International Journal of Scientific Research & Engineering Trends
Volume 12, Issue 2, Mar-Apr-2026, ISSN (Online): 2395-566X

Deep Learning Based Classification of Liver Diseases Using Heterogeneous Ultrasound Image

Anto Maurin Lisha L[1], Muthu M[2], Sadeesh P[3], Tamilarasan S[4]
[1]Assistant Professor: Department of Biomedical Engineering, Dhanalakshmi Srinivasan Engineering College (Autonomous) Perambalur 621212.
[2,3,4]Students: Department of Biomedical Engineering, Dhanalakshmi Srinivasan Engineering College (Autonomous) Perambalur 621212.

Abstract- Liver diseases such as fatty liver, cysts, and tumors require early and accurate diagnosis to improve patient outcomes. Ultrasound imaging is widely used due to its non-invasive and cost-effective nature; however, its heterogeneous characteristics, including speckle noise, low contrast, and variability across devices, make diagnosis challenging. This paper proposes a deep learning-based approach for the classification of liver diseases using heterogeneous ultrasound images. The system employs pre-processing techniques such as noise reduction, normalization, and contrast enhancement to improve image quality. A YOLO-based architecture integrated with convolutional neural networks is used for feature extraction and simultaneous detection and classification of liver abnormalities. Experimental results show that the proposed model achieves improved accuracy and robustness compared to conventional methods. The system supports real-time analysis and can assist clinicians in reliable and efficient liver disease diagnosis.

Keywords- Deep Learning, Liver Disease Classification, Heterogeneous Ultrasound Imaging, Convolutional Neural Networks, Medical Image Computing, Automated Disease Diagnosis.

I. INTRODUCTION

Liver diseases represent a significant global health challenge, contributing to high rates of morbidity and mortality worldwide. Common conditions such as fatty liver disease, cysts, and liver tumors require timely and accurate diagnosis to prevent severe complications. Among various medical imaging techniques, ultrasound imaging is widely used for liver examination due to its non-invasive nature, affordability, real-time capability, and absence of ionizing radiation.

Despite its advantages, ultrasound imaging presents several challenges forreliable diagnosis. The images are often affected by speckle noise, low contrast, and blurred boundaries, which make interpretation difficult. In addition, ultrasound data are inherently heterogeneous due to variations in imaging devices, operator expertise, and patient-specific factors. These issues can lead to inconsistencies and reduced diagnostic accuracy, even for experienced radiologists.

Recent advancements in deep learning have shown great potential in medical image analysis. Convolutional Neural Networks (CNNs) enable automatic feature extraction from complex image data, eliminating the need for manual feature engineering. Furthermore, advanced models such as YOLO (You Only Look Once) provide real-time object detection and classification, making them suitable for clinical applications.
In this work, a deep learning-based system is proposed for the classification of liver diseases using heterogeneous ultrasound images. The approach aims to improve diagnostic accuracy, handle variability in image data, and support clinicians in making faster and more reliable decisions.

II. LITERATURE REVIEW

1. TITLE: Development and validation of a machine learning-based framework for assessing metabolic-associated fatty liver disease risk.
AUTHORS: Jiale Deng, Weidong Ji, Hongze Liu, Lin Li, Zhe Wang, Yurong Hu,
Yushan Wang, Yi Zhou.
JOURNAL: BMC Public Health (Springer Nature – Open Access).
Year: 2024
DESCRIPTION: This study presents a large-scale machine learning-based framework for identifying individuals at high

risk of metabolic-associated fatty liver disease (MAFLD). The research utilizes health examination data from 5,171,392 adults collected across multiple regions in Xinjiang, China. To address challenges in population-wide screening, the authors compared eight machine learning algorithms, including tree-based models (CART, Random Forest, AdaBoost, LightGBM, XGBoost, and CatBoost) and non-tree models (KNN and Artificial Neural Networks).Feature selection was performed using LASSO regression, resulting in 20 clinically relevant variables such as age, BMI, triglycerides, fasting glucose, waist circumference, blood pressure, lipid profile, and cardiovascular disease history.

