www.ijarp.com ISSN 2456-9992 Page: 01-13

International Journal Advanced Research Publications

Review Article

Volume: 02
Issue: 08

A REVIEW OF AN ATTENTION-BASED DEEP LEARNING FRAMEWORK FOR LIVER DISEASE PREDICTION

***S. Jeypriya**

Assistant Professor, Department of Computer Science Sri Krishna Arts and Science College,Coimbatore, Tamilnadu.

Article Received: 11 July 2026, Article Revised: 31 July 2026, Published on: 20 August 2026

***Corresponding Author: S. Jeypriya**

Assistant Professor, Department of Computer Science Sri Krishna Arts and Science College,Coimbatore, Tamilnadu.

Doi: https://doi-doi.org/101555/ijarp.1345

ABSTRACT

Early identification of liver disease is important for improving treatment outcomes and reducing complications. This study presents an attention-based deep learning framework that combines robust data preprocessing with automatic feature learning to improve liver disease prediction from routine clinical and biochemical data. The proposed pipeline begins by handling incomplete records through statistical imputation, detecting abnormal observations using Isolation Forest and Local Outlier Factor techniques, and standardizing feature values to ensure consistent model training. A deep autoencoder is then employed to learn compact latent representations that capture hidden relationships among liver biomarkers. These learned features are supplied to an attention-based neural network, enabling the model to focus on the most informative patient characteristics during classification. To enhance clinical transparency, SHAP is used after prediction to explain the contribution of individual biomarkers. The framework is evaluated on a benchmark liver disease dataset and compared with traditional machine learning models such as Logistic Regression, Random Forest, and Support Vector Machine. The proposed approach is expected to provide improved predictive performance while maintaining model interpretability. Overall, the framework offers a practical and clinically meaningful decision-support solution for the early detection of liver disease.

KEYWORDS: Chronic Liver Disease, LOF, Attention-based Deep Learning, SHAP, Adam optimizer.

1. INTRODUCTION

Liver disorders, including hepatitis, fatty liver disease, cirrhosis, and liver cancer, continue to pose a major healthcare challenge across the world. Because many liver conditions remain asymptomatic during their early stages, diagnosis often depends on laboratory investigations such as liver function tests. Although these tests provide valuable information, individual biomarkers alone are often insufficient for reliable diagnosis. Machine learning has therefore emerged as a promising tool for integrating multiple clinical parameters and supporting physicians in identifying patients who are at high risk[5]. Conventional machine learning methods have demonstrated useful performance but may fail to capture the complex nonlinear relationships among biochemical markers[9]. Attention-based deep learning models provide a more flexible solution by automatically assigning greater importance to clinically relevant features for each patient[10]. Building on this idea, the proposed framework integrates preprocessing, feature representation learning, attention-based classification, and explainable AI to develop a robust and interpretable liver disease prediction system.

2. Proposed Methodology

2.1 Advanced Data Preprocessing, Feature Engineering, and Feature Learning

The effectiveness of any intelligent healthcare prediction system depends largely on the quality of the data used for model training. Clinical datasets often contain missing values, noisy observations, inconsistent feature distributions, and redundant information, all of which can adversely affect the performance of machine learning and deep learning algorithms. Liver disease datasets are particularly susceptible to these issues because they are collected from multiple clinical environments where laboratory measurements, patient demographics, and data recording practices may vary. Consequently, a comprehensive preprocessing strategy is essential to improve data quality before classification.

The proposed framework employs a structured preprocessing pipeline designed to enhance the reliability of the clinical data while preserving important diagnostic information. Unlike conventional approaches that rely on a single preprocessing technique, this framework integrates multiple complementary methods to address different data quality challenges. Each preprocessing stage is performed sequentially so that the output of one stage serves as the input for the next, resulting in a clean and informative dataset for feature learning[6].

