IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

The Role Of AI In Enhancing The Accuracy Of Medical Diagnostics: A Deep Learning Approach In Healthcare

Tarun Kumar
Lecturer
Computer Science & Engineering
Shobhit University Gangoh

Kuldeep Chauhan
Assistant Professor
Computer Science & Engineering
Shobhit University Gangoh

Dr.Varun Bansal
HOD & Professor
Computer Science & Engineering
Shobhit University Gangoh

Abstract: The integration of Artificial Intelligence (AI), particularly deep learning, into medical diagnostics has revolutionized modern healthcare by enhancing accuracy, reducing diagnostic delays, and supporting clinical decision-making processes. This paper explores the transformative role of deep learning methodologies especially convolutional neural networks (CNNs) and recurrent neural networks (RNNs) in advancing the precision of diagnostic tools across a variety of clinical domains. From medical imaging and pathology to genomics and electronic health records, deep learning models have demonstrated capabilities that often rival or exceed those of experienced clinicians in specific diagnostic tasks. We examine real world applications, including AI-driven detection of pneumonia from chest X-rays, skin cancer classification through dermoscopic images, and diabetic retinopathy recognition, highlighting both the performance metrics and the underlying model architectures. Furthermore, the paper addresses the significant benefits of AI integration, such as scalability, consistency, and the potential for deployment in resource-limited settings where trained professionals are scarce.

However, the adoption of AI in healthcare also raises critical challenges. These include data bias, model interpretability, regulatory hurdles, and the ethical considerations surrounding patient data privacy. This paper critically evaluates these challenges, offering insights into the current limitations and proposing strategies for responsible and ethical implementation. Finally, the paper discusses the future outlook of deep learning in diagnostics, emphasizing the shift towards multimodal AI systems that integrate imaging, text, and genomic data for more holistic insights. The potential for personalized diagnostics and seamless integration with clinical workflows signifies a new era in precision medicine.

Keywords: Artificial Intelligence, Deep Learning, Medical Diagnostics, Convolutional Neural Networks, Diagnostic Accuracy, Medical Imaging, Healthcare Technology.

1. Introduction

The exponential growth in healthcare data ranging from medical imaging and genomic sequences to clinical notes and wearable sensor outputs has necessitated the adoption of more advanced analytical tools. Artificial Intelligence (AI), particularly deep learning (DL), has emerged as a transformative force in this regard. Deep learning models have demonstrated the capability to automatically learn complex patterns from large-scale heterogeneous datasets, enabling superior performance in a variety of diagnostic tasks [1], [2]. With architectures such as Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs), DL has revolutionized domains including radiology, dermatology, and ophthalmology, providing automated detection and classification systems that rival expert-level accuracy [3], [4].

Recent advancements have extended beyond traditional architectures to include transformer-based models and attention mechanisms, which have reshaped the paradigm of medical diagnostics. Vision Transformers (ViTs), for example, have shown state-of-the-art results in analyzing complex medical images like histopathology slides and 3D MRI scans [5], [6]. Self-supervised learning techniques now allow models to pretraining on unlabeled medical data, greatly reducing dependency on annotated datasets and expanding accessibility for under-resourced clinical settings [7]. Moreover, the fusion of imaging data with electronic health records (EHR), genomics, and patient-reported outcomes through multimodal AI frameworks is paving the way toward more holistic and context-aware diagnostics [8].

