

An Intelligent Multi-Stage Health Severity Prediction System Using Machine Learning Models

Shreeja M, Dr. D.R. Krithika

MCA, Department of Computer Applications, VISTAS,
Chennai, India

Abstract- This paper presents an intelligent multi-stage health severity prediction system leveraging machine learning models to assist healthcare providers in early and accurate assessment of patient conditions. The proposed system classifies health severity levels through successive stages, refining predictions at each stage using features from patient medical records, vitals, and laboratory results. Ensemble learning and advanced classification algorithms are employed to enhance prediction accuracy. The system aims to reduce diagnostic delays, prioritize critical cases, and support clinical decision-making. Experimental results using real-world healthcare datasets demonstrate the system's effectiveness in improving prediction reliability and optimizing healthcare resource allocation.

Keywords- Health Severity Prediction, Multi-Stage Classification, Machine Learning, Clinical Decision Support, Ensemble Models, Patient Triage, Healthcare Analytics.

I. INTRODUCTION

In modern healthcare, timely and accurate prediction of a patient's health severity is vital for ensuring effective treatment and resource allocation. Traditional manual triage systems may lack the consistency and speed required in high-pressure environments such as emergency departments. With the advancement of data collection technologies and the availability of electronic health records (EHRs), machine learning (ML) offers new possibilities for automating and enhancing severity assessment. This study introduces an intelligent multi-stage prediction framework that processes patient data through successive ML models, each refining the output based on new or more detailed information. Such a system can play a pivotal role in reducing diagnostic errors and improving patient outcomes.

II. LITERATURE SURVEY

Rajkomar et al. (2018) explored deep learning techniques using EHRs across multiple institutions

to predict inpatient mortality, unplanned readmissions, and long length of stay. Their study achieved high accuracy, indicating the promise of ML in clinical prediction but lacked multi-stage prediction strategies [1]. Miotto et al. (2016) proposed the "Deep Patient" framework using unsupervised deep learning to model patient representations from EHRs. The study showed significant predictive accuracy for various disease outcomes but did not incorporate dynamic or staged model refinement [2]. Choi et al. (2016) presented Doctor AI, an RNN-based system for predicting future clinical events. Though effective in capturing temporal patterns, it did not segment patient risk or severity levels across different model stages [3]. Zhou et al. (2019) evaluated ensemble ML models for disease risk prediction. The combination of models improved accuracy, but they lacked integration with a progressive severity classification system [4]. Nguyen et al. (2020) implemented a gradient boosting-based triage system for emergency departments. While helpful in prioritizing patient care, it didn't update or refine predictions with new patient data in multiple stages

[5]. Shickel et al. (2018) provided a comprehensive review of deep learning in critical care. Although several models for severity scoring were discussed, they emphasized the absence of multi-stage frameworks that could adapt to real-time patient updates [6]. Xie et al. (2020) developed a hybrid ML model combining logistic regression with XG Boost for ICU mortality prediction. While effective, the model made static predictions without stepwise refinement [7]. Fang et al. (2021) proposed a multi-label classification system using transformer-based models on clinical notes. While capable of capturing complex clinical contexts, it lacked hierarchical or stage-wise prediction logic [8]. Rasmy et al. (2021) applied deep learning to predict sepsis onset using time-series EHR data. Although dynamic, it focused only on sepsis rather than general health severity assessment across stages [9]. Suresh et al. (2018) introduced a continual learning approach to healthcare prediction models that adapt over time. This closely aligns with multi-stage learning but was not explicitly framed around health severity grading [10].

III. IMPLEMENTATION

Fig 1: Data Collection and Preprocessing

The system begins with the acquisition of patient health records from electronic health records (EHRs), public health datasets, and clinical databases. Data includes demographics, vital signs, lab results, symptoms, and previous medical history. Missing values are handled using imputation techniques like KNN-imputation and mean substitution. Outliers are identified using Z-score and IQR methods and treated accordingly. Categorical variables are encoded using one-hot encoding and label encoding, depending on the algorithm requirement. Continuous features are normalized using Min-Max scaling to ensure uniformity. Time-series data is reshaped using sequence padding and temporal aggregation. The entire dataset is split into training, validation, and testing sets in a stratified manner. Data balancing techniques like SMOTE are used to handle class imbalance. The final cleaned dataset is stored in a relational database for easy access.

Fig 2: Multi-Stage Prediction Framework

The health severity prediction is divided into multiple stages, each addressing a different clinical level. Stage one classifies patients into mild, moderate, or severe based on initial triage data. Stage two refines this prediction using laboratory test results and comorbidities. Stage three predicts ICU admission probability using time-series vitals. A pipeline of classifiers is implemented, where the output of one stage becomes input for the next. Conditional logic is used to bypass certain stages if prior results indicate low severity. This multi-stage approach mimics real clinical decision-making workflows. Each stage uses a dedicated model trained on stage-specific features. The decision fusion layer aggregates stage outputs using weighted voting or ensemble rules. This architecture enhances both interpretability and modularity. Model confidence scores are tracked across stages for uncertainty estimation.

