



iJRASET

International Journal For Research in
Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** X **Month of publication:** October 2026

DOI: <https://doi.org/10.22214/ijraset.2026.85087>

www.ijraset.com

Call: ☎ 08813907089

E-mail ID: ijraset@gmail.com

Artificial Intelligence for Intelligent and Privacy-Preserving Healthcare: A Survey of Medical Image Analysis, Multimodal Diagnosis, Semi-Supervised Learning, and Personalized Healthcare Systems

Abhinav Krishna¹, Akshay G², Abhishek Sreeji³, Abiya F R⁴, Mrs. Sagara M R⁵

^{1,2,3,4}Dept. of Computer Science (Data Science), St. Thomas Institute for Science & Technology, Trivandrum, India

⁵Asst. Professor, Head of Dept. of Computer Science (Data Science), St. Thomas Institute for Science & Technology, Trivandrum, India

Abstract: Artificial intelligence (AI) is increasingly embedded in healthcare, from image-based diagnosis to patient-facing management tools, yet the supporting literature is fragmented across communities that rarely read one another. Motivated by the need for a coherent view, this survey examines four recent peer-reviewed studies that together span multimodal COVID-19 diagnosis, privacy-preserving learning on encrypted medical images, graph-refined semi-supervised pathology classification, and an AI-enabled mobile application for personalized healthcare management and doctor recommendation. Using a structured narrative methodology, the reviewed studies are classified into a five-category taxonomy, compared along technical dimensions, healthcare capabilities, and reported limitations, and synthesized into a layered view that treats imaging, multimodal fusion, privacy preservation, label-efficient learning, and personalization as complementary layers of intelligent healthcare rather than independent technologies.

The synthesis indicates that each study advances one layer under its own experimental setting, and that none addresses privacy, label efficiency, multimodality, explainability, and clinical validation together. Recurring gaps include limited clinical validation, dataset constraints, unquantified generalization across institutions, privacy-performance trade-offs, and inconsistent evaluation protocols that prevent direct numerical comparison. Future directions include privacy-preserving multimodal learning, reduced reliance on annotations, explainable and robust models, integration of imaging with clinical records, and multi-site clinical validation. This review is not exhaustive, and no new experiments are presented.

Index Terms: Artificial intelligence in healthcare, medical image analysis, multimodal fusion, privacy-preserving deep learning, learnable image encryption, semi-supervised learning, graph-based segmentation, personalized healthcare, clinical decision support, survey.

I. INTRODUCTION

A. Research Context

Artificial intelligence (AI) and machine learning (ML) are now applied across the healthcare pipeline, including image-based diagnosis, risk prediction, patient-facing management, and administrative support. Deep learning in particular has become the dominant methodological choice for analyzing medical and histology images [5]. The four studies examined in this survey illustrate how broad this landscape is: a multimodal COVID-19 detection model that fuses imaging and laboratory information [1], a learnable image-encryption scheme that allows deep networks to be trained on protected medical images [2], a semi-supervised framework that refines pseudo-labels for whole-slide pathology images through graph segmentation [3], and a mobile application that combines symptom-based specialist recommendation, personalized diet and exercise planning, and record management [4].

B. Motivation

Each of these studies is typically read within its own community: diagnostic imaging, security and privacy, computational pathology, or health informatics. Yet the problems they address are linked. Diagnosis from a single source can be ambiguous, which motivates fusion of complementary evidence [1]. Training complex networks often requires cloud resources that are treated as semi-honest, which raises privacy concerns for sensitive images [2].

Dense annotations are costly to obtain, particularly for gigapixel slides [3]. Patient-facing tools must handle fragmented care delivery and sensitive records [4]. Reading these strands together exposes dependencies that no single paper addresses.

C. Problem Statement

The literature on intelligent healthcare lacks a consolidated account of how multimodal, privacy-preserving, label-efficient, and personalized approaches relate to one another, what each has demonstrated under its own experimental conditions, and which shared limitations remain. This survey addresses that gap for the studies in scope.