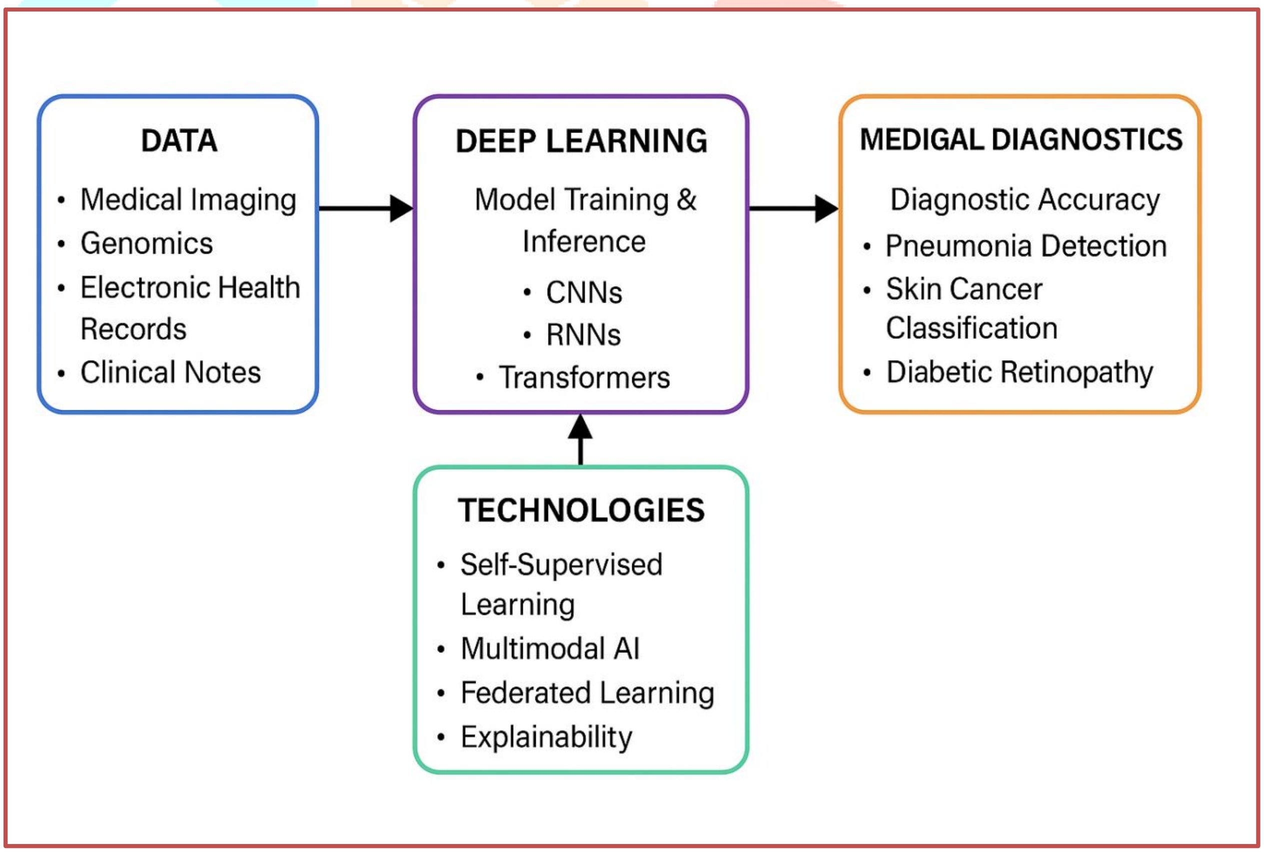

Figure 1: Block Diagram Illustrating the Integration of Deep Learning and AI Technologies in Medical Diagnostics

Despite these advances, the integration of AI into clinical workflows presents multifaceted challenges. One major concern is model generalizability; models trained on homogeneous datasets may underperform when deployed across diverse patient populations, leading to biased or inaccurate results [9]. Additionally, black-box model behavior raises issues of interpretability and clinician trust. The introduction of explainable AI

(XAI) and post-hoc interpretability tools aims to mitigate these concerns, but their clinical acceptance remains a work in progress. Furthermore, federated learning has emerged as a secure paradigm that enables collaborative model training across decentralized medical institutions without the need to share raw patient data, preserving data privacy and regulatory compliance [10].Given the convergence of these cutting-edge technologies, this paper investigates the evolving role of deep learning in medical diagnostics with a focus on clinical accuracy, scalability, and ethical deployment. We provide a critical analysis of current DL models applied in diagnostics, review real-world implementations, and identify key opportunities and risks in integrating these systems into routine healthcare. By exploring applications from AI-based pneumonia detection to diabetic retinopathy and skin cancer classification, this study presents a comprehensive overview of the state-of-the-art and proposes future directions for personalized, explainable, and multimodal diagnostic frameworks[11].

2.Literature Review

The application of AI in medical diagnostics has been widely explored over the last decade, with deep learning models becoming increasingly central in this domain. One of the earliest and most impactful demonstrations was by Rajpurkar et al. [10], who developed **Chex Net**, a deep CNN that outperformed radiologists in pneumonia detection using chest X-ray images. Similarly, Esteva et al. [11] used a CNN trained on over 100,000 dermoscopic images to classify skin lesions; achieving performance comparable to dermatologists. The evolution of deep learning architectures brought forward attention-based and transformer models. Dosovitskiy et al. [12] introduced Vision Transformers (ViTs), which were later adapted for medical imaging tasks due to their ability to capture long-range dependencies in high-resolution images. Building upon this, Chen et al. [13] proposed TransUNet, a hybrid CNN-transformer architecture for medical image segmentation, demonstrating superior performance over traditional U-Net models on various organ and tumor segmentation benchmarks [13].

Multimodal data integration has emerged as another powerful direction, aiming to combine imaging, clinical, and genomic data. Lee et al. [14] reviewed this trend, highlighting how multimodal AI systems offer a more comprehensive understanding of patient health and increase diagnostic reliability. These systems often integrate EHR data and imaging inputs using late-fusion or co-attention mechanisms. In the context of diabetic retinopathy detection, Ting et al. [15] incorporated patient metadata with fundus photographs to significantly improve diagnostic outcomes.

