



AI-Augmented Clinical Diagnostics: Integrating Deep Learning with Electronic Health Records and Imaging Data

Ravi Kumar Amaresam

Minisoft Technologies, USA

ABSTRACT: AI-enhanced clinical diagnostics will make the healthcare revolution by integrating deep learning with the rich contextual information stored in electronic health records (EHRs) and high-resolution imaging data. Fusion Multimodal fusions of pixel-level results of radiology or pathology images with structured and unstructured EHR data have been demonstrated to have higher diagnostics accuracy, risk stratification, and disease subtyping compared to single-modality models. The current literature has shown that combining convolutional neural networks (CNNs) with imaging and deep models of EHR, like recurrent or transformer, provides clinically relevant predictions of diseases like Alzheimer, breast cancer, and cardiometabolic diseases. Nevertheless, their implementation in practice in clinical settings has not yet been widespread since there is difficulty in terms of data quality, interoperability, algorithmic bias, interpretability, and EHR integration.

This paper suggests a single AI-enhanced clinical diagnostics system that incorporates deep learning pipelines of imaging and EHR data into an end-to-end and clinically consistent architecture. These frameworks include (1) ingestion of data, preprocessing, representation learning, multi-modal fusion (early, late, and hybrid), and explainability layers (inherently embedded in a clinical decision support and reporting interface). It is created to provide potential support to validation, transparent model behavior, and EHR seam, to enable clinician trust and adoption. We also describe a strategy of evaluation that integrates the conventional machine learning measures with the workflow and user-focused measures. The framework suggested will be beneficial in helping the researchers and healthcare organizations operationalize AI-enhanced, multimodal diagnostic systems that should advance beyond proof-of-concept models and be reliable, equitable, and scalable clinical applications.

KEYWORDS: Deep learning; Electronic health records; Medical imaging; Multimodal fusion; Clinical decision support; AI-augmented diagnostics

I. INTRODUCTION

Clinical diagnostics is conditionally becoming more dependent on large amounts of heterogeneous data, including high-resolution medical images all the way to granular longitudinal data of patient interactions, laboratory findings, prescribed medications and free-text clinical encounters. EHRs have become a hub of such information whereas developments in imaging technologies, such as digital radiology, computed tomography, magnetic resonance imaging, and digital pathology, have increased visual information available to clinicians by a very significant margin. Simultaneously, pattern recognition in imaging and tabular or textual clinical data has become a potent branch of deep learning, which allows detection of subtle patterns of diseases that a human mind can be blind to [1] [2].

So far, numerous successful medical AI applications are still limited by being designed to operate in a single modality, usually to address particular imaging problems as lung nodule detection, diabetic retinopathy grade, or breast cancer screening. These convolutional neural network based models are able to achieve performance on par or better than expert clinicians on limited benchmarks. Nevertheless, images are seldom used by human diagnosticians; instead, they use the results of the imaging as an interpretive process in relation to the age of a patient, comorbidity, laboratory values, previous diagnoses, treatment, and social determinants of health. This abundant information is coded in EHR systems and can significantly change the probability of disease in the pre-test phase and subsequent clinical decision-making. The growing sophistication of this paradigm of fusion has encouraged studies involving AI systems that reason collaboratively on imaging and EHR data to generate more accurate, robust and clinical meaningful diagnostic results [3].



According to the recent systematic and scoping reviews, there is an increasing number of studies where imaging information is intertwined with structured and unstructured EHR elements to be used in such tasks as diagnosis, prognosis, and prediction of treatment response. In these works, the multimodal models always perform better than the single-modality baselines, and frequently multimodal models exploit the complementary information provided by both the modalities. To give an example, images of radiology or pathology result in precise morphological or functional patterns, whereas EHR data add demographic, clinical, and temporal contextualizing the borderline results and allowing risk mitigation. Image-fusion models with clinical data have better discrimination than image-only models and EHR-only models in neurological disorders, breast cancer, and other diseases, and yield better calibration. Such returns highlight the potential of AI-enhanced clinical diagnostics: machines that are not only able to recognize images but also to carry out clinical reasoning in a way that even suggests holistic human reasoning [4] [5].