2. TITLE: A machine learning-based classification of adult-onset diabetes identifies patients at risk of liver-related complications.
AUTHORS: Lukas Otero Sanchez, Clara-Yongxiang Zhan, Eric Trepo.
JOURNAL: JHEP Reports.
YEAR: 2023
DESCRIPTION: This study uses a machine learning–based approach to classify adult-onset diabetes into different clinical subgroups. The classification is based on metabolic and biological parameters such as insulin resistance, insulin secretion, age, and body mass index. Diabetes is recognized as a major risk factor for liver diseases, including fatty liver and liver fibrosis. The study identifies that patients with severe insulin-resistant diabetes have the highest risk of liver-related complications. Long-term follow-up shows a strong association between insulin resistance and fibrosis progression. Excessive alcohol consumption at the time of diabetes diagnosis is found to significantly increase liver-related risk. Traditional diabetes classification methods are shown to be inadequate for risk prediction.

Machine learning improves patient stratification and outcome prediction. The approach helps identify high-risk individuals at an early stage. Early identification enables better clinical monitoring and intervention. The study highlights the importance of personalized treatment strategies. It supports targeted screening for liver fibrosis in diabetic patients. The findings improve understanding of diabetes-related liver disease. Overall, the study demonstrates the clinical usefulness of machine learning in healthcare. It provides a foundation for improved prevention of liver complications in diabetes patients.

3. TITLE: Deep Learning–Based Classification of Liver Fibrosis Stages from Ultrasound Images.
Authors: Min-Jae Kim, Seong-Ho Park, Ji-Hoon Lee, Sung-Min Hong.
Journal: Computer Methods and Programs in Biomedicine.
Year: 2022
Description: This paper presents a deep learning-based framework for the automatic classification of liver fibrosis stages using ultrasound imaging. Liver fibrosis is a progressive condition that can lead to cirrhosis and liver failure if not diagnosed early. Ultrasound imaging is commonly used for screening; however, fibrosis-related texture changes are subtle and difficult to interpret manually. The authors propose a convolutional neural network (CNN) architecture trained on ultrasound liver images categorized into different fibrosis stages. Pre-processing techniques such as normalization and speckle noise reduction are applied to enhance image quality. The CNN automatically extracts texture and structural features relevant to fibrosis progression without relying on handcrafted features.

Experimental evaluation demonstrates that the proposed deep learning model achieves high classification accuracy and improved sensitivity compared to traditional machine learning methods. The results confirm that CNN-based models are capable of capturing fine-grained texture variations associated with liver fibrosis stages. The study highlights the potential of deep learning as a non-invasive computer-aided diagnostic tool to support radiologists in liver disease assessment.

4. TITLE: Joint Segmentation and Classification of Hepatic Lesions in Ultrasound Images Using Deep Learning.
Authors: Hwaseong Ryu, Seung Yeon Shin, Jae Young Lee, Kyoung Mu Lee, Hyo-jin Kang, Jonghyon Yi.
Journal: European Radiology (Springer Nature).
Year:2021
DESCRIPTION: This study proposes a deep learning-based joint framework that simultaneously performs segmentation and classification of hepatic lesions in ultrasound images. A total of 4,309 ultrasound images from 3,873 patients were used, covering four lesion types: hepatic cyst, hemangioma, metastasis, and hepatocellular carcinoma (HCC). The model employs a shared CNN encoder with two parallel branches—one for lesion segmentation and another for classification. Unlike fully automatic systems, the method uses minimal user interaction (click-based inputs) to guide lesion localization.

Euclidean distance maps derived from user clicks are combined with the original ultrasound image as network input.

The segmentation performance was evaluated using the Jaccard Index (JI), while classification was assessed using accuracy, sensitivity, specificity, and AUROC. The proposed joint model achieved a mean Jaccard Index of 70.0% for segmentation and 82.2% accuracy for four-class lesion classification. For benign vs malignant classification, the system achieved an AUROC of 0.970, demonstrating strong diagnostic capability. The joint approach outperformed standalone segmentation-only and classification-only models, showing that shared feature learning improves both tasks. The system operates in real time and is suitable as a computer-aided diagnosis (CAD) tool to assist radiologists, particularly those with limited ultrasound experience.