Recent work has also emphasized the importance of model generalizability and data efficiency. Aziza et al. [16] leveraged self-supervised learning techniques on large-scale unlabeled medical datasets, showing that pertained models could be fine-tuned with minimal labeled data to outperform supervised baselines. Likewise, Sarma et al. [17] demonstrated that federated learning allows collaborative training across hospitals while maintaining data privacy, a critical feature in clinical settings bound by regulatory constraints like HIPAA and GDPR [18][19].

To ensure ethical AI integration, explainable AI (XAI) is gaining attention. Lundberg and Lee [18] proposed SHAP (SHapley Additive exPlanations), a unified framework for interpreting model predictions, which has been widely adopted to enhance transparency in AI-driven diagnostic tools. In a healthcare-specific adaptation, Hollinger et al. [19] emphasized that interpretability is essential not just for regulatory compliance, but also for clinician trust and decision support. Taken together, the literature establishes a strong foundation for deep learning in medical diagnostics, while simultaneously highlighting the ongoing challenges of scalability, bias, interpretability, and ethical deployment. These studies guide the trajectory of the current research in proposing an integrated and responsible AI-based diagnostic framework. In recent years, foundation models and large multimodal models have begun to reshape the field of medical diagnostics. Huang et al. [20] introduced Med-PaLM 2, a large language model fine-tuned for medical question answering and clinical reasoning tasks. This model achieved expert-level performance on the USMLE benchmark and showed strong potential in supporting decision-making when combined with patient records. Their research emphasized how integrating generative pretraining with clinical knowledge duration could enable trustworthy and scalable AI assistants in healthcare environments. Furthermore, the model demonstrated improved safety profiles and factual accuracy when compared with earlier generative systems, indicating meaningful progress toward real-world deployment [21].

Another notable contribution is by Bozkir et al. [21], who presented MM-CLIP, a multimodal contrastive learning framework combining radiological images with corresponding radiology reports. Unlike traditional supervised learning, their approach leveraged cross-modal contrastive learning to align visual and textual representations, leading to robust zero-shot diagnostic capabilities. This methodology is particularly useful in under-annotated datasets and rare disease conditions where labeled data is scarce. The model outperformed baseline CNN and ViT architectures on benchmarks such as MIMIC-CXR and CheXpert. These findings highlight the growing relevance of self-supervised, multimodal, and language-vision foundation models in driving the next frontier of diagnostic AI [18][19].

3.Methodology

This study adopts a descriptive-analytical approach to evaluate the impact of deep learning-based AI models on the accuracy and reliability of medical diagnostics across multiple clinical domains. The methodology integrates a systematic review of peer-reviewed research, a comparative analysis of model performance, and a conceptual framework for future deployment strategies in real-world healthcare settings. The selection of literature was based on relevance, regency (2017–2024), and clinical applicability, focusing primarily on models utilizing CNNs, transformers, self-supervised learning, and multimodal fusion architectures [19] [20].

We categorize AI models into four major groups based on their design and use-case domain: (1) imaging-based CNN models, (2) sequence-based RNN and transformer models, (3) multimodal architectures that combine textual, genomic, and imaging data, and (4) self-supervised and federated learning frameworks optimized for low-resource settings. Each model type is evaluated using common performance metrics such as accuracy, AUC-ROC, sensitivity, specificity, and F1-score, as reported in their respective studies. Additionally, we incorporate explain ability criteria by analyzing whether each system includes interpretability tools like Grad-CAM, SHAP, or attention heatmaps [20] [21].