Regardless of this promise, multimodal AI translation into everyday clinical practice is not even. There are multiple challenges that have been reported again and again both in reviews and implementation-oriented research. To start with, the interoperability and data quality concerns with the EHRs such as missingness, variations in coding, and heterogeneous schemes across institutions hamper the reliability and generalizability of learned representations. Second, lots of models are trained on retrospective and single-center data without effective external validation, which is worrisome because they may be subject to overfitting, bias and drift when applied in new locations or groups of people. Third, when algorithmic performance is good, workflow integration and usability issues, including the absence of seamless integration into current EHR interfaces and inadequate clinician participation in the design, prevent adoption. Lastly, deep multimodal architectures are often opaque such that the insight into the way of generating predictions is very limited, losing clinician trust and facing trouble with regulatory approval [6].

To close these gaps, it is necessary to go beyond the isolated models development to full, system-wide models that explicitly view the entire lifecycle of AI-enhanced diagnostics. These structures should define the way EHR and imaging system data is ingested; the way data is preprocessed, harmonized and encoded; the way networks are designed on a modality basis and the way multimodal fusion is accomplished and assessed. They should also specify how prediction and explanations are provided to clinicians in time-constrained work methods, how the feedback is obtained, and how the models are observed and fashioned as maladaptive data distributions and clinical guidelines alter. Notably, they ought to incorporate the tenets of transparency, fairness, and safety, which complies with new requirements of reliable AI in healthcare.

In this respect, the term of AI-augmented clinical diagnostics focuses on the augmentation and not substitution of clinician knowledge. In this perspective, deep learning models are formulated as decision support systems that should be used to complement clinical judgment and provide probabilistic ratings, prioritized differentials, and explanations at the case-level that may be questioned and overruled by clinicians. This paradigm could be reinforced by multimodal integration which enables AI systems to expose cross-modal configurations, e.g., discordances between the results of the imaging and clinical courses, as such mistakes in diagnosis, data spoling, or unusual disease manifestations. Such models embedded into EHR-integrated clinical decision support systems could be used to aid in triage, highlight high-risk patients, and standardize patient interpretation across providers and, ultimately, help to minimize errors in diagnostics and unnecessary variations in care.

The present research article is focused on AI-enhanced clinical diagnostics, which combine deep learning and EHR data, as well as imaging. It contains three major objectives. First, it will synthesize and organize the literature to multimodal fusion of EHR and imaging data in clinical diagnostics in 2016 to 2023, with a specific emphasis on deep learning-based methods, fusion techniques, and clinical applications. Second, it suggests a general but instantiated architecture that encompasses the end to end pipeline of data ingestion to model deployment and monitoring of AI-enhanced diagnostics in health systems with heterogeneous EHR and imaging infrastructure. Third, it provides a description of an evaluation plan, which integrates algorithmic performance measures with user-friendly and workflow-based measures, and presents opportunities and ongoing challenges of the research and practice.

This article will take a conceptual form by outlining a systematic framework based on recent empirical and review findings to give researchers, clinicians, and health IT stakeholders a conceptual blueprint to design, implement, and evaluate AI-augmented clinical diagnostic systems. Instead of proposing a disease-specific model, it dwells on reusable architectural elements and design principles that may be adapted to various clinical situations. That way, it aims to dismantle the divide between the realities of the promising multimodal AI prototypes and the sustainable, and reliable deployment in an actual healthcare setting.



II. RELATED WORK

The fast development of artificial intelligence (AI) in the medical field has been reported in interdisciplinary studies. Preliminary general surveys laid the conceptual and technical groundwork of AI in medicine, whereas newer research has focused on practical use and domain-specific uses and interoperability issues. This chapter summarizes previous research in the area of AI foundations, clinical applications, disease-specific applications, predictive modeling, and health data integration.

Jiang et al. [1] give a background overview of AI in healthcare and its historical development of being a rule-based expert system to contemporary machine learning (ML) and deep learning methods. Their overview leads to the transformative nature of AI in the fields of diagnostics, treatment planning, and precision medicine, as well as outlines available challenges that include the heterogeneity of data, ethical issues, and regulatory limitations. These authors place AI in the context of healthcare delivery as not just a technological revolution but a revolutionary change within the system.

Continuing this general direction, He et al. [2] discuss the development of AI technologies in clinical medicine in practice. They determine several essential conditions of successful implementation, such as annotated datasets of high quality, strict validation, compliance with regulations, and their inclusion in clinical processes. Notably, the study highlights the discrepancy between the work of algorithms in controlled research contexts and their practical implementation in the clinical environment, highlighting the problem of generalizability, interpretability, and trust in the work of the clinician.

