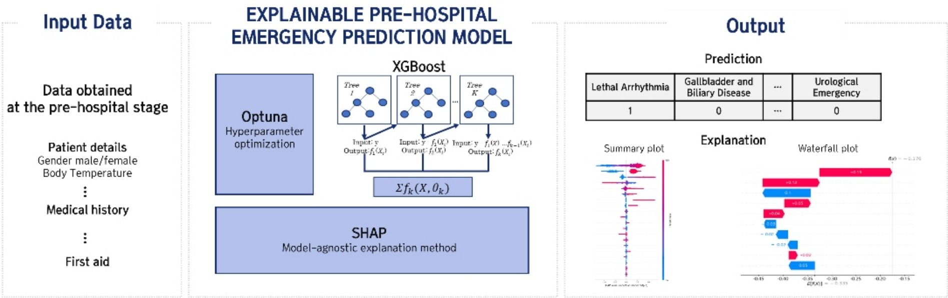

Figure 3: Diagram of explainable pre-hospital emergency prediction.

The integration of SHAP with our XGBoost model allows for the visualization of key reasons behind the model's predictions, which are communicated to medical staff. This ensures that the model's outputs are not just accurate but also meaningful and actionable for those making clinical decisions. Two main types of visualizations are used to interpret the model's results:

Summary Plot: This provides an overview of the impact of different variables on the predictions for the condition being analyzed. It helps in understanding the relative importance of each feature at a global level.

Waterfall Plot: For individual patient predictions, waterfall plots detail the contribution of each variable to the specific outcome. This granular insight is crucial for medical professionals assessing the urgency and specific needs of a patient, allowing them to understand how and why the model has reached its conclusion.

Through the use of XGBoost and SHAP, complemented by Optuna for hyperparameter optimization, our model not only predicts severe emergency conditions with high accuracy but also ensures that these predictions are transparent and interpretable. This approach bridges the gap between advanced machine learning techniques and clinical applicability, facilitating the adoption of AI tools in emergency medical settings and enhancing patient care by providing clear, actionable insights derived from pre-hospital information.

EXPERIMENTAL RESUTLS

To evaluate the model's predictive capabilities, we use Accuracy and F1 score. Accuracy measures the proportion of total predictions that are correct. While intuitive, accuracy might not be the best measure in imbalanced datasets where the event of interest (e.g., emergencies) is rare. F1 score is a harmonic mean of Precision and Recall. Precision indicates the proportion of positive identifications that are actually correct. In the context of acute hospital care, it reflects how many of the predicted emergency disease are genuine. Recall (Sensitivity) measures the proportion of actual emergency diseases that are correctly identified. Upon examining the performance results summarized in Table 3, it appears necessary to improve the model by increasing Recall to enhance the F1 score. However, from the perspective of Accuracy, it was observed that the model secured quite accurate results.

Table 2. Prediction results (top 5 diseases based on F1 score).

Disease	Accuracy	F1score
Cardiac Arrest	98.30	0.56
Acute Renal Failure, Diabetic Coma	85.43	0.53
Gastrointestinal Bleeding/Foreign Body	89.23	0.52
Poisoning	86.87	0.52
Severe Burns	98.28	0.50

Figure 4 presents a summary plot, one of the SHAP results for the cardiac arrest prediction model. The x-axis shows the SHAP values, while the y-axis is organized in descending order based on the importance of the variables. The color coding of the dots represents the actual values of the variables, with redder hues indicating higher values and bluer hues indicating lower values. Dots located in the negative region of the x-axis negatively influence the prediction, whereas those in the positive region have a positive impact. A pattern of blue dots on the left and red dots on the right suggests a strong positive correlation with the target variable, whereas the reverse pattern implies a negative correlation.

From this analysis, we deduced the variables the AI model relies on to predict cardiac arrest. The model was more likely to predict cardiac arrest when the paramedics assessed the patient's symptoms as cardiac or respiratory arrest, and emergency interventions such as CPR, oxygen delivery via Bag Valve Mask (BVM), AED Monitoring, and airway management were implemented, particularly when the patient's state of consciousness was not Alert (A) but Unresponsive (Coma).

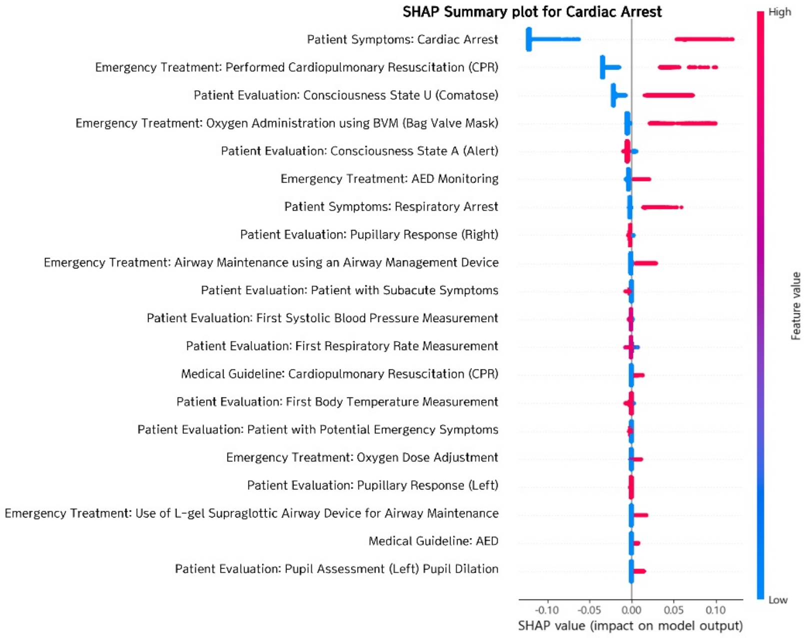

Figure 4: SHAP summary plot for cardia arrest.