D. Scope of the Survey

The survey covers AI and ML in healthcare with emphasis on (i) medical image analysis and deep learning, (ii) multimodal diagnosis, (iii) privacy-preserving analysis and learnable image encryption, (iv) semi-supervised learning with graph-based segmentation and pseudo-labeling for pathology, and (v) personalized healthcare systems, including symptom-based doctor and specialist recommendation, natural language processing (NLP), electronic health record (EHR) access, and AI-enabled applications. Explainability, reliability, and deployment challenges are examined across all five areas. Coverage of each topic is bounded by what the four seed studies and the literature they cite actually report; topics such as NLP and EHR management are represented by a single seed study.

E. Research Questions

RQ1. Which AI and ML approaches do the reviewed studies use for intelligent healthcare and medical diagnosis?

RQ2. How do multimodal, privacy-preserving, and semi-supervised approaches address limitations of conventional medical AI?

RQ3. How do the reviewed studies differ in data requirements, learning strategy, model architecture, privacy treatment, and evaluation?

RQ4. Which limitations and research gaps recur across the reviewed studies?

RQ5. Which research directions, reported by the original studies or derived from this synthesis, could improve reliability, privacy, generalization, explainability, and clinical applicability?

F. Contributions of the Survey

This survey (1) proposes a five-category taxonomy that organizes the reviewed studies; (2) provides a thematic rather than paper-by-paper synthesis of problem, method, data, results, limitations, and reported future scope; (3) compares the studies in four tables covering overview, technical design, capability coverage, and challenges; (4) identifies research gaps and links each to evidence in the reviewed literature; and (5) offers a conceptual, non-implemented view of how the layers of intelligent healthcare relate. The survey does not propose a new medical system and reports no new experiments; all quantitative values are attributed to their original studies.

G. Organization of the Paper

Section II describes the survey methodology. Section III summarizes foundational concepts. Section IV presents the taxonomy, and Section V synthesizes the literature thematically. Section VI gives comparative tables, Section VII discusses research gaps, and Section VIII provides a cross-study synthesis. Section IX outlines future directions, Section X states the limitations of this survey, and Section XI concludes.

II. SURVEY METHODOLOGY

This survey follows a structured narrative (scoping) design. It is not a systematic review, and no PRISMA-style identification, screening, or counting of database records was performed, because no independent database search was conducted. The workflow is shown in Fig. 3.




Fig. 3. Survey workflow (conceptual figure; not an experimental result).

A. Literature Sources

The primary sources are four full-text papers supplied by the author team, all published by IEEE (IEEE Access and the IEEE ICICCS conference). Secondary sources are works cited inside those papers; they are used for background and are described only as characterized by the citing paper, using the bibliographic details given there. A separate sample survey document was used only as a structural and stylistic reference and contributes no content.

B. Search Strategy

No keyword database search was run. The seed set was assembled by the team around medical AI, and secondary literature was reached by following the reference lists of the seed papers. Consequently, the coverage of this survey depends on the seed set and is not exhaustive.

C. Inclusion Criteria

A study was included if it (i) applies AI or ML to a healthcare task, (ii) is a peer-reviewed publication with available full text, (iii) describes its method and data, and (iv) fits at least one of the five taxonomy areas.

D. Exclusion Criteria

Sources were excluded or restricted to background use if they were outside healthcare, lacked verifiable bibliographic details in the citing paper, or could not be matched to the claim attributed to them. In one seed paper [1], several related-work citations could not be matched to relevant entries in its own reference list; the performance figures attributed to them there are therefore not reused in this survey.

E. Study Selection

All four seed papers met the inclusion criteria and are retained. They are labeled [1]–[4] throughout. Secondary references were selected only where they clarify a concept discussed by a seed paper.

F. Data Extraction

For each seed paper the following were extracted: healthcare problem, objective, methodology and models, data and modality, evaluation metrics, key results, stated limitations, and stated future scope. Items not reported are recorded as “not reported” rather than inferred.