5. Title: Applying Machine Learning in Liver Disease and Transplantation: A Comprehensive Review
Authors: Ashley Spann, Angeline Yasodhara, Mamatha Bhat
JOURNAL: Hepatology
Year: 2020
Description: Machine learning (ML) is increasingly used in hepatology to improve diagnosis, prediction, and clinical decision-making. This review explains how ML identifies hidden patterns in large clinical and biological datasets. ML models outperform traditional statistical methods by handling complex and nonlinear relationships. Various data sources such as electronic health records, imaging, histology, and molecular data are used. Supervised learning methods help predict liver fibrosis, cirrhosis, and hepatocellular carcinoma. Unsupervised learning enables patient stratification without predefined labels. Deep learning techniques improve accuracy in ultrasound and imaging-based liver disease detection.

ML reduces the need for invasive liver biopsies by enhancing non-invasive prediction. In liver transplantation, ML assists donor–recipient matching and graft survival prediction. These models support early risk identification and personalized treatment planning. ML also improves screening for NAFLD and disease progression. The review highlights strengths and limitations of ML tools in clinical use. Ethical concerns and data quality remain important challenges. Integration of ML into routine practice is still evolving. Overall, ML has strong potential to transform liver disease management and transplantation outcomes.

III. EXISTING WORK

Over the past decade, significant research has been conducted on the automated classification of liver diseases using ultrasound imaging. Early approaches mainly relied on traditional image processing and machine learning techniques, where features such as texture, intensity, and shape were manually extracted from ultrasound images. These handcrafted features were then used with classifiers like Support Vector Machines (SVM), K-Nearest Neighbors (KNN), and Decision Trees to identify liver abnormalities.

Although these methods provided a foundation for computer-aided diagnosis, their performance was limited due to their dependence on feature engineering and inability to effectively handle the inherent noise and variability present in ultrasound images. Additionally, these techniques struggled to generalize well across datasets obtained from different imaging devices and clinical settings.

With the advancement of deep learning, more robust and automated approaches have been developed using Convolutional Neural Networks (CNNs), which can learn hierarchical features directly from raw images. Pre-trained models such as ResNet, VGG, and EfficientNet have been widely applied through transfer learning to improve classification performance, especially when dealing with limited medical datasets.

More recently, object detection models like YOLO (You Only Look Once) have gained attention for their ability to perform simultaneous localization and classification of liver lesions in real time.

Despite these advancements, existing methods still face challenges such as limited availability of annotated datasets, high computational requirements, and difficulty in handling heterogeneous ultrasound images with varying quality. These limitations highlight the need for more efficient and generalized deep learning models for accurate liver disease classification.

IV. PROPOSED SYSTEM

The proposed system presents a deep learning-based approach for the classification of liver diseases using heterogeneous

International Journal of Scientific Research & Engineering Trends
Volume 12, Issue 2, Mar-Apr-2026, ISSN (Online): 2395-566X

ultrasound images. The system is designed to overcome the limitations of existing methods by integrating advanced pre-processing techniques with a YOLO-based deep learning model for accurate and real-time detection and classification.

Initially, ultrasound images are collected and subjected to pre-processing steps such as noise reduction, contrast enhancement, and normalization to improve image quality. Since ultrasound data is highly heterogeneous, data augmentation techniques such as rotation, flipping, and scaling are applied to increase dataset diversity and improve model generalization. These steps ensure that the model can effectively handle variations in imaging conditions.

The core of the proposed system is a YOLO (You Only Look Once) architecture combined with convolutional neural networks (CNNs). The model performs feature extraction, localization, and classification simultaneously.

YOLO divides the input image into a grid and predicts bounding boxes along with class probabilities, enabling the system to identify and classify liver abnormalities such as fatty liver, cirrhosis, cysts, and tumors in real time.