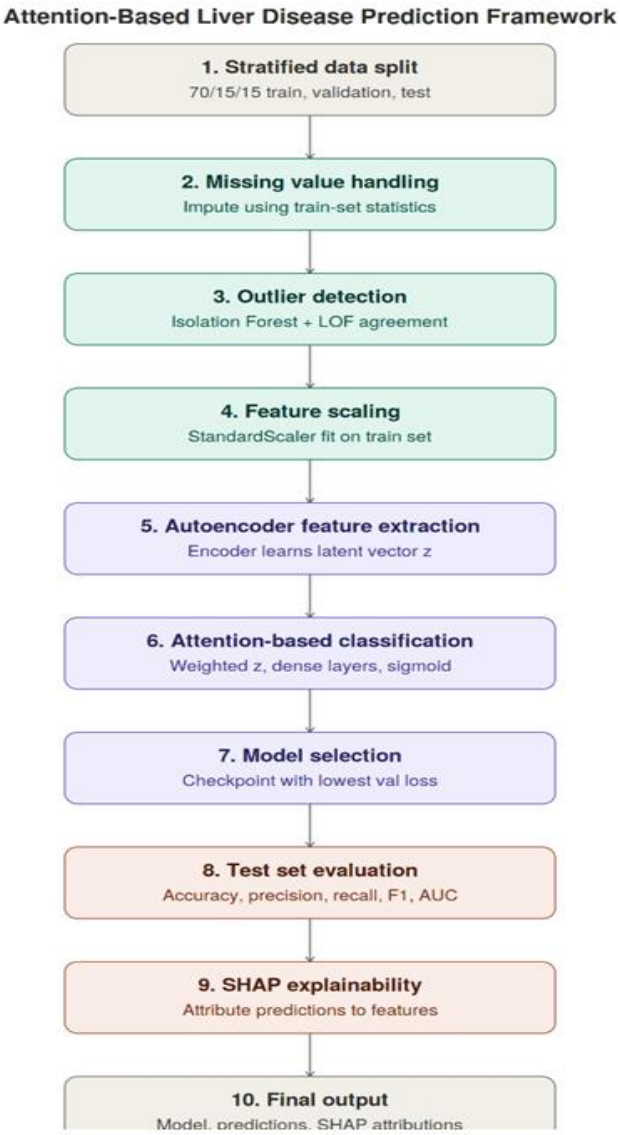

Fig. 1. Proposed Method

2.1 Advanced Data Preprocessing, Feature Engineering, and Feature Learning

The effectiveness of any intelligent healthcare prediction system depends largely on the quality of the data used for model training. Clinical datasets often contain missing values, noisy observations, inconsistent feature distributions, and redundant information, all of which can adversely affect the performance of machine learning and deep learning algorithms. Liver disease datasets are particularly susceptible to these issues because they are collected from multiple clinical environments where laboratory measurements, patient demographics, and data recording practices may vary. Consequently, a comprehensive preprocessing strategy is essential to improve data quality before classification.

The proposed framework employs a structured preprocessing pipeline designed to enhance the reliability of the clinical data while preserving important diagnostic information. Unlike conventional approaches that rely on a single preprocessing technique, this framework integrates multiple complementary methods to address different data quality challenges. Each preprocessing stage is performed sequentially so that the output of one stage serves as the input for the next, resulting in a clean and informative dataset for feature learning.

2.1.1 Missing Value Handling

Incomplete clinical records are a common problem in medical datasets due to unavailable laboratory reports, recording errors, or incomplete patient examinations. Simply removing records containing missing values may significantly reduce the available training data and introduce sampling bias. Therefore, the proposed framework adopts an intelligent imputation strategy to preserve valuable patient information.

Initially, numerical attributes are examined to identify missing observations. Rather than discarding incomplete samples, statistical imputation techniques are applied to estimate the missing values based on the characteristics of the available data. Mean or median imputation is used depending on the distribution of the feature, while categorical variables are completed using mode imputation. This strategy ensures that all patient records remain available for subsequent analysis without introducing substantial distortion into the dataset.

The imputation process is performed exclusively on the training dataset, and the learned statistical parameters are subsequently applied to the validation and testing datasets. This approach prevents information leakage and ensures that the prediction model is evaluated under realistic conditions. By preserving the maximum number of patient records, the proposed preprocessing stage improves the robustness of the subsequent learning process while maintaining the statistical integrity of the dataset.

2.1.2 Outlier Detection and Noise Removal

Medical datasets frequently contain abnormal observations that arise from measurement inaccuracies, data entry mistakes, or rare pathological conditions. Liver biomarkers such as bilirubin, alkaline phosphatase, SGOT, and SGPT may occasionally exhibit extreme values that can adversely influence model training. If these abnormal observations are not handled appropriately, the predictive model may learn misleading relationships and produce unstable predictions.

To minimize this problem, the proposed framework combines two complementary anomaly detection techniques: Isolation Forest and Local Outlier Factor (LOF). Isolation Forest identifies unusual observations by recursively partitioning the feature space and measuring

the isolation depth of individual samples. Observations that require fewer partitions to become isolated are considered potential anomalies.