G. Synthesis Method

Qualitative thematic synthesis was used. Studies were grouped by taxonomy category, compared along common dimensions, and examined for recurring limitations. Statements are tagged conceptually as (a) proposed by the original study, (b) evaluated by the original study, (c) observed by this survey from comparison, or (d) an inferred gap. No meta-analysis was performed.

III. BACKGROUND AND FOUNDATIONAL CONCEPTS

A. Artificial Intelligence in Healthcare

AI in healthcare spans clinical tasks such as image interpretation and decision support, and non-clinical tasks such as appointment management and navigation. The seed studies cover both ends: pathology and radiology-style classification [3], [1], [2] and a patient-facing application [4].

B. Machine Learning and Deep Learning in Healthcare

Classical ML models such as logistic regression and random forests are used for tabular and symptom data in [4], whereas convolutional neural networks (CNNs) dominate imaging work, including ResNet-50 in [3] and DenseNet-121 and XceptionNet in [2] [6], [7]. Deep models are data-hungry and computationally demanding, which creates the annotation and privacy pressures discussed below.

C. Medical Image Analysis

The reviewed studies use chest X-ray and CT images [1], MRI and chest X-ray images [2], and hematoxylin-and-eosin whole-slide images (WSIs) [3]. WSIs may exceed $100,000 \times 100,000$ pixels, so networks are generally trained on extracted patches, and region-level annotation requires pathologists to inspect tissue at several magnifications [3].

D. Multimodal Learning

Multimodal learning combines complementary evidence. In [1], features from chest X-ray, CT, and RT-PCR information are combined by concatenation or weighted averaging before classification. The same study cites multimodal fusion for kidney pathology as related work [8].

E. Privacy-Preserving AI

Cloud servers used for training are considered semi-honest in [2]: they follow the protocol but whoever can access them may inspect uploaded data. Regulations such as HIPAA make medical data sharing complex [2]. Learnable image encryption transforms images so that networks can be trained on the encrypted form without decryption, as in the Tanaka and SKK schemes [9], [10].

F. Semi-Supervised Learning

Semi-supervised learning (SSL) uses a small labeled set with a large unlabeled set [11]. Pseudo-labeling retrain a model on its own predictions for unlabeled data [12], but it depends on pseudo-label reliability and can overfit to wrong labels through confirmation bias [13]. Uncertainty-aware selection has been proposed as a remedy [14].

G. Graph-Based Learning

Graphs represent spatial and feature relationships among image regions. In [3], patches of a WSI form graph nodes connected to their eight neighbors and to tumor and normal terminal nodes, and a minimum cut is computed with the max-flow algorithm of classical graph-cut segmentation [15].

H. Personalized Healthcare and Clinical Decision Support

Personalized systems tailor outputs to an individual's symptoms, metrics, location, or history. [4] combines symptom-based specialist recommendation, calorie-based meal planning, prescription handling, and QR-gated record access. Related health-informatics systems cited there address disease prediction, chatbots, and diet and exercise advice [16], [17]. Such tools support, but do not replace, clinical decision-making.

IV. TAXONOMY OF AI-BASED HEALTHCARE APPROACHES

Fig. 1 organizes the reviewed studies into five categories. A study may touch several categories; it is placed in the category that best reflects its central contribution, with secondary roles noted in Table III.




Fig. 1. Taxonomy of AI-based healthcare approaches derived from the reviewed studies (conceptual figure).

1) Medical image-based deep learning matters because imaging is the main data source for diagnosis and because deep networks can learn discriminative features directly. All three imaging studies belong here at the task level [1], [2], [3]. 2) Multimodal medical diagnosis addresses the ambiguity of single-source evidence; it is represented by [1]. 3) Privacy-preserving medical AI addresses exposure of sensitive images to semi-honest servers; it is represented by [2]. 4) Semi-supervised and graph-based learning addresses annotation scarcity and structural context; it is represented by [3]. 5) Personalized and intelligent healthcare applications address access, navigation, and individualized guidance; it is represented by [4].