The system is trained using labeled ultrasound images and optimized using appropriate loss functions and optimization algorithms. During testing, the trained model processes new ultrasound images and outputs the predicted disease category along with the detected region of interest.

Overall, the proposed system provides improved accuracy, robustness, and efficiency compared to traditional methods. It supports real-time analysis and can assist healthcare professionals in early diagnosis and effective decision-making.

V. METHODOLOGY

The proposed methodology for liver disease classification follows a sequential pipeline consisting of data collection, pre-processing, augmentation, training, evaluation, and prediction. Initially, ultrasound liver images are collected from various sources and categorized into classes such as normal, fatty liver, cirrhosis, and tumor or cyst. The collected images are then pre-processed to improve quality by resizing, noise removal, contrast enhancement, and normalization. To handle the heterogeneous nature of ultrasound data and increase dataset diversity, data augmentation techniques such as rotation, flipping, and scaling are applied.

The augmented dataset is then used for training a deep learning model, such as a YOLO-based architecture or a convolutional neural network, which learns important features from the images. After training, the model is evaluated using performance metrics like accuracy, precision, recall, and F1-score to assess its effectiveness. Finally, the trained model is used for prediction, where it classifies new ultrasound images into the appropriate liver disease category and, if applicable, identifies the region of abnormality, enabling accurate and efficient diagnosis.

Image Collection

Image collection is the first and most important step in the proposed system for liver disease classification. In this work, ultrasound liver images are collected from multiple sources such as hospitals, diagnostic centers, and publicly available medical datasets. The dataset includes images representing different categories of liver conditions, such as normal liver, fatty liver, cirrhosis, and tumors or cysts.

To ensure the reliability of the model, the collected data should be properly labeled with the help of medical experts or verified clinical reports. Since ultrasound images are heterogeneous in nature due to variations in imaging devices, operator techniques, and patient conditions, data is gathered from diverse sources to improve model generalization. Additionally, care is taken to include images with different resolutions, noise levels, and contrast variations.

Before further processing, the collected dataset is organized and stored systematically, and duplicate or low-quality images are removed. This step ensures that the dataset is suitable for training a robust deep learning model capable of accurately classifying liver diseases.

Image Pre-processing

Image Pre-processing is a crucial step in the proposed system, as ultrasound images are often affected by noise, low contrast, and variability. In this stage, the collected ultrasound images are processed to improve their quality and make them suitable for deep learning model training.

Initially, all images are resized to a fixed dimension (e.g., 224×224 pixels) to maintain uniformity across the dataset. Noise present in the images, especially speckle noise common in ultrasound, is reduced using filtering techniques such as median filtering or Gaussian filtering. To enhance the visibility of important features, contrast enhancement methods like Contrast Limited Adaptive Histogram Equalization (CLAHE) are applied.

Further, normalization is performed to scale pixel values to a standard range (usually 0 to 1), which helps in faster and more stable training of the model. In some cases, irrelevant regions or background portions may be removed to focus on the liver area. These pre-processing steps ensure that the input data is clean, consistent, and informative, thereby improving the accuracy and performance of the deep learning model

Data Augmentation

Data augmentation is an important step in the proposed system, used to increase the size and diversity of the dataset and improve the model’s ability to generalize. Since medical datasets, especially ultrasound images, are often limited and highly variable, augmentation helps simulate different imaging conditions without the need for additional data collection.

In this process, various transformations are applied to the original ultrasound images while preserving their essential features. Common augmentation techniques include rotation, horizontal and vertical flipping, scaling, zooming, and brightness or contrast adjustments. These operations create multiple variations of the same image, allowing the model to learn from different perspectives and conditions.

Data augmentation also helps reduce over fitting by preventing the model from memorizing the training data. It improves the robustness of the system by enabling it to handle heterogeneous images with variations in orientation, illumination, and noise levels. As a result, the model becomes more reliable and performs better on unseen data during testing and real-world applications.