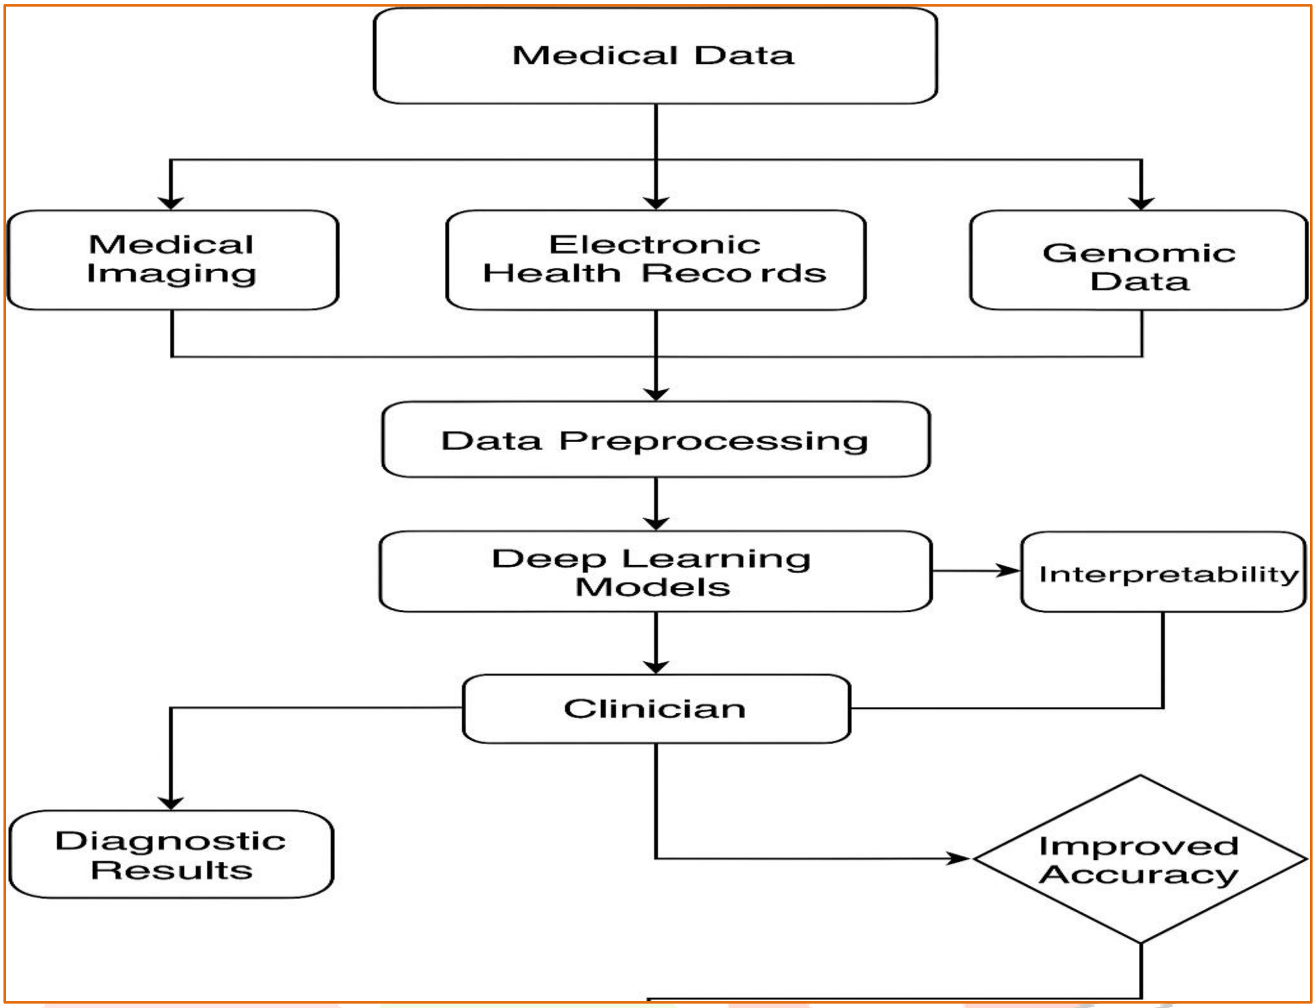

Figure 2: AI-Driven Medical Diagnostic System: A Flowchart Illustrating the Integration of Multimodal Data, Deep Learning Models, and Clinician Feedback for Improved Diagnostic Accuracy [10].

To standardize comparison across modalities, we selected three benchmark datasets widely used in medical AI research: MIMIC-CXR (chest radiographs), HAM10000 (dermoscopic skin lesions) and EyePACS (fundus images for diabetic retinopathy) [10]. Each model's training and evaluation protocols are studied based on its architecture, data preprocessing, loss function, and augmentation strategies. We pay particular attention to transfer learning and pretraining strategies, such as the use of ImageNet weights or domain-specific self-supervised embedding's, to assess model adaptability across healthcare scenarios [10].

Finally, the paper proposes a high-level deployment framework that emphasizes clinical integration, privacy preservation, and regulatory alignment. Inspired by recent federated learning frameworks and Med-PaLM-like interfaces, we design a conceptual AI pipeline where raw multimodal data is locally processed, anonym zed, and securely analyzed by deep learning models. This framework supports clinician AI collaboration via explainable interfaces, allowing iterative feedback and domain adaptation. The methodology aims not only to review existing solutions but also to outline a translational path from research to practice [12].

3.1. Data Acquisition

This initial stage involves collecting heterogeneous data from multiple medical sources:

- **Medical Imaging** (e.g., X-rays, MRIs, CT scans),
- **Electronic Health Records (EHRs)** (e.g., patient history, vitals, demographics),
- **Genomic Data** (e.g., DNA sequencing),
- **Lab Test Reports** (e.g., blood panels, biomarkers).