Davenport and Kalakota [3] also divide AI applications in the health sector into three areas: support of administrative workflow, clinical decision support, and patient engagement. Their article underlines the fact that AI has proven to be highly diagnostic in particular roles, but with a wider scope its value resides in supplementing medical practitioners and not replacing them. This augmentation model corresponds with the existing views that propose the collaboration of humans and AI to enhance their efficiency and quality of care.

Man-disease interactions receive a lot of attention. Oncology Bi et al. [4] are a literature review on AI application in cancer imaging, specifically on radiomics and deep learning-based image analysis. They show that AI methods can identify high-dimensional quantitative characteristics of imaging data, which is useful in the detection, classification and prediction of tumor response to treatment. Nevertheless, they also address clinical issues like reproducibility, standardization of data and necessity of prospective validation research.

On the same note, Johnson et al. [5] discuss the application of AI in the field of cardiology, where predictive modeling and risk stratification are possible through large-scale structured and unstructured data, including electrocardiograms, imaging, and electronic health records (EHRs) in cardiology. Their review addresses the topic of ML models to detect arrhythmia, predict heart failure, and interpret the image, stating the necessity to combine multimodal data sources and improve cardiovascular care.

Screening devices that are AI-based in ophthalmology have also been demonstrated as promising. Rajalakshmi et al. [6] prove that automated detection of diabetic retinopathy is possible with the help of AI algorithms on smartphone-based fundus photography. Using their work as an example, AI-powered mobile health solutions can open up access to screening in the resource-constrained environment. An example of how digital innovation can ensure access to healthcare by bridging the gap is the incorporation of AI and portable imaging technologies.

In addition to diagnostic imaging, machine learning-based predictive analytics has developed clinical risk assessment. Ye et al. [7] also use ML methods that are integrated with clinical notes to forecast mortality among critically ill diabetic patients. Their article emphasizes the importance of natural language processing (NLP) to extract meaningful clinical text characteristics of an unstructured clinical text. The study shows better predictive performance than the conventional statistical models through the incorporation of structured and unstructured EHR data.

Health data is vital to the functioning of AI systems because of its accessibility and quality. The article by Casey et al. [8] offers a literature review on the utilisation of electronic health records in population health research, with a description of methodological methods of managing large-scale clinical data. They write about the strengths, like the longitudinal coverage and real-world representation whereas the weaknesses, such as the missing data, differences in

codification and selection bias, are taken into consideration. Their results highlight the need in AI-based studies to have solid data processing and validation.

Besides the clinician generated data, patient generated health data (PGHD) has become a useful source of real-time health information. Cohen et al. [9] review the experience of Project HealthDesign that examined the implementation of PGHD into clinical care. Their results indicate that the challenges that are technical and workflow-related are the amount of data, its lack of standardization, and the issue of clinical liability. Nevertheless, with an appropriate organization into clinical systems, the opportunities of personalized and preventive care are provided by PGHD.

Standard interoperability is an important element that can facilitate the smooth sharing of data required in AI-based healthcare solutions. The article by Mandel et al. [10] proposes the SMART on FHIR which is a standards-based platform enabling interoperable applications to interface with EHR systems. SMART enables modular applications development and improves innovation in health IT ecosystems by using the Fast Healthcare Interoperability Resources (FHIR) standard. This platform offers an infrastructure that is scalable to provide AI applications into clinical workflows.

Braunstein [11] further explains how FHIR has the potential of succeeding in the implementation of healthcare interoperability. The author suggests that the key to the full potential of the digital health technologies, including AI, lies in the standardized data exchange mechanisms. Better interoperability increases data liquidity, which facilitates a cross-institutional study, big data analytics, and coordinated care delivery.

In line with such initiatives, Plastiras and O’Sullivan [12] suggest standard information model in the sharing of personal health data and observations of the daily living with the EHR systems. The work is concerned with semantic consistency and structural consistency of the relationship between PGHD and institutional records whereby data provided by patients can exert some form of meaning to the clinical decision-making process.

III. PROPOSED FRAMEWORK

The AI-enhanced clinical diagnostics framework suggested is a modular and end-to-end architecture aimed to implement deep learning models trained on EHR data and imaging data into standard clinical practices. It includes six fundamental layers, data acquisition and integration, data preprocessing and harmonization, modality-specific representation learning, multimodal fusion, explainability and uncertainty quantification, and clinical deployment and monitoring.