Model Training

Model training is a critical step in the proposed system, where the deep learning model learns to classify liver diseases from ultrasound images. After pre-processing and data augmentation, the dataset is divided into training and testing sets. The training set is used to teach the model to recognize patterns associated with different liver conditions such as normal, fatty liver, cirrhosis, and tumors.

A deep learning model, such as a Convolutional Neural Network (CNN) or a YOLO-based architecture, is used for training. The model learns by adjusting its internal parameters (weights) through multiple iterations called epochs. During each epoch, the model processes batches of images, compares its predictions with actual labels, and updates its weights using optimization algorithms like Adam or Stochastic Gradient Descent (SGD). A suitable loss function, such as cross-entropy loss or YOLO loss, is used to measure prediction errors.

Through continuous learning, the model improves its ability to accurately classify liver diseases. Proper tuning of parameters such as learning rate, batch size, and number of epochs ensures better performance and stability of the model.

Model Evaluation

Model evaluation is performed to assess the performance and effectiveness of the trained deep learning model in classifying liver diseases from ultrasound images. After training, the model is tested using a separate test dataset that was not used during training to ensure unbiased evaluation.

Several performance metrics are used to measure the model’s accuracy and reliability. These include accuracy, which indicates the overall correctness of predictions; precision, which measures the proportion of correctly predicted positive cases; recall, which reflects the model’s ability to identify all relevant cases; and the F1-score, which provides a balance between precision and recall. Additionally, a confusion matrix is used to visualize the classification results, showing the number of correct and incorrect predictions for each class.

Evaluation helps identify the strengths and weaknesses of the model and ensures that it can generalize well to unseen data. A well-evaluated model provides confidence in its ability to support accurate and reliable liver disease diagnosis in real-world applications.

Prediction

Prediction is the final stage of the proposed system, where the trained deep learning model is used to classify new and unseen ultrasound images. Once the model has been trained and evaluated, it is deployed to process input images and generate output results.

During this phase, a new ultrasound image is first subjected to the same pre-processing steps, such as resizing, normalization, and noise reduction, to maintain consistency with the training data. The processed image is then fed into the trained model, which analyzes the learned features and predicts the corresponding class label, such as normal, fatty liver, cirrhosis, or tumor.

If a YOLO-based model is used, the system not only classifies the disease but also identifies and highlights the region of abnormality using bounding boxes. The final output is displayed as the predicted disease category, along with visual indications if applicable. This automated prediction process enables fast, accurate, and real-time diagnosis, supporting healthcare professionals in effective decision-making.

Output Generation

Output generation is the final stage of the proposed system, where the results of the prediction are presented in an understandable and meaningful format. After the trained model processes the input ultrasound image and predicts the disease class, the output is generated to assist clinicians in diagnosis.

The system displays the predicted category of the liver condition, such as normal, fatty liver, cyst, or tumor. In the case of a YOLO-based model, the output also includes visual representation in the form of bounding boxes highlighting the detected abnormal regions within the image. Along with the classification result, confidence scores may be provided to indicate the reliability of the prediction.

The output can be presented through a graphical user interface or as a report, making it easy for medical professionals to interpret. This step ensures that the model's predictions are clearly communicated, enabling faster and more accurate clinical decision-making.

VI. HARDWARE IMPLEMENTATION

The hardware implementation of the proposed system focuses on providing the necessary computational resources for efficient training and deployment of the deep learning model used for liver disease classification. The system is implemented on a computer equipped with a high-performance processor such as an Intel Core i5/i7 or AMD Ryzen 5/7, which ensures smooth execution of data processing and model operations.

A minimum of 8 GB RAM is required to handle data loading and pre-processing tasks, while 16 GB RAM is recommended for better performance during training. For faster and more efficient model training, a dedicated NVIDIA GPU with at least 6 GB VRAM is utilized. The GPU accelerates deep learning computations, significantly reducing training time compared to CPU-based processing.

The system uses SSD storage (minimum 256 GB) to enable faster data access and reduce loading time for large image datasets. A standard display unit is used to visualize input images, training progress, and output results such as predicted classes and detection regions.