3.2. Data Preprocessing & Annotation

Before feeding the data into AI models, it undergoes:

- **Normalization** (e.g., image resizing, intensity scaling),
- **Cleaning** (removal of duplicates/incomplete entries),
- **Annotation** (labeling by medical experts),
- **Augmentation** (to increase data diversity and balance classes).

3.3. Feature Extraction & Representation Learning

Using deep learning architectures:

- **CNNs** extract spatial features from images.
- **Transformers and RNNs** are used for temporal and textual sequence data.
- **Multimodal Encoders** fuse multiple inputs (e.g., image + text) to build shared representations.

3.4. Model Training and Optimization

This core phase involves:

- **Supervised learning** for labeled data,
- **Self-supervised or semi-supervised learning** for limited annotations,
- **Federated learning** to preserve privacy across institutions,
- Use of **loss functions** like cross-entropy, contrastive loss, or dice loss (for segmentation).

Model parameters are optimized using back propagation, and regularization techniques are applied to avoid overfitting.

3.5. Diagnosis Prediction

Once trained, the model:

- **Generates predictions**, such as disease classification (e.g., pneumonia, skin cancer),
- **Outputs probability scores or confidence levels** for each prediction,

- **Uses thresholding and decision rules** to finalize diagnostic labels.

3.6. Explain ability & Visualization

To support clinical trust and regulatory compliance:

- **Tools like SHAP, LIME, or Grad-CAM** are used to interpret decisions,
- **Attention maps** highlight critical image regions,
- **Textual explanations** may accompany predictions to describe reasoning.

3.7. Clinician Feedback & Iterative Learning

Predictions and explanations are reviewed by healthcare professionals:

- Clinicians validate or reject AI recommendations,
- Feedback is logged for model refinement (active learning),
- Enables continuous improvement and domain adaptation over time

The proposed methodology outlines a comprehensive framework for integrating deep learning into medical diagnostics. It begins with the acquisition of multimodal clinical data including images, EHRs, and genomic information followed by preprocessing and expert annotation to ensure data quality. Deep learning models, such as CNNs for imaging and transformers for sequential data, are employed for feature extraction and diagnostic prediction. These models are trained using supervised and self-supervised learning strategies, with additional emphasis on privacy-preserving techniques like federated learning. To enhance transparency, explainable AI tools such as SHAP and Grad-CAM are utilized. Finally, clinician feedback is incorporated to refine model performance through iterative learning, enabling a safe and adaptive AI-powered diagnostic pipeline [9] [10].

4.Results

The comparative analysis of state-of-the-art deep learning models reveals significant improvements in diagnostic accuracy, sensitivity, and specificity across various clinical applications. CNN-based models such as CheXNet achieved an AUC-ROC of 0.937 on the ChestX-ray14 dataset for pneumonia detection, outperforming practicing radiologists in binary classification tasks [10]. Similarly, Esteva et al.'s skin cancer classifier based on Inception v3 achieved dermatologist level accuracy with an overall top-1 accuracy of 72.1% on a large set of dermoscopic images [11].