Figure 1. AI-Augmented Clinical Diagnostics Architecture Integrating EHR and Imaging Data



3.1 Data acquisition and integration

In the base, the framework presupposes the connection to imaging repositories, e.g., picture archiving and communication systems, and EHR platforms which contain structured and unstructured clinical data. A level of data integration executes safe extract transform load pipelines that compromise imaging metadata, pixel data, demographic characteristics, procedure and diagnosis codes, laboratory findings, medication orders and narrative notes on a recurring or on request basis. This layer imposes privacy and access control policies, such that data utilized in model training and inference do not violate institutional and regulatory limits. Where feasible, common terminologies and formats (e.g. DICOM to imaging data, HL7 FHIR to clinical data) are used in order to support interoperability and post-harmonization further.

3.2 Data preprocessing and harmonization

In the second layer, preprocessing of each modality is done, and the heterogeneous data are reconciled into model-ready forms. In the case of imaging data, common processes are DICOM parsing, de-identification, intensity standardization, resolution standardization and where necessary region-of-interest (where suitable) extraction or region-of-interest segmentation mask. The enhancement of data augmentation plans, including rotations, flips, and contrast adjustments is actively used under selection to enhance the robustness without violating the clinical plausibility. In the case of EHR data, preprocessing entails work on local codes to standardised vocabularies, dealing with missingness by imputation or explicit missingness indicators, temporal correspondence of events, and feature sequences or tabular structuring. NLP pipelines which include tokenization, entity recognition and more recently, transformer-based embeddings are used to process free-text notes to extract subtle clinical ideas.

One of the sub-modules is harmonization whereby imaging and EHR data are harmonized at the patient and encounter level, so that images and clinical variables belong to the same diagnostic episode or period. This subject-level and time correspondence is essential in preventing the label leakage, and in accurately representing the information present to clinicians during the diagnosis stage. The result of this layer is a series of coordinated multimodal instances, each containing one or more pictures and a set of organized and unstructured EHR features, in addition to outcome labels or targets of supervised learning.

3.3 Modality-specific representation learning

The third tier is composed of independent deep learning pipelines which learn rich representations of any given modality. In the case of imaging, convolutional neural networks/ vision transformers are the main feature extractors that generate high-dimensional embeddings that represent the spatial and contextual patterns of the clinical task. Depending on the task, architectures could be two-dimensional or three-dimensional, single or multi-view and they can also use attention mechanisms to concentrate on salient parts. Learning on large-scale natural or medical data sets can be used to transfer and increase the speed of learning in a data-sparse context.

In the case of EHR data, numbers of architectures can be implemented as per the input structure. Transformer-based or recurrent neural networks are appropriate with longitudinal events where the temporal dependence is important in diagnosis, procedure, and labs. Autoencoders in the form of factorization machines or autoencoders can be educated when given high-dimensional, sparse feature space and when using CNNs or transformers can encode tokenized clinical note spaces. Mostly, several sub-networks can process various sub-modalities of EHR structured codes, numeric labs, and text, and their results can be concatenated or combined through attention to create a single EHR embedding.

The major design concept is that every modality-specific pipeline should be trained to generate both downstream task discriminative and cross-modal fusion compatible embeddings. Joint training with fusion layers, multi-task goals, or contrastive learning which aligns representations within one modality with another can be promoted to the same patient.

3.4 Multimodal fusion strategies

The fourth layer enforces multimodal fusion and may take early or late and hybrid approaches. At the early stage of fusion, both imaging and EHR pipeline embeddings are redirected or fused through learned transformations and these are then sent through downstream classification or regression heads. This method performs direct learning of cross-modal interactions on the model, however it is highly balanced to avoid overfitting particularly when one modality becomes predominant in dimensionality or quality. In contrast, late fusion takes modality specific predictions or risk scores and either use weighted averaging or stacking or learned gating schemes. Late fusion may be more resistant to the absence of modalities and be more modular, with imaging and EHR models being able to be revised separately.

Hybrid approaches combine the two levels of fusion, e.g. with early fusion of strongly coupled sub-modalities (e.g. imaging and key labs) and late fusion of the fused signal with independent risk scores or specialist models. Attention-based fusion mechanisms are able to gain dynamical weight of modalities and features according to patient specific context which may enhance performance and interpretability. These strategies differ according to the availability of data, the needs of clinical tasks and the constraints of the systems.