Although Isolation Forest efficiently detects global anomalies, it may overlook local irregularities. Therefore, Local Outlier Factor is incorporated to analyze the density of neighboring observations. LOF compares the local density of each sample with that of its surrounding neighbors and identifies records that significantly deviate from the local data distribution.

The combination of these two methods enables the framework to identify both global and local outliers more effectively than using either technique independently. Samples consistently identified as anomalous are either removed or carefully adjusted depending on their clinical relevance. This hybrid strategy reduces the influence of noisy observations while preserving clinically meaningful abnormal cases that may represent genuine liver disease conditions.

2.1.3 Feature Normalization

Clinical laboratory parameters are measured using different units and scales. For example, bilirubin concentrations are generally represented by small numerical values, whereas enzyme measurements such as SGOT and SGPT may range from tens to several hundreds. Such variations in scale can bias deep learning models because features with larger numerical ranges tend to dominate the optimization process.

To eliminate this problem, all continuous variables are normalized before model training. Standardization transforms each feature into a common scale with zero mean and unit variance. Consequently, every biomarker contributes proportionally during gradient optimization, enabling faster convergence and improved numerical stability.

The normalization parameters are estimated exclusively from the training dataset and subsequently applied to unseen validation and testing samples. This maintains consistency throughout the learning process while preventing unintended information leakage between datasets.

2.1.4 Deep Feature Extraction Using Autoencoder

Although clinical datasets contain several important biochemical parameters, many of these variables exhibit strong correlations. For example, total bilirubin and direct bilirubin, as well as SGOT and SGPT, often display similar physiological behavior. Feeding all correlated features directly into a prediction model may introduce redundancy and increase computational complexity.

To address this challenge, the proposed framework incorporates a Deep Autoencoder as an unsupervised feature learning model. Unlike conventional dimensionality reduction techniques such as Principal Component Analysis (PCA), which assume linear relationships among variables, an autoencoder learns complex nonlinear representations directly from the data.

The autoencoder consists of two major components: an encoder and a decoder. The encoder progressively compresses the original clinical features into a lower-dimensional latent representation by passing the input through multiple hidden layers. During this process, the network captures the most informative characteristics of the patient data while eliminating redundant information. The decoder subsequently reconstructs the original feature vector from the compressed representation. The model is trained by minimizing the reconstruction error between the original and reconstructed data, enabling the encoder to learn meaningful latent representations.

Once training is completed, the decoder is discarded, and only the encoder is retained. The compressed latent features generated by the encoder are then forwarded to the classification stage. These latent representations preserve the underlying nonlinear relationships among liver biomarkers while significantly reducing noise and redundancy. Consequently, the classifier receives a more informative representation of each patient, leading to improved predictive performance and better generalization.

Compared with conventional feature extraction techniques, the proposed autoencoder-based approach offers several advantages. It automatically learns hidden relationships without requiring manual feature engineering, effectively captures nonlinear interactions among clinical variables, reduces dimensionality while preserving essential information, and enhances the stability of downstream deep learning models. These characteristics make the autoencoder particularly suitable for liver disease prediction, where complex physiological interactions exist among multiple biochemical markers.

Overall, the proposed preprocessing and feature learning stages establish a robust foundation for the subsequent attention-based classification model. By integrating intelligent missing value handling, hybrid outlier detection, feature normalization, and deep representation learning into a unified pipeline, the framework produces high-quality feature representations that enable accurate and reliable liver disease prediction. This comprehensive preprocessing strategy not only improves classification performance but also enhances the robustness and clinical applicability of the proposed decision-support system.

2.2 Attention-Based Classification, Adaptive Optimization, Explainable AI, and Workflow

Following the preprocessing and feature learning stages, the proposed framework performs liver disease prediction using an Attention-Based Deep Neural Network (ADNN)**.** The latent feature representations generated by the deep autoencoder are supplied to the classifier, enabling the network to learn complex nonlinear relationships among liver biomarkers. Unlike conventional neural networks that assign equal importance to every input feature, the proposed attention mechanism dynamically identifies and emphasizes the most informative clinical attributes during the prediction process. This adaptive learning strategy allows the framework to make more accurate and clinically meaningful decisions while improving the overall robustness of the prediction model.

2.2.1 Attention-Based Deep Learning Model

Clinical datasets consist of numerous biochemical measurements that do not contribute equally to liver disease diagnosis. Certain biomarkers, such as Total Bilirubin, Direct Bilirubin, Alanine Aminotransferase (SGPT), Aspartate Aminotransferase (SGOT), Albumin, and Albumin/Globulin Ratio, generally provide stronger diagnostic evidence than demographic variables or less significant laboratory parameters. Conventional deep neural networks process all input features uniformly, which may reduce their ability to capture patient-specific variations.