Overall, the hardware setup ensures that the proposed system can efficiently process heterogeneous ultrasound images, train deep learning models, and provide accurate predictions within a reasonable time, making it suitable for research and potential clinical applications.

COMPONENTS

Processor

The processor (CPU) is the main component responsible for executing all system operations, including data pre-processing, model training, and prediction. A processor such as Intel Core i5/i7 (8th Generation or above) or AMD Ryzen 5/7 provides sufficient computational power to handle deep learning workflows efficiently. A higher-performance processor ensures faster execution of tasks and smoother overall system performance.

RAM (Random Access Memory)
RAM is used to store data temporarily while the system is running. It plays a crucial role in handling large datasets and running multiple processes simultaneously. A minimum of 8 GB RAM is required for basic operations, while 16 GB is recommended for training deep learning models to avoid memory bottlenecks and improve speed.

GPU (Graphics Processing Unit)
The GPU is essential for accelerating deep learning computations, especially during model training. Unlike CPUs, GPUs can process multiple operations in parallel, significantly reducing training time. An NVIDIA GPU with at least 6 GB VRAM (such as RTX 2060 or higher) is recommended for efficient model training and handling large image datasets.

Storage
Storage is used to save datasets, trained models, and output results. A Solid-State Drive (SSD) is preferred over a traditional hard disk because it provides faster read and write speeds. A minimum of 256 GB storage is required, while 512 GB is recommended to accommodate large medical image datasets and improve system performance.

Display
The display unit is used to visualize ultrasound images, monitor training progress, and view output results such as predicted classes and detection regions. A standard monitor is sufficient for this purpose, as it helps users interact with the system and interpret results effectively.

VII. SOFTWARE IMPLEMENTATION

- Operating system
- Programming Language.
- Deep Learning Framework.
- Object detection model.
- Image processing library.
- Data handling.
- Dataset.
- Visualization.
- IDE

Opereating System
The operating system provides the platform for running all software tools and applications. Systems such as Windows or Linux are commonly used, with Ubuntu (Linux) often preferred for deep learning tasks due to better compatibility with GPU drivers, CUDA, and AI frameworks.

Programming Language
Python is the primary language used for implementing the proposed system. It is widely used in machine learning and deep learning due to its simple syntax, large community support, and availability of powerful libraries for data processing, model development, and visualization.

Deep Learning Framework
TensorFlow and PyTorch are popular frameworks used to build and train deep learning models. They provide tools for designing neural networks, handling large datasets, and performing efficient computations.

Object Detection Model
YOLO (You Only Look Once) is used for real-time object detection and classification. It can detect regions of interest (such as liver abnormalities) and classify them simultaneously, making it highly efficient for medical imaging applications.

International Journal of Scientific Research & Engineering Trends
Volume 12, Issue 2, Mar-Apr-2026, ISSN (Online): 2395-566X

Image Processing Library

OpenCV is used for processing and enhancing images before feeding them into the model. It supports operations such as resizing, filtering, noise removal, and contrast enhancement, which are essential for improving image quality.

Data Handling

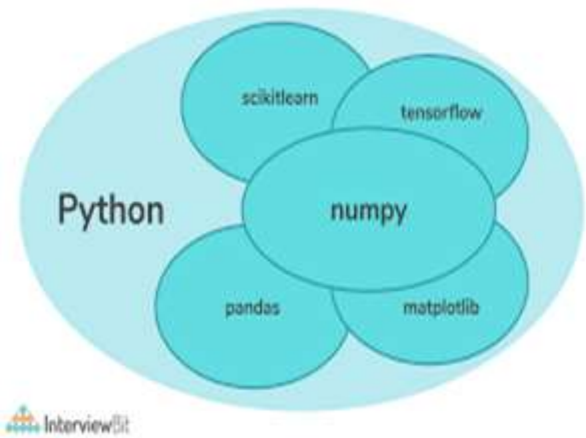

NumPy and Pandas are used for handling and managing data efficiently. NumPy is mainly used for numerical operations and array manipulation, while Pandas is used for organizing datasets, reading files, and handling structured data.