3.5 Explainability and uncertainty quantification

The framework also has an explainability and uncertainty quantification layer to facilitate the trust of clinicians and regulatory acceptance. To imaging parts, gradient-based saliency maps, class activation maps, or attention heatmaps indicate image parts that contributed the most to the prediction, and aid the clinician to evaluate image consistency with known pathology. In the case of EHR components, feature attribution techniques like SHAP or integrated gradients, can be used to measure the impact of a particular diagnosis, lab, or medication, whereas attention weights in sequence models point to the time or events that are influential.

Monte Carlo dropout or deep ensembles are uncertainty estimation methods that can give calibrated uncertainty information like confidence intervals or risk groupings, allowing conservative deployment with cases of low uncertainty being given further consideration. Aggregating the explainability output at the multimodal level provides clinicians with a clear account of the evidence-concordant or evidence-discordant results of imaging and labs, with an example of presenting clinicians with a coherent narrative of the result set (a successful outcome, an inconclusive finding, etc).

3.6 Clinical deployment, feedback, and monitoring

The last layer deals with entry into clinical practices. The EHR-integrated interfaces including dashboards, alerts, or inherent visualizations on radiology and pathology viewers surface model outputs including predictions, explanations, and confidence measures. Human factors principles govern interaction design, and it makes sure that information is displayed at the right moment and in easy to digest formats to busy clinicians.

It provides a two-way feedback: clinicians are allowed to give corrections, annotations, or overrides, which are recorded and are used to improve models either by retraining models on a regular basis or by active learning mechanisms. Constant monitoring of performance indicators, data drift and fairness indicators across subpopulation, the alerts are caused by inequalities in performance or biases. Governance forms determine positions and duties of model stewardship, documentation as well as adherence to regulatory and institutional provisions.

The framework offers a reusable template through the arrangement of the AI-enhanced diagnostic system into these modular layers to be instantiated to a variety of clinical applications without losing data handling, modeling, explainability, and deployment practices.




Figure 2. Clinical Workflow Integration of AI-Augmented Diagnostics within EHR Systems

IV. FRAMEWORK EVALUATION

Assessment of the suggested framework should be a multidimensional one that incorporates performance of the algorithm, clinical usefulness, integration of workflow, and safety. Conventional machine learning scores are necessary but not enough to identify actual change in healthcare contexts.




Figure 3. Multi-Level Evaluation and Monitoring Framework for AI-Augmented Diagnostics

4.1 Algorithmic performance and ablation

The essence of multimodal models is evaluated based on conventional measures relevant to the task, including area under the ROC curve, precision-recall curves, calibration plots, and decision-curve analysis of diagnostic classification. To measure the incremental value of fusion, the complete multimodal system, as compared to single-modality baselines (imaging-only and EHR-only) should be evaluated. Ablation studies are carried out to investigate the value of each modality, fusion strategy and architectural element, such as the disabling of attention mechanisms or the removal of certain EHR sub-modalities or simplification of imaging backbones whilst tracking changes in performance.

Robust validation protocols are very critical. Internal cross-validation can be used to determine the baseline performance, and temporal validation (training on previous data, testing on subsequent periods) can be used to determine the resilience to temporal drift. Generalizability and the detection of domain shift and bias are assessed by external validation on datasets in other institutions, scanners, or demographics of the patients. In cases where possible, federated or multi-center studies can be used to make evidence even stronger such that performance gains are not exclusive to one site.



4.2 Explainability, trust, and usability

Considering the intricacy of multimodal deep learning systems, they should be evaluated by explaining their features as well as enticing clinician confidence. By carrying out user studies on clinicians, it is possible to evaluate how understandable, actionable and consistent saliency maps, feature attributions and uncertainty estimates are to clinical reasoning. The quantitative measures can involve time to decision, perceived work load, and congruency in interpretations with different users when they are given explanations compared with when they are withheld.

Mismatches between model explanations and clinician expectations may be demonstrated in the process of cognitive walkthroughs, think-aloud sessions, which can inform visualization/interaction design refinement efforts. The metrics of trust calibration, which refers to whether clinicians weight model recommendations in the appropriate proportion to a score of confidence, and previous performance, are also significant so that overreliance or excessive skepticism does not occur. Qualitative feedback on the types of explanations and the incorporation into the tools used can be useful in making sure that the framework does not interfere with diagnostic reasoning but, instead, helps to support it.