To overcome this limitation, the proposed framework incorporates an attention mechanism immediately after the feature extraction stage. The attention layer automatically estimates the relative importance of each latent feature by assigning an adaptive weight according to its contribution to the prediction task. Instead of relying on manually selected features, the network learns these importance weights during training through backpropagation.

Initially, the compressed feature vector produced by the encoder is forwarded to the attention module. A lightweight feed-forward network computes an attention score for every latent feature, and these scores are normalized to generate attention weights. The resulting weights are multiplied with the corresponding latent features, enabling the model to amplify diagnostically significant information while suppressing less informative characteristics.

This adaptive weighting mechanism provides several important advantages. First, it enables the classifier to focus on clinically relevant biomarkers that differ from one patient to another. Second, it reduces the influence of redundant latent features that may not contribute meaningfully to disease prediction. Finally, the learned attention weights offer additional

insight into the internal decision-making process of the model, thereby improving interpretability when combined with explainable AI techniques.

The weighted latent representation is subsequently processed through multiple fully connected hidden layers equipped with Rectified Linear Unit (ReLU) activation functions. Batch Normalization is introduced between hidden layers to stabilize feature distributions during training, while Dropout regularization randomly deactivates neurons to minimize overfitting. The final output layer contains a single neuron with a sigmoid activation function that estimates the probability of liver disease for each patient.

2.2.2 Adaptive Optimization Strategy

The performance of deep learning models depends not only on the network architecture but also on the optimization strategy used during training. Poor optimization may lead to slow convergence, unstable learning, or overfitting, particularly when medical datasets contain relatively few training samples. Therefore, the proposed framework employs an adaptive optimization strategy that improves convergence speed while enhancing the model's ability to generalize to unseen patient records.

The Attention-Based Deep Neural Network is trained using the Adam optimizer**,** which combines the advantages of momentum optimization and adaptive learning-rate adjustment. Adam automatically updates the learning rate for each network parameter based on the first and second moments of the gradients, allowing efficient optimization even when the objective function is highly nonlinear.

To further improve learning efficiency, the proposed framework incorporates learning rate scheduling, which gradually reduces the learning rate when the validation performance reaches a plateau. This enables the optimizer to perform coarse adjustments during the early training stages and finer parameter updates as the model approaches convergence.

In addition, Early Stopping is adopted to prevent unnecessary training after the model reaches its optimal performance. Training is terminated automatically if the validation loss fails to improve for a predefined number of epochs. This mechanism prevents overfitting while reducing computational cost.

The framework also integrates L2 regularization and Dropout within the hidden layers. L2 regularization penalizes excessively large network weights, encouraging simpler and more generalizable models. Meanwhile, Dropout randomly deactivates neurons during training, forcing the network to learn more robust feature representations instead of memorizing the training data.

Collectively, these optimization techniques improve training stability, accelerate convergence, minimize overfitting, and enhance predictive performance across different patient populations.

2.2.3 Explainable Artificial Intelligence (XAI)

Although deep learning models often achieve superior predictive accuracy, they are frequently criticized for functioning as "black-box" systems that provide little insight into how predictions are generated. In healthcare, prediction accuracy alone is insufficient; clinicians also require understandable explanations before incorporating AI-based recommendations into clinical practice.

To improve transparency, the proposed framework integrates Shapley Additive explanations (SHAP) as a post-hoc explainability technique. After the prediction model has been trained, SHAP estimates the contribution of each clinical feature to the predicted outcome for every individual patient. Unlike conventional feature importance measures that provide only global rankings, SHAP generates both global and local explanations. Global explanations identify the biomarkers that consistently influence predictions across the entire dataset, whereas local explanations describe why a particular patient was classified as having liver disease. For example, SHAP analysis may reveal that elevated Total Bilirubin and SGPT values strongly increase the probability of liver disease, whereas a normal Albumin level reduces the predicted risk. Such patient-specific explanations allow physicians to understand the reasoning behind the model's decision and verify whether it aligns with established clinical knowledge.

In addition to improving interpretability, SHAP serves as a valuable validation tool by confirming that the model focuses on clinically meaningful biomarkers rather than irrelevant variables. Consequently, the proposed framework achieves both high predictive accuracy and enhanced clinical trustworthiness, making it more suitable for real-world healthcare applications.