V. REVIEW AND SYNTHESIS OF EXISTING LITERATURE

A. Deep Learning for Medical Image Diagnosis

Deep learning for histopathology has been surveyed extensively [5]. [3] notes that earlier patch-based CNNs, including deep architectures such as ResNet and Inception variants, reached pathologist-level or better performance on specific tasks, but under fully supervised training with large annotated sets. [1] pursues COVID-19 detection from imaging with a CNN, and [2] treats brain-tumor MRI and COVID-19 chest X-ray classification as downstream tasks for studying privacy. Across these studies the models are standard CNNs; the novelty lies in what surrounds the classifier (fusion, encryption, or label refinement) rather than in a new backbone. Accuracy, F1, and AUC are used respectively, so the numerical results are not directly comparable.

B. Multimodal Medical Diagnosis

Problem and objective. COVID-19 is hard to diagnose from a single source because symptoms overlap with other respiratory diseases, RT-PCR can be slow, and imaging quality varies [1]. The objective is to improve detection and classification by fusing chest X-ray, CT, and RT-PCR information. **Methodology.** The study preprocesses and augments images, extracts features with a CNN, fuses them by concatenation or weighted averaging, and classifies with a softmax layer. **Data.** A Kaggle-sourced collection of 500 positive and 500 negative samples was split 80%/20% for training and testing. **Results.** The study reports 92.1% accuracy for the fused model, against 80.3% for chest X-ray, 85.6% for CT, and 89.2% for RT-PCR alone. **Contribution.** It demonstrates, within its setting, that fusing modalities outperformed each single modality.

Limitations. The authors state that the model has not been validated in a clinical setting, that comparison with other COVID-19 methods was limited, and that larger and more varied data are needed. This survey additionally observes several reporting inconsistencies: the fusion equations list only X-ray and CT features, while the text and abstract describe RT-PCR as an additional input; the result-analysis text describes a second model that extracts features from clinical variables such as age and gender; accuracy is the only metric reported although precision, recall, and F1 are named; the ablation is described without numerical

values; and the outcome histogram in the paper's figure appears unbalanced relative to the stated 500/500 composition. Because the claims in its correlation analysis are not tied to a validated outcome measure, they are treated as descriptive only. Its future scope is evaluation on larger, more diverse data and other populations and settings.

C. Privacy-Preserving Medical Image Analysis

Problem and objective. Training complex networks in the cloud exposes images to a semi-honest server, and earlier medical-image encryption methods require decryption before training [2]. The objective is a learnable encryption scheme that preserves training utility while resisting reconstruction attacks. **Background.** The Tanaka scheme and the SKK scheme with independent pixel-wise keys [9], [10] enable training on encrypted images, but attack algorithms, notably the leading-bit and minimum-difference attacks, can recover substantial image features from SKK-encrypted images [18]. GAN-based transformation removes key management [19], and DeepEDN encrypts and decrypts medical images but, as characterized in [2], is unsuitable when decryption must precede cloud training [20].

Methodology. The reviewed study adds block-based statistical smoothing (median, mean, maximum, or minimum filter, chosen per block by a secret key) after the SKK negative-positive transformation and color shuffling. The block size acts as a parameter that trades accuracy for security. **Data and models.** Combined open MRI brain-tumor data (6,426 images, 80:20 split) and a COVID-19 chest X-ray set (6,939 images; 5,898 training, 1,041 testing) were used with DenseNet-121 and XceptionNet [6], [7] under five-fold cross-validation. **Results.** In F1 score, DenseNet on brain-tumor data reached 98.36 with plain images, 93.36 with SKK, and 89.30 with the proposed scheme (4×4 blocks); on COVID-19 X-ray data the corresponding values were 96.25, 95.01, and 94.71 (6×6 blocks). XceptionNet showed the same ordering. In the filter-size ablation, larger blocks lowered F1 and, qualitatively, reduced recoverable features under both attacks. Grad-CAM [21] visualizations indicated that encrypted-image models remained learnable, with a larger region of interest for the proposed scheme.