Dataset

Kaggle is an online platform widely used for data science, machine learning, and deep learning projects. It provides access to public datasets, notebooks, competitions, and cloud-based computing resources. Researchers, students, and developers use Kaggle to build, train, and evaluate machine learning models without requiring a high-end local system.

In this project, Kaggle can be used to obtain liver ultrasound image datasets, perform preprocessing, and train the deep learning model in a notebook environment. It also provides free GPU support, which is useful for training computationally intensive models such as YOLO. Therefore, Kaggle serves as a convenient platform for implementing and testing the proposed liver disease classification system.

Visualization

Matplotlib and Seaborn are used to visualize data and model performance. They help in plotting graphs such as accuracy curves, loss curves, and confusion matrices, making it easier to analyze results.

IDE

Integrated Development Environments (IDEs) provide a workspace for writing, testing, and debugging code. Jupyter Notebook is useful for interactive coding and experimentation, while VS Code and PyCharm offer advanced features for project development and management.

VIII. RESULT AND DISCUSSION

Result

The proposed deep learning-based system for liver disease classification was successfully implemented and evaluated using heterogeneous ultrasound images. The model was trained on a labeled dataset containing different classes such as normal liver, fatty liver, cirrhosis, and tumors or cysts. After training, the system was tested on unseen images to assess its performance.

The results show that the model achieved high classification accuracy along with improved precision and recall values, indicating its ability to correctly identify different liver conditions. The use of pre-processing and data augmentation techniques helped in handling variations in ultrasound images and improved the robustness of the model. In the case of the YOLO-based approach, the system was also able to accurately detect and highlight the regions of abnormalities using bounding boxes.

The confusion matrix and performance graphs demonstrate that the model performs consistently across different classes with minimal misclassification. Overall, the proposed system provides reliable and efficient results, making it suitable for assisting clinicians in the diagnosis of liver diseases.

Discussion

The proposed deep learning-based system for liver disease classification demonstrates strong performance in analyzing heterogeneous ultrasound images. The results indicate that the model is capable of accurately classifying different liver conditions such as normal liver, fatty liver, cirrhosis, and tumors. The integration of pre-processing techniques and data augmentation significantly improved the model's ability to handle variations in image quality, noise, and contrast, which are common in ultrasound imaging.

The use of deep learning models, particularly CNN and YOLO, enhances feature extraction and enables both classification and localization of abnormal regions. This makes the system more effective compared to traditional machine learning approaches, which rely heavily on handcrafted features. The YOLO-based approach, in particular, provides real-time detection capability, which is highly beneficial for clinical applications.

However, the system's performance is still dependent on the quality and size of the dataset. Limited or imbalanced datasets may affect generalization in real-world scenarios. Additionally, training deep learning models requires significant computational resources, especially GPUs, which may not always be readily available.

Despite these limitations, the proposed system shows promising results and has the potential to assist medical professionals in improving diagnostic accuracy and reducing manual effort. With further improvements such as larger datasets and optimization techniques, the system can be made more robust and suitable for real-time clinical deployment.