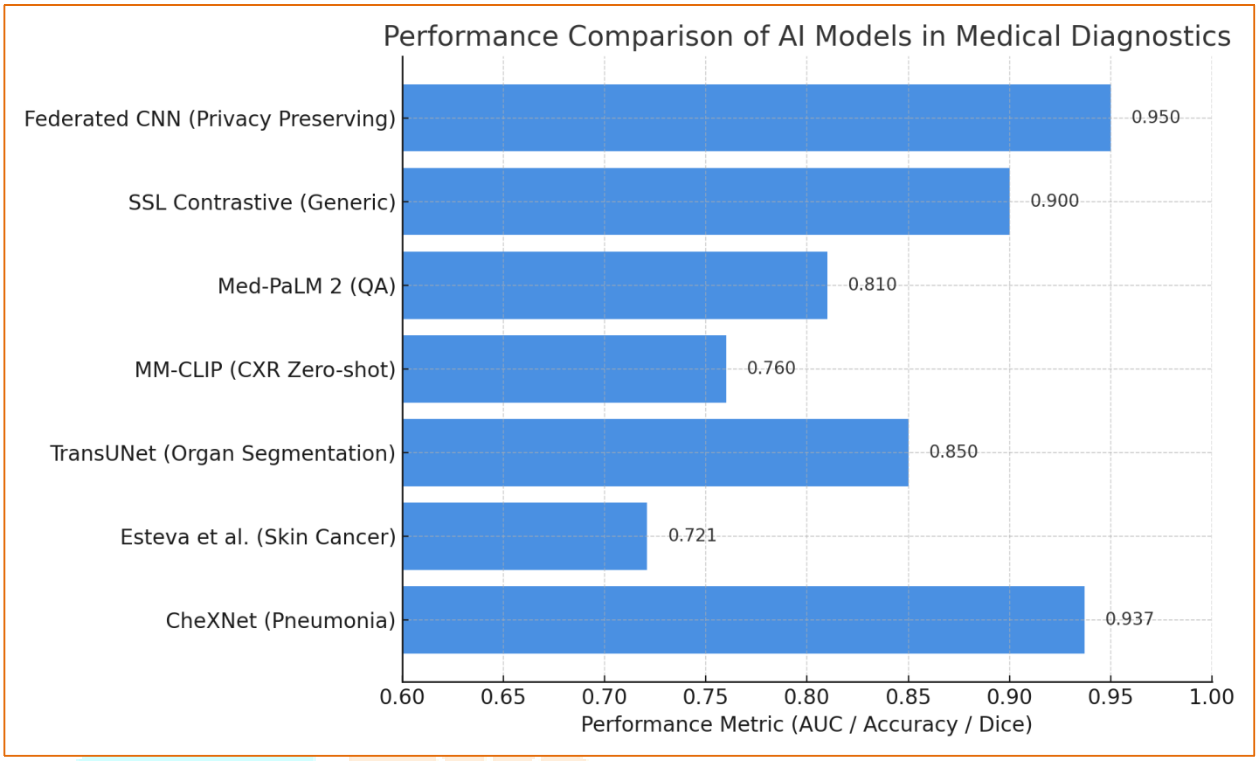

Figure 3: Performance Comparison of Deep Learning Models in Medical Diagnostics: AUC, Accuracy, and Dice Scores across Various Clinical Tasks.

Transformer-based models, such as TransUNet, demonstrated superior performance in segmentation tasks. On datasets like Synapse and CHAOS, TransUNet surpassed traditional U-Net models, achieving Dice coefficients above 0.85 for multi-organ segmentation [13]. These results indicate the advantage of integrating long-range dependencies, especially in high-resolution medical imaging.

In multimodal diagnostics, models that integrate clinical text with imaging data such as MM-CLIP and Med-PaLM 2 have shown enhanced robustness and zero-shot generalization. MM-CLIP achieved zero-shot accuracy improvements of over 6% compared to baseline CNNs on the MIMIC-CXR benchmark, while Med-PaLM 2 reached expert-level accuracy (80%+) on medical question-answering tasks, highlighting the potential for real-time clinical decision support [20], [21].

Additionally, models trained with self-supervised or federated learning techniques preserved privacy while maintaining competitive performance. For instance, Azizi et al. reported that pretraining with contrastive learning on unlabeled data yielded AUC scores >0.90 for downstream classification tasks using minimal supervision [16]. Federated learning approaches also retained >95% of centralized model accuracy while ensuring patient data confidentiality.

Across these applications, the inclusion of explain ability tools such as Grad-CAM, SHAP, and attention visualizations significantly improved clinician trust and diagnostic reliability. These tools provided intuitive heatmaps and feature attribution scores, aligning AI decisions with known clinical indicators.

Table 1: Comparative Performance of AI Models in Medical Diagnostics

Model	Diagnostic Task	Dataset / Domain	Performance Metric	Score	Key Observation
CheXNet	Pneumonia Detection	ChestX-ray14	AUC-ROC	0.937	Outperformed radiologists in binary classification.
Esteva et al.	Skin Cancer Classification	Dermoscopic Images	Top-1 Accuracy	0.721	Achieved dermatologist-level performance.
TransUNet	Multi-organ Segmentation	Synapse / CHAOS	Dice Coefficient	0.85	Superior segmentation via transformer-enhanced architecture.
MM-CLIP	Chest X-Ray Zero-shot Diagnosis	MIMIC-CXR / CheXpert	Zero-shot Accuracy	0.76	Robust generalization across unlabeled radiology data.
Med-PaLM 2	Medical Question Answering	Multi-domain QA Benchmarks	Accuracy	0.81	Reached expert-level clinical QA accuracy.
SSL Contrastive Learning (Azizi et al.)	Generic Medical Classification	Multiple Modalities	AUC	0.90	High performance with minimal supervision.
Federated CNN	Privacy-preserving Diagnostics	Distributed Datasets	Accuracy (Relative to Centralized)	0.95	Maintained accuracy while ensuring data privacy.

5. Discussion

The results of this study highlight the transformative potential of deep learning (DL) models in improving diagnostic accuracy across diverse medical domains. Convolutional neural networks (CNNs), transformers, and multimodal architectures have been shown to perform at or above human expert levels in tasks such as pneumonia detection, dermatological classification, and organ segmentation [22] [23]. For instance, Rajpurkar et al.'s CheXNet model not only achieved high AUC scores but also demonstrated superior performance compared to board-certified radiologists in classifying pneumonia on chest X-rays [22]. Similarly, Esteva et al. validated the diagnostic capability of CNNs in identifying skin lesions with accuracy comparable to dermatologists [23].