Limitations and future scope. The security evaluation is based on visual comparison of attacked images rather than a quantitative security metric, and accuracy declines as protection increases, most noticeably on the MRI task. The authors note unsuitability for low-resolution images. Their stated future work includes additional smoothing methods, other datasets, attention modules, and GAN-generated encrypted images.

D. Semi-Supervised Medical Image Classification

Problem and objective. Fully supervised WSI classification needs exhaustive region-level annotation, while pseudo-label SSL suffers when few labels yield erroneous pseudo-labels [3]. The objective is to refine pseudo-labels using local and global context.

Methodology. A ResNet-50 patch classifier is first trained on labeled slides. Its predictions on unlabeled slides seed a graph-cut segmentation over a slide-level patch graph. The energy function combines a Gaussian-mixture regional term and a boundary term, and a regional-term correction weights confident predictions more heavily. The refined labels are added to training data for retraining. **Data.** Camelyon16 breast-cancer slides [22] and a collection of TCGA kidney-cancer slides were used, with labeled fractions of 1% to 20%.

Results. The study reports mean AUC over five random samplings. With 5% labeled slides, the proposed method reached 89.76 on Camelyon16 (full-supervision upper bound 94.62) and 92.77 on TCGA (upper bound 98.62); with 20% labeled slides it reached 92.92 and 96.15, respectively. It outperformed the compared baselines in most settings; one baseline, Self-Path [23], was slightly better on TCGA at 2% labeled data. Standard deviations were large at 1% labeled data, indicating instability under extreme label scarcity. **Limitations and future scope.** The patch feature is only the mean RGB color, and only breast and kidney cancer were tested. The authors plan pathology-specific features such as nuclear shape and morphology and evaluation on other cancers. Related pathology SSL work cited there includes teacher-student and self-supervised approaches [23], [24].

E. Personalized AI Healthcare Systems

Problem and objective. Fragmented care delivery makes it hard to find suitable providers, book appointments, track prescriptions, and maintain records [4]. The reviewed application aims to integrate these functions with AI-driven recommendations.

Methodology. A Flutter/Dart front end, Flask back end, and Firebase storage host logistic regression for symptom-to-specialist recommendation, a random forest for calorie estimation feeding an LLM-generated meal plan, doctor verification against the Indian Medical Register by web automation, OpenStreetMap-based navigation and pharmacy search, a chatbot (Allevia), voice input, and QR-code-gated EHR summaries. **Data.** A disease-prediction dataset from IEEE DataPort was split 80/20; synonym augmentation and LLM-assisted specialist mapping were used.

Results. The study reports that logistic regression reached nearly 97.24% precision after augmentation, after earlier attempts of 62%–70% accuracy, and that the calorie-prediction random forest reached 84.76% precision. It also reports qualitative comparison with Ada Health, Babylon Health, and MyFitnessPal. **Limitations.** The application had not been publicly deployed, and testing used simulated user data. The authors report map-rendering delays, synchronization lags, single-language support, generalized diet and exercise advice, and the absence of advanced encryption. This survey notes that the paper does not state whether augmented samples were separated from the test set before evaluation, that its feature-comparison table credits end-to-end encrypted sharing while its conclusion states that advanced encryption is absent, and that the 97.24% figure is a precision value rather than accuracy. **Related work.** Systems cited by [4] include symptom-based disease prediction with doctor recommendation [16], small language models for multimodal risk prediction with high computational cost [25], chatbots with transfer learning [26], NLP chatbots [27], [28], and diet and exercise guidance for chronic disease [17]. Their reported limitations, as summarized in [4], include small datasets, limited scope, and the need for real-time data.

VI. COMPARATIVE ANALYSIS

The reported performance values should be interpreted within the respective experimental settings because the studies use different datasets, healthcare tasks, modalities, and evaluation protocols (accuracy, F1, AUC, and precision). Table II therefore lists results next to their own metric and does not rank the studies. Entries marked “not reported” indicate information absent from the source paper.