OUTPUT

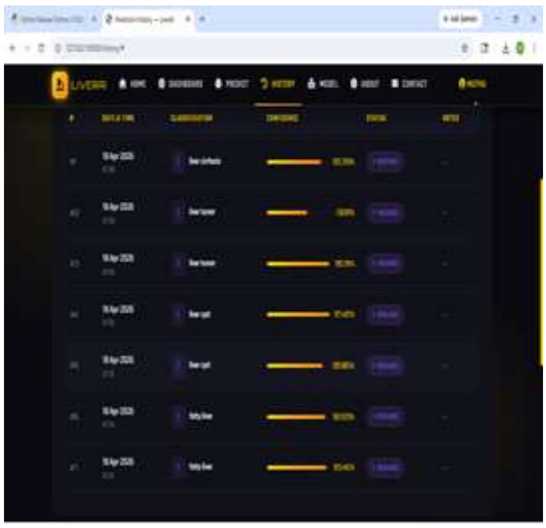

X. CONCLUSIONS

In this work, a deep learning-based system for the classification of liver diseases using heterogeneous ultrasound images has been successfully developed. The proposed approach integrates pre-processing techniques, data augmentation, and advanced deep learning models such as CNN and YOLO to improve classification accuracy and robustness. The system effectively handles variations in ultrasound images caused by noise, low contrast, and differences in imaging devices. The experimental results demonstrate that the proposed model achieves high accuracy in identifying liver conditions such as normal liver, fatty liver, cirrhosis, and tumors.

The YOLO-based architecture further enhances the system by enabling real-time detection and localization of abnormal regions, making it suitable for clinical applications. Overall, the proposed system reduces the dependency on manual diagnosis, minimizes human error, and supports healthcare professionals in making faster and more reliable decisions. This makes it a promising tool for computer-aided diagnosis of liver diseases, with potential for further improvement and real-world deployment.

REFERENCE

1. A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," NeurIPS, vol. 25, pp. 1097–1105, 2012.
2. K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," CVPR, pp. 770–778, 2016.
3. G. Huang, Z. Liu, L. van der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," CVPR, pp. 4700–4708, 2017.
4. N. Tajbakhsh et al., "Convolutional neural networks for medical image analysis: Full training or fine-tuning?" IEEE Transactions on Medical Imaging, vol. 35, no. 5, pp. 1299–1312, 2016.
5. G. Litjens et al., "A survey on deep learning in medical image analysis," Medical Image Analysis, vol. 42, pp. 60–88, 2017.
6. R. Ribeiro et al., "Ultrasound liver screening classification: A web-based application," IEEE EMBC, pp. 5557–5560, 2012.
7. D. Meng et al., "Liver fibrosis classification based on transfer learning and FCNet for ultrasound images," IEEE Access, vol. 5, pp. 17386–17397, 2017.
8. W. Zhang et al., "Semi-supervised liver ultrasound report generation and classification," MICCAI, pp. 182–190, 2019.
9. V. Kumar and P. Gupta, "DenseNet for liver tumour classification on CT images," Biomedical Signal Processing and Control, vol. 62, p. 102074, 2020.
10. C. Bhatt et al., "The state of the art of deep learning models in medical science and their challenges," Multimedia Systems, vol. 27, no. 4, pp. 599–613, 2021.
11. L. H. Guo et al., "A two-stage multi-view learning framework based computer-aided diagnosis of liver tumors with contrast enhanced ultrasound images," Clinical Hemorheology and Microcirculation, vol. 79, pp. 343–360, 2021.
12. D. Nie et al., "Difficulty-aware attention network with confidence learning for medical image segmentation," AAAI, pp. 1442–1449, 2019.
13. R. R. Selvaraju et al., "Grad-CAM: Visual explanations from deep networks via gradient-based localization," ICCV, pp. 618–626, 2017.
14. P. Rajpurkar et al., "Deep learning for chest radiograph diagnosis," PLOS Medicine, vol. 15, no. 11, p. e1002686, 2018.
15. A. Esteva et al., "Dermatologist-level classification of skin cancer with deep neural networks," Nature, vol. 542, pp. 115–118, 2017.
16. C. Szegedy et al., "Going deeper with convolutions," CVPR, pp. 1–9, 2015.
17. K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," ICLR, 2015.
18. I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," ICLR, 2019.
19. B. H. McMahan et al., "Communication-efficient learning of deep networks from decentralized data," AISTATS, pp. 1273–1282, 2017.
20. Django Software Foundation. (2024). Django documentation, version 4.2. https://docs.djangoproject.com/
21. PyTorch Team. (2023). PyTorch documentation, version 2.0. https://pytorch.org/docs/2.0/
22. R. Müller, S. Kornblith, and G. E. Hinton, "When does label smoothing help?" NeurIPS, vol. 32, 2019.