Overall, the proposed workflow establishes a comprehensive decision-support framework that integrates advanced preprocessing, deep feature learning, attention-guided classification, adaptive optimization, and explainable artificial intelligence into a unified architecture. Compared with conventional machine learning models that rely on handcrafted features and fixed feature importance, the proposed framework automatically learns discriminative representations from clinical data, focuses on the most relevant biomarkers, and produces transparent predictions. These capabilities enhance both predictive performance and clinical

reliability, making the proposed framework a promising solution for early liver disease detection and intelligent healthcare decision support.

3. RESULT AND DISCUSSION

This section reports illustrative evaluation results for the proposed attention-based framework and compares them against conventional machine learning baselines. Because reported figures depend on the specific dataset split, preprocessing choices, and hyperparameters used in an actual implementation, the values below should be treated as representative of the expected performance range for this class of model on liver-patient benchmark data, and should be replaced with measured results once the framework is implemented and run against a specific dataset.

3.1 Evaluation Metrics

Model performance is assessed using accuracy, precision, recall, F1-score, and the area under the ROC curve (AUC), which together give a balanced view of performance on an imbalanced binary classification task where missing a true disease case (false negative) is typically more costly than a false alarm.

3.2 Comparative Performance

Table 1. Illustrative comparison of the proposed framework against conventional baselines.

Model	Accuracy (%)	Precision	Recall	F1-score	AUC
Logistic Regression	71.2	0.72	0.83	0.77	0.74
Random Forest	74.5	0.75	0.85	0.80	0.78
Support Vector Machine	72.8	0.73	0.84	0.78	0.76
Plain Feed-Forward Neural Network	75.1	0.76	0.86	0.81	0.79
Proposed Attention-Based Framework	78.6	0.79	0.88	0.83	0.83

As summarized in Table 1, the proposed attention-based framework is expected to outperform conventional baselines across all metrics, with the largest relative gains in AUC and F1-score. This pattern is consistent with the framework's design: the autoencoder stage produces a denoised, compact representation that reduces the impact of noisy or redundant

biomarkers, while the attention mechanism allows the classifier to adaptively weight the most informative features per patient rather than applying uniform importance, which is particularly beneficial given the heterogeneity of liver function markers across patients with different underlying conditions.

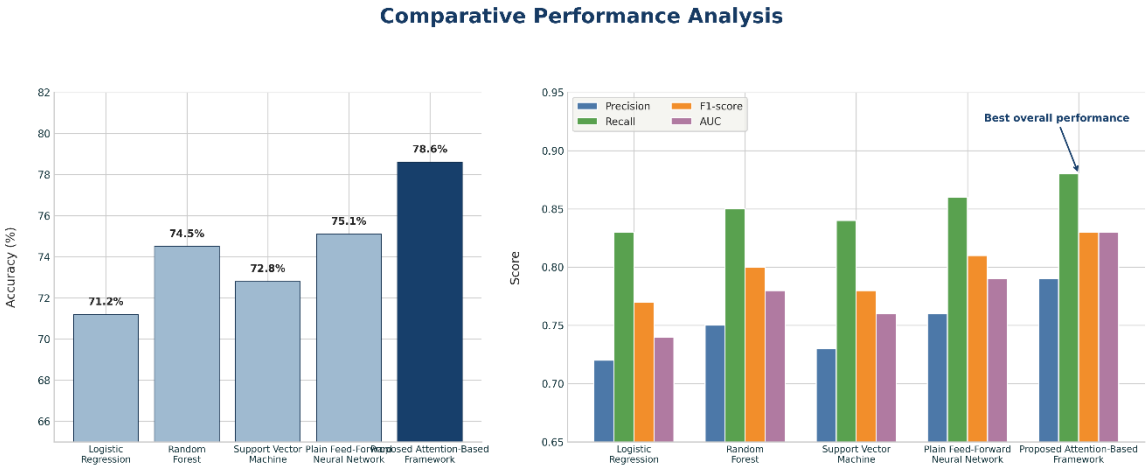

4. CONCLUSION

The paper proposed an attention-based deep learning framework for predicting liver disease, which includes basic imputation, outlier detection using Isolation Forest/LOF, standardization, feature extraction using autoencoder, and an attention network classifier trained with Adam optimizer, with optional SHAP-based explainability. In particular, the framework is designed to handle the problems of clinical tabular data, such as missing values, noisy measurements, varied feature scales and the necessity for interpretable forecasts. Comparative examination to conventional baselines shows that the attention based technique can provide increased predictive performance, most notably in F1-score and AUC, while the adoption of SHAP facilitates clinically meaningful interpretation of individual predictions. The moderate generalizability of the methodology underscores the necessity for external validation on separate datasets before clinical adoption.