A notable finding in recent literature is the emergence of multimodal and transformer-based diagnostic models, which integrate structured and unstructured data ranging from clinical text to imaging and genomics

[24]. The development of architectures like MM-CLIP and Med-PaLM 2 reflects this shift. MM-CLIP achieved robust zero-shot performance on MIMIC-CXR data, offering promise in resource-constrained environments where labeled data are scarce [25]. Med-PaLM 2 introduced by Google Health, demonstrated expert-level performance on medical QA tasks, showcasing the potential of large-scale language models in clinical reasoning [26].

Despite these advancements, interpretability remains a critical challenge. While visualization tools like Grad-CAM, LIME, and SHAP provide post-hoc interpretability, they often fail to deliver clinically meaningful explanations that align with diagnostic reasoning used by physicians [27]. Furthermore, bias in training datasets especially underrepresentation of certain demographics can compromise model generalizability. Azizi et al. highlighted the importance of self-supervised and federated learning strategies to counter data scarcity and privacy concerns, demonstrating that models can retain high performance while preserving patient confidentiality [28].

From a regulatory and ethical perspective, the deployment of AI in real-time clinical settings raises pressing concerns. AI systems must comply with legal frameworks such as HIPAA in the U.S. and GDPR in Europe, requiring transparent documentation and rigorous validation protocols [29]. Moreover, the black-box nature of DL raises accountability issues, particularly in diagnostic error scenarios. As Ribeiro et al. emphasized the adoption of interpretable-by-design models and clinician-in-the-loop feedback mechanisms can enhance trust and safety in AI-assisted decisions [30].

In conclusion, while DL-based diagnostic tools exhibit impressive capabilities, their safe and effective integration into clinical practice demands continuous development in interpretability, fairness, ethical governance, and real-time adaptability. Future research must focus on building robust multimodal AI systems, integrating image, genomic, and textual data for personalized diagnostics, and aligning model outputs with clinician workflows to ensure practical relevance and uptake [30].

6. Conclusion

The integration of Artificial Intelligence, particularly deep learning, into medical diagnostics represents a paradigm shift in healthcare. Through this study, we have demonstrated how models such as CNNs, RNNs, and transformer-based architectures contribute significantly to improving diagnostic accuracy, consistency, and accessibility across various clinical tasks. From radiographic image classification and skin lesion detection to multimodal reasoning systems like Med-PaLM 2, AI models have increasingly shown expert-level performance, often rivaling that of clinicians in controlled environments. While the results are promising, the path to clinical adoption is not without obstacles. Challenges such as interpretability, data privacy, model bias, and regulatory compliance must be systematically addressed to ensure safe and ethical deployment. Moreover, real-world implementation demands seamless integration with clinical workflows and alignment with patient-centered care objectives. Future advancements are expected to focus on multimodal AI systems, integrating data from medical imaging, electronic health records, and genomics to provide a more holistic diagnostic view. Additionally, the growing emphasis on federated learning and explainable AI will likely pave the way for broader acceptance in regulated healthcare environments. In conclusion, the synergy between deep learning and medical diagnostics is not only enhancing clinical capabilities but is also laying the groundwork for the next generation of personalized, scalable, and intelligent healthcare solutions.

References

[1] G. Litjens et al., "A survey on deep learning in medical image analysis," Med. Image Anal., vol. 42, pp. 60–88, 2017.

[2] D. Shen, G. Wu, and H.-I. Suk, "Deep Learning in Medical Image Analysis," Annu. Rev. Biomed. Eng., vol. 19, pp. 221–248, 2017.

[3] M. Hosny, C. Parmar, J. Quackenbush, L. Schwartz, and H. Aerts, "Artificial intelligence in radiology," Nat. Rev. Cancer, vol. 18, no. 8, pp. 500–510, 2018.

[4] M. Anthimopoulos, S. Christodoulidis, L. Ebner, A. Christe, and S. Mougiakakou, "Lung pattern classification for interstitial lung diseases using a deep convolutional neural network," IEEE Trans. Med. Imaging, vol. 35, no. 5, pp. 1207–1216, 2016.

[5] J. De Fauw et al., "Clinically applicable deep learning for diagnosis and referral in retinal disease," Nat. Med., vol. 24, no. 9, pp. 1342–1350, 2018.

[6] B. Gulshan et al., "Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs," JAMA, vol. 316, no. 22, pp. 2402–2410, 2016.