4.3 Workflow and clinical impact

In addition to personal forecasting, the framework should be assessed by measuring its effect on clinical process and patient outcome. The changes that implementation studies can monitor include the turnaround time of a diagnostic test, the percentage of cases that are high-risk and properly prioritized, and the diagnostic discrepancy rate prior to and following a system implementation. To illustrate, time-motion can be used to study the impact of AI-enhanced triage on radiologist reading order, consultation patterns, or rework rates. Future studies or pragmatic trials can address the question of whether AI-assisted diagnostics can decrease unnecessary testing, decrease hospitalization, or evidence-based guideline adherence.

In cases where the outcome data exist, comparative effectiveness studies may also examine whether utilization of the multimodal system is related to better clinical outcomes e.g. mortality, complication or readmissions with confounder adjustments. In early phases of deployment, interim proxies can include surrogate measures which include diagnostic confidence scores, inter-rater agreement and frequency of second opinions, as long-term outcomes are established.

4.4 Safety, fairness, and governance

Lastly, the framework should be considered safe and just. Such systematic failure modes as misclassification in rare but critical circumstances are monitored in safety tests, as well as the possibility of automation bias: clinicians overly trust model outputs. Vulnerabilities could be discovered during stress testing under distribution shifts such as during outbreaks, technology upgrades or other changes and so on.

Fairness assessment implies subgroups stratification of the metrics of the performance to reveal inequalities by demographic and clinical categories. In case of large performance differences, the mitigation measures, including reweighting, subgroup-specific calibration, or directed data augmentation, might be necessary. The governance processes must be evaluated in terms of accountability clarity, model documentation transparency, and institutional and regulatory AI requirements in healthcare.

Through this holistic assessment approach, stakeholders would be able to leave behind the monolithic accuracy criteria to get a realistic view of the performance of AI-enhanced, multimodal diagnostic systems in practice, its impact on clinician behavior, and how it can be improved by iteration to enable the delivery of safe, equitable, and effective care.

V. FUTURE OPPORTUNITIES

There is ongoing development of deep learning and EHR being combined with imaging data to provide clinical diagnostics and numerous opportunities in this area arise based on the existing limitations and trends.

To start with, the progress of foundation models and large multimodal transformers makes it possible that one may have unified architectures, where images, structured data, and text can be fed into a single model and hopefully make fusion easier and enhance cross-modal reasoning. The need to have task-specific architectures and the rapid adaptation to novel diseases and modalities such as the case of small clinical corpora and imaging datasets could be reduced by pretraining such models on large-scale, de-identified clinical corpus and imaging datasets, and then fine-tuning them on task-specific data.




Figure 4. Interoperability Stack Linking EHR, Imaging Systems, and AI Services

Second, self-supervised and loosely supervised methods of learning provide opportunities to exploit the large quantities of unlabeled clinical information available in EHRs and imaging archives. Multi-modal contrastive learning, e.g., matching the representations of images with related reports or EHR snapshots, may produce more informative patient embeddings and minimise the need to use costly expert annotations. This is especially so with rare diseases, the pediatric population, and resource limited environment where the labeled data is limited.

Third, multimodal clinical AI explainability research is on its knees. The future work can test hierarchical and narrative forms of explanation where imaging and EHR evidence are incorporated into the coherent case-based narratives with respect to clinical patterns of reasoning. Simulation Interfaces There might be improved understanding and trust with interactive methods where clinicians can query model behaviour and simulate counter-factuals and cross-modal evidence. It will be necessary to combine the research in the field of human-computer interaction with the formal assessment of interpretability and trust calibration.

Fourth, the framework could truly be longitudinal, context-aware diagnostics by combining real-time and ongoing streams of data, including bedside monitoring, wearable sensors, and patient-generated health data with the framework. Imaging, EHR, and patient-generated data AI systems can potentially have the capacity to identify early warning signs of worsening, individualize the screening schedule, and facilitate proactive care organization. This will involve powerful streaming architectures, in device inferences, and privacy preserving analytics.



Last but not least, the development of AI-enhanced diagnostics will be determined by policy, regulatory, and educational factors. Standards and guidelines in the documentation, validation, and monitoring of multimodal models can be created to facilitate the safe deployment as well as interoperability across institutions. The introduction of AI, data literacy, and critical assessment of algorithmic-based clinician training programs will enable health care providers to take a significant role in the design, testing, and regulation of these systems. It will be essential that multidisciplinary teams of clinicians, data scientists, ethicists, and patients address the issue of ensuring that AI-enhanced diagnostics will promote equity, respect patient autonomy, and coexist with the values of the society.