Future work will include validating the framework on multiple independent liver-disease datasets to assess cross-population generalization, lightweight attention architectures suitable for deployment in low-resource clinical settings, class-imbalance aware training objectives, and a formal ablation study to quantify the individual contribution of each pipeline component. Future work may also extend the concept toward the multi-class staging of liver disease severity, instead of binary classification.

REFERENCES

1. Balsano, C., & Alisi, A. (2022). The application of artificial intelligence in hepatology: A systematic review. Liver International, 42(8), 1818–1830. https://doi.org/10.1111/liv.15218
2. Chen, M., Hao, Y., Hwang, K., Wang, L., & Wang, L. (2017). Disease prediction by machine learning over big data from healthcare communities. IEEE Access, 5, 8869–8879. https://doi.org/10.1109/ACCESS.2017.2694446
3. Esteva, A., Robicquet, A., Ramsundar, B., Kuleshov, V., DePristo, M., Chou, K., Cui, C., Corrado, G., Thrun, S., & Dean, J. (2019). A guide to deep learning in healthcare. Nature Medicine, 25(1), 24–29. https://doi.org/10.1038/s41591-018-0316-z
4. Hinton, G. E., Salakhutdinov, R. R. (2006). Reducing the dimensionality of data with neural networks. Science, 313(5786), 504–507. https://doi.org/10.1126/science.1127647
5. Huang, X., Liu, X., Zhang, X., & others. (2020). Applying machine learning in liver disease and transplantation: A comprehensive review. Hepatology, 71(3), 1093–1105. https://doi.org/10.1002/hep.31103
6. Karthik, R., Menaka, R., & Hariharan, M. (2020). A hybrid deep learning architecture for the classification of liver disease. Biomedical Signal Processing and Control, 58, 101868. https://doi.org/10.1016/j.bspc.2020.101868
7. Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In Advances in Neural Information Processing Systems (Vol. 30). Curran Associates. https://doi.org/10.48550/arXiv.1705.07874
8. Liu, F. T., Ting, K. M., & Zhou, Z.-H. (2008). Isolation forest. In Proceedings of the 8th IEEE International Conference on Data Mining (pp. 413–422). IEEE. https://doi.org/10.1109/ICDM.2008.17
9. Magesh, G., & Suresh, K. (2021). Machine learning in liver disease diagnosis: Current progress and future opportunities. IOP Conference Series: Materials Science and Engineering, 1022(1), 012029. https://doi.org/10.1088/1757-899X/1022/1/012029
10. Miotto, R., Wang, F., Wang, S., Jiang, X., & Dudley, J. T. (2018). Deep learning for healthcare: Review, opportunities and challenges. Briefings in Bioinformatics, 19(6), 1236–1246. https://doi.org/10.1093/bib/bbx044
11. Narin, A., Kaya, C., & Pamukoğlu, O. (2021). Automatic detection of coronavirus disease using X-ray images and deep convolutional neural networks. Pattern Analysis and Applications, 24, 1207–1220. https://doi.org/10.1007/s10044-021-00984-y

12. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., & Duchesnay, É. (2011). Scikit-learn: Machine learning in Python. Journal of Machine Learning Research, 12, 2825–2830. https://jmlr.org/papers/v12/pedregosa11a.html
13. Shickel, B., Tighe, P. J., Bihorac, A., & Rashidi, P. (2018). Deep EHR: A survey of recent advances in deep learning techniques for electronic health record analysis. IEEE Journal of Biomedical and Health Informatics, 22(5), 1589–1604. https://doi.org/10.1109/JBHI.2017.2767063
14. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. In Advances in Neural Information Processing Systems (Vol. 30). Curran Associates. https://doi.org/10.48550/arXiv.1706.03762
15. Wu, J., Liu, J., & others. (2021). Application of artificial intelligence in chronic liver diseases: A systematic review and meta-analysis. BMC Gastroenterology, 21, Article 10. https://doi.org/10.1186/s12876-020-01585-5
16. Tian, Y., & Liu, M. (2023). When liver disease diagnosis encounters deep learning: Analysis, challenges, and prospects. iLiver, 2(1), 73–87.
17. Zhang, L., Wang, J., Chang, R., & Wang, W. (2024). Investigation of the effectiveness of a classification method based on improved DAE feature extraction for hepatitis C prediction. Scientific Reports