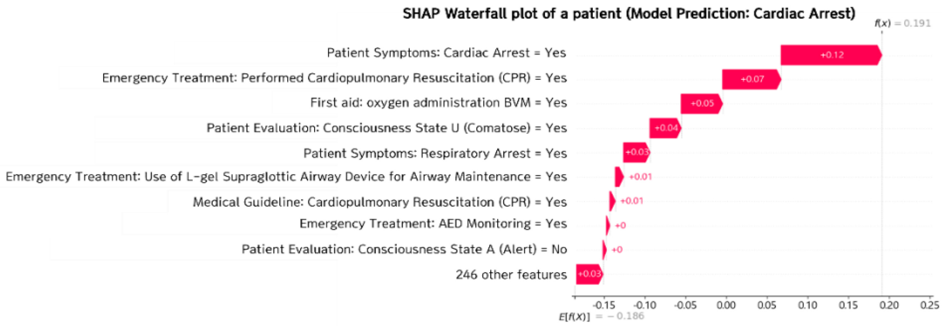

Figure 5: SHAP waterfall plot of a patient (model prediction: cardiac arrest).

The waterfall plot allows for a detailed analysis of how each variable, including symptoms assessed and interventions performed by paramedics, impacts the model's prediction. This feature is instrumental for hospital physicians, as it provides a visual confirmation of which pre-hospital stage information most accurately describes the patient's condition upon arrival. Consequently, this aids in bridging the informational gap between pre-hospital and in-hospital data, ensuring that treatment decisions are informed by a comprehensive understanding of the patient's pre-hospital care.

In the scenario depicted by Figure 5, paramedics assessed a patient exhibiting severe symptoms indicative of cardiac and respiratory arrest. Immediate pre-hospital emergency interventions were administered, including Cardiopulmonary Resuscitation (CPR), oxygen delivery using a Bag Valve Mask (BVM), comprehensive airway management, and monitoring with an Automated External Defibrillator (AED). These interventions were critical in stabilizing the patient's condition en route to the hospital. Our model predicted the patient's condition as a cardiac arrest, influenced significantly by the patient being in a U (Coma) state.

CONCLUSION

This study introduces an innovative model of explainable artificial intelligence that predicts emergency medical conditions from pre-hospital information, markedly enhancing emergency medical services' continuity and quality. Integrating advanced ensemble models with SHAP for explainability, our model accurately forecasts severe emergency conditions and provides insight into the reasoning behind its predictions. The model's transparency and interpretability crucially bridge the gap between paramedic-provided pre-hospital care and in-hospital treatments, allowing medical professionals to make informed decisions with a comprehensive understanding of the patient's pre-admission condition and interventions.

Experimental results underscore the model's ability to identify life-threatening conditions, including cardiac arrest and acute renal failure, accurately. Detailed visualizations, such as summary and waterfall plots, clarify the contributions of various variables to these predictions. Adopting this XAI approach represents a significant advancement in emergency medical care, promising improved patient outcomes through seamless information flow between pre-hospital and in-hospital care phases. This study establishes a precedent for integrating explainable AI into healthcare, highlighting AI's potential to enhance decision-making processes and build medical professionals' trust in AI-assisted predictions.

ACKNOWLEDGMENT

This research was supported by CONNECT-AI platform(CONnected Network for EMS Comprehensive Technical-support using Artificial Intelligence) of the Korea Planning & Evaluation Institute of Industrial Technology (KEIT) funded by the Korea National Fire Agency (No. RS-2023-00237687).

REFERENCES

Akiba, T., Sano, S., Yanase, T., Ohta, T., & Koyama, M. (2019, July). Optuna: A next-generation hyperparameter optimization framework. In Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining (pp. 2623–2631).

Chen, T., & Guestrin, C. (2016, August). Xgboost: A scalable tree boosting system. In Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining (pp. 785–794).

Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. Advances in neural information processing systems, 30.

Wilson, M. H., Habig, K., Wright, C., Hughes, A., Davies, G., & Imray, C. H. (2015). Pre-hospital emergency medicine. The Lancet, 386(10012), 2526–2534.

Zhang, Z., Sarcevic, A., Joy, K., Ozkaynak, M., & Adelgais, K. (2021). User needs and challenges in information sharing between pre-hospital and hospital emergency care providers. In AMIA Annual Symposium Proceedings (Vol. 2021, p. 1254). American Medical Informatics Association.