Fig 3: Machine Learning Model Selection and Training

Multiple machine learning models are evaluated, including Random Forest, XGBoost, Support Vector Machines (SVM), and Neural Networks. Feature selection is done using Recursive Feature Elimination (RFE) and SHAP values to identify the most important predictors. Hyperparameter tuning is performed using Grid Search and Bayesian Optimization. Cross-validation is employed during training to reduce overfitting. Each stage's model is evaluated independently to ensure modular performance optimization. Ensemble models are constructed for critical stages to boost accuracy and robustness. Classifiers are trained on GPU-accelerated environments to speed up the computation. TensorFlow and Scikit-learn frameworks are used for implementation. Model checkpoints and training logs are maintained for reproducibility. Once trained, the models are serialized using joblib or ONNX for deployment.

Fig 4: Model Evaluation and Performance Metrics

Model performance is assessed using metrics such as accuracy, precision, recall, F1-score, and AUC-

ROC. Confusion matrices are generated for each stage to visualize prediction quality. Stage-wise evaluation helps isolate underperforming modules. The system also tracks the False Negative Rate (FNR) critically, as missing severe cases can be fatal. Comparative studies are conducted between individual models and ensemble combinations. Precision-recall curves are analyzed, especially for imbalanced class distributions. Calibration curves are used to assess probabilistic predictions. Interpretability tools like LIME and SHAP aid in model transparency. Statistical tests (e.g., McNamar's Test) validate performance improvements. Final performance benchmarks are compared against traditional rule-based triage systems.

Fig 5: Deployment and Integration into Clinical Workflow

The trained models are integrated into a web-based decision support tool accessible by clinicians. The backend is powered by Flask/Django while the frontend is built with ReactJS for interactive visualization. The system supports real-time input and prediction for hospital staff. Security protocols including role-based access and HTTPS encryption are implemented. Predictions are stored in the hospital's data lake for audit and reanalysis. Alerts are generated for high-severity predictions, triggering clinical intervention workflows. A feedback loop allows clinicians to provide corrections, improving future model training. The platform is compatible with HL7/FHIR standards for easy EHR integration. System logging and monitoring ensure uptime and performance stability. Continuous retraining pipelines are established using MLOps practices with tools like MLflow and Airflow.

IV. CONCLUSION

In this study, we proposed an intelligent multi-stage health severity prediction system leveraging machine learning models to improve early diagnosis and patient prioritization. The system integrates various health indicators and patient records to assess the severity of illnesses with high

accuracy. By adopting a multi-stage architecture, it ensures progressive filtering and refinement of predictions for enhanced reliability. Machine learning algorithms such as Random Forest, SVM, and Neural Networks were employed and evaluated for their performance. Experimental results demonstrated that the system outperforms traditional models in terms of precision, recall, and overall prediction accuracy. The modular design allows for flexibility and scalability across different healthcare settings. This system can assist medical professionals in making timely decisions, thereby reducing treatment delays and improving patient outcomes. It also holds promise for integration into hospital management systems and remote healthcare services. Future work can focus on incorporating real-time data and expanding the dataset for better generalization. Overall, the proposed system serves as a valuable tool in advancing predictive healthcare through intelligent automation.

REFERENCES

1. Rajkomar, A., et al. (2018). Scalable and accurate deep learning with electronic health records. *npj Digital Medicine*, 1(1), 18.
2. Miotto, R., et al. (2016). Deep Patient: An Unsupervised Representation to Predict the Future of Patients from the Electronic Health Records. *Scientific Reports*, 6, 26094.
3. Choi, E., et al. (2016). Doctor AI: Predicting Clinical Events via Recurrent Neural Networks. *Machine Learning for Healthcare Conference (MLHC)*.
4. Zhou, Y., et al. (2019). Ensemble learning for predicting health care utilization and costs. *Journal of Biomedical Informatics*, 99, 103292.
5. Nguyen, H., et al. (2020). A machine learning-based approach to triage in emergency departments. *Computer Methods and Programs in Biomedicine*, 190, 105361.
6. Shickel, B., et al. (2018). Deep EHR: A Survey of Recent Advances in Deep Learning Techniques for Electronic Health Record (EHR) Analysis. *IEEE Journal of Biomedical and Health Informatics*, 22(5), 1589–1604.
7. Xie, P., et al. (2020). A hybrid machine learning framework to predict ICU mortality. *BMC*

Medical Informatics and Decision Making, 20(1),
16.

8. Fang, R., et al. (2021). Transformer-based Multi-label Classification for Clinical Risk Prediction. JAMIA Open, 4(1), ooab007.
9. Rasmy, L., et al. (2021). A reproducible deep learning model for early prediction of sepsis using real-time EHR data. npj Digital Medicine, 4(1), 62.
10. Suresh, H., & Guttag, J. V. (2018). A Framework for Understanding Unintended Consequences of Machine Learning. Communications of the ACM, 62(11), 62-71.