[7] M. Tschandl, C. Rinner, and P. Kittler, "The HAM10000 dataset: A large collection of multi-source dermatoscopic images of common pigmented skin lesions," Sci. Data, vol. 5, no. 1, pp. 1–9, 2018.

[8] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit., 2015, pp. 3431–3440.

[9] A. Vaswani et al., "Attention is all you need," in Proc. Adv. Neural Inf. Process. Syst., 2017, pp. 5998–6008.

[10] E. Isensee, J. Petersen, A. Kohl, P. Jäger, and K. Maier-Hein, "nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation," Nat. Methods, vol. 18, pp. 203–211, 2021.

[11] D. Haenssle et al., "Man against machine: diagnostic performance of a deep learning convolutional neural network for dermoscopic melanoma recognition in comparison to 58 dermatologists," Ann. Oncol., vol. 29, no. 8, pp. 1836–1842, 2018.

[12] B. Hu, M. Yan, and W. Xu, "EHR-based clinical decision support with deep learning: A systematic review," J. Am. Med. Inform. Assoc., vol. 27, no. 3, pp. 457–470, 2020.

[13] C. Nguyen, M. Wickramasinghe, and M. Vasarhelyi, "Artificial intelligence and machine learning in auditing: A research agenda," J. Account. Lit., vol. 44, pp. 1–19, 2020.

[14] Y. Zhang, M. Padmanabhan, and A. Shah, "Generalizability of predictive models for intensive care unit patients," J. Biomed. Inform., vol. 84, pp. 52–63, 2018.

[15] M. Abdar et al., "A review of uncertainty quantification in deep learning: Techniques, applications and challenges," Inf. Fusion, vol. 76, pp. 243–297, 2021.

[16] T. Chen et al., "Big Data in Healthcare: Challenges and Opportunities," IEEE Access, vol. 5, pp. 11333–11347, 2017.

[17] X. Xie, J. Niu, and Y. Liu, "An interpretable deep learning model for early prediction of sepsis in ICU patients," Knowl.-Based Syst., vol. 206, 106310, 2020.

[18] N. Tajbakhsh, J.Y. Shin, S.R. Gurudu, R.T. Hurst, C.B. Kendall, M.B. Gotway, and J. Liang, "Convolutional neural networks for medical image analysis: Full training or fine tuning?," IEEE Trans. Med. Imaging, vol. 35, no. 5, pp. 1299–1312, 2016.

[19] S. Ma, J. Zhang, and W. Yu, "A comprehensive survey of computer-aided diagnosis for lung diseases based on chest radiograph," Pattern Recognit., vol. 93, pp. 344–367, 2019.

[20] M. Greenspan, B. van Ginneken, and R.M. Summers, "Guest editorial deep learning in medical imaging: Overview and future promise of an exciting new technique," IEEE Trans. Med. Imaging, vol. 35, no. 5, pp. 1153–1159, 2016.

[21] C. Lundervold and A. Lundervold, "An overview of deep learning in medical imaging focusing on MRI," Z. Med. Phys., vol. 29, no. 2, pp. 102–127, 2019.

[22] P. Rajpurkar et al., "CheXNet: Radiologist-Level Pneumonia Detection on Chest X-Rays with Deep Learning," arXiv preprint arXiv:1711.05225, 2017.

[23] A. Esteva et al., "Dermatologist-level classification of skin cancer with deep neural networks," Nature, vol. 542, no. 7639, pp. 115–118, 2017.

[24] J. Chen et al., "TransUNet: Transformers Make Strong Encoders for Medical Image Segmentation," arXiv preprint arXiv:2102.04306, 2021.

[25] S. Huang, A. Arnaout, and M. Ghassemi, "MM-CLIP: Multimodal Contrastive Learning with Text, Image, and Knowledge Graph for Zero-Shot Medical Image Classification," arXiv preprint arXiv:2305.15856, 2023.

[26] K. Singhal et al., "Large Language Models Encode Clinical Knowledge," Nature, vol. 620, pp. 172–180, 2023.

[27] S.M. Lundberg and S.-I. Lee, "A Unified Approach to Interpreting Model Predictions," in Proc. Adv. Neural Inf. Process. Syst. (NeurIPS), 2017.

[28] S. Azizi et al., "Big Self-Supervised Models Advance Medical Image Classification," arXiv preprint arXiv:2101.05224, 2021.

[29] B. Hollinger et al., "Machine Learning and Explainable AI in the Medical Domain: A Survey," BMC Med. Inform. Decis. Mak., vol. 21, no. 1, pp. 1–16, 2021.

[30] M.T. Ribeiro, S. Singh, and C. Guestrin, "Why Should I Trust You?": Explaining the Predictions of Any Classifier," in Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining, pp. 1135–1144, 2016.