VI. CONCLUSION

AI-enhanced clinical diagnostics combining deep learning with EHR and imaging data are an appealing prospect of improving the accuracy of diagnostics, its consistency, and contextualization. In 2016-2023 research has shown that in imaging and EHR, deep learning has performed well and multimodal fusion generally has better performance than single-modality models. Nonetheless, clinical scalability is limited by the quality of data, generalizability, workflow, and explainability issues. The paper has suggested a framework that takes the entire lifecycle of AI-enhanced diagnostics, including data collection, preprocessing, learning of modality-specific representations, multimodal fusion, elucidation, and clinical implementation and tracking. Through the description of clear architecture elements and assessment plans, the framework shall be used to steer the process of designing technically sound, clinically significant and operationally viable systems. To achieve the potential of AI-enhanced clinical diagnostics, it will be necessary to perform multidimensional assessment arduously, refining it based on a clinical response; to consider fairness, safety and governance. Future research and innovation can be fruitful in the domain of multimodal foundation models, self-supervised learning, interactive explainability, and incorporation of patient-generated data. AI systems integrating clinical intervention and imaging data in a cautious and responsible manner can become reliable collaborators in clinical decision making, aiding more accurate, timely and equitable care.

REFERENCES

1. Jiang F, Jiang Y, Zhi H, et al. Artificial intelligence in healthcare: past, present and future. *Stroke and Vascular Neurology*. 2017;2(4). doi: 10.1136/svn-2017-000101.
2. He J, Baxter SL, Xu J, Xu J, Zhou X, Zhang K. The practical implementation of artificial intelligence technologies in medicine. *Nature Medicine*. 2019;25(1):30–36. doi: 10.1038/s41591-018-0307-0.
3. Davenport T, Kalakota R. The potential for artificial intelligence in healthcare. *Future Healthcare Journal*. 2019;6(2):94. doi: 10.7861/futurehosp.6-2-94.
4. Bi WL, Hosny A, Schabath MB, et al. Artificial intelligence in cancer imaging: clinical challenges and applications. *CA: A Cancer Journal for Clinicians*. 2019;69(2):127–157. doi: 10.3322/caac.21552.
5. Johnson KW, Torres Soto J, Glicksberg BS, et al. Artificial intelligence in cardiology. *Journal of the American College of Cardiology*. 2018;71(23):2668–2679. doi: 10.1016/j.jacc.2018.03.521.
6. Rajalakshmi R, Subashini R, Anjana RM, Mohan V. Automated diabetic retinopathy detection in smartphone-based fundus photography using artificial intelligence. *Eye*. 2018;32(6):1138–1144. doi: 10.1038/s41433-018-0064-9.
7. Ye J, Yao L, Shen J, Janarthanam R, Luo Y. Predicting mortality in critically ill patients with diabetes using machine learning and clinical notes. *BMC Medical Informatics and Decision Making*. 2020;20(11):1–7. doi: 10.1186/s12911-020-01318-4.
8. Casey JA, Schwartz BS, Stewart WF, Adler NE. Using electronic health records for population health research: a review of methods and applications. *Annual Review of Public Health*. 2016;37. doi: 10.1146/annurev-publhealth-032315-021353.
9. Cohen DJ, Keller SR, Hayes GR, Dorr DA, Ash JS, Sittig DF. Integrating patient-generated health data into clinical care settings or clinical decision-making: lessons learned from Project HealthDesign. *JMIR Human Factors*. 2016;3(2):e5919. doi: 10.2196/humanfactors.5919.
10. Mandel JC, Kreda DA, Mandl KD, Kohane IS, Ramoni RB. SMART on FHIR: a standards-based, interoperable apps platform for electronic health records. *Journal of the American Medical Informatics Association*. 2016;23(5):899–908. doi: 10.1093/jamia/ocv189.
11. Braunstein ML. Healthcare in the age of interoperability: the promise of fast healthcare interoperability resources. *IEEE Pulse*. 2018;9(6):24–27. doi: 10.1109/MPUL.2018.2869317.
12. Plastiras P, O’Sullivan D. Exchanging personal health data with electronic health records: a standardized information model for patient generated health data and observations of daily living. *International Journal of Medical Informatics*. 2018;120:116–125. doi: 10.1016/j.ijmedinf.2018.10.006