TABLE I
OVERVIEW OF REVIEWED STUDIES

Study	Year	Healthcare Problem	Data / Modality	AI/ML Method	Main Contribution	Major Limitation
Arul Raj & Sugumar [1]	2023	COVID-19 detection and classification	Chest X-ray, CT, RT-PCR (Kaggle-sourced; 500 positive, 500 negative)	CNN feature extraction; concatenation / weighted-average fusion; softmax	Showed fused modalities outperforming single modalities in its setting	No clinical validation; limited comparison; reporting inconsistencies; small dataset
Huang et al. [2]	2022	Privacy of medical images in cloud DNN training	MRI brain tumor (6,426 images); COVID-19 chest X-ray (6,939 images)	SKK encryption plus block-based statistical smoothing; DenseNet-121, XceptionNet	Parameterized privacy-accuracy trade-off with large keyspace	Accuracy drop; attack evaluation visual only; unsuitable for low-resolution images
Shin et al. [3]	2022	Pathology classification with few annotations	WSIs: Camelyon16 (breast), TCGA (kidney)	ResNet-50 patch classifier; graph-cut pseudo-label refinement with GMM regional term	Graph-based refinement of pseudo-labels with confidence-weighted regional term	Mean-RGB features only; two cancer types; high variance at 1% labels
Saritha et al. [4]	2025	Access to care, specialist recommendation, personalized management	Symptom data (IEEE DataPort); health metrics; diet data; text and voice input	Logistic regression, random forest, NLP augmentation, LLM-based meal plan and chatbot (Groq API)	Integrated mobile platform: recommendation, booking, prescriptions, QR-gated records	Not publicly deployed; simulated-user testing; no advanced encryption; single language

TABLE II
TECHNICAL COMPARISON

Study	Learning Type	Model / Algorithm	Input Data	Fusion / Encryption / Graph Technique	Evaluation Metric	Key Result (as reported)
[1]	Supervised	CNN (VGGNet mentioned for feature extraction)	X-ray, CT, RT-PCR (clinical variables also described)	Concatenation; weighted averaging	Accuracy (precision, recall, F1 named but not reported)	Fused 92.1%; X-ray 80.3%; CT 85.6%; RT-PCR 89.2%
[2]	Supervised	DenseNet-121; XceptionNet	Encrypted MRI and chest X-ray images, 256×256	SKK + median/mean/max/min block smoothing	F1 score; qualitative attack recovery; Grad-CAM	DenseNet brain tumor: 98.36 plain, 93.36 SKK, 89.30 proposed; COVID-19: 96.25, 95.01, 94.71
[3]	Semi-supervised (pseudo-labeling)	ResNet-50 patch classifier	128×128 WSI patches	Patch graph; min-cut with GMM regional term and boundary term	AUC (mean of five samplings)	5% labeled: Camelyon16 89.76 (upper bound 94.62); TCGA 92.77 (upper bound 98.62)
[4]	Supervised, plus API-based LLM generation	Logistic regression; random forest	Binary symptom vectors, synonym-augmented text, height/weight/age/gender	Not applicable (NLP synonym augmentation; IMR web verification)	Precision (recall, F1 named)	LR about 97.24% precision; calorie RF 84.76% precision

TABLE III
HEALTHCARE CAPABILITY COMPARISON

Capability	[1] Arul Raj & Sugumar	[2] Huang et al.	[3] Shin et al.	[4] Saritha et al.
Medical imaging	Yes (X-ray, CT)	Yes (MRI, X-ray)	Yes (WSI)	No
Multimodal data	Yes (imaging + RT-PCR)	No	No	Partial: text/voice and records; no fusion model reported
Privacy protection	Not addressed	Yes (central contribution)	Not addressed	Partial: IMR verification, QR-gated access; encryption absent per authors
Semi-supervised learning	No	No	Yes (central contribution)	No
Graph-based learning	No	No	Yes (graph cut)	No

Capability	[1] Arul Raj & Sugumar	[2] Huang et al.	[3] Shin et al.	[4] Saritha et al.
NLP	No	No	No	Yes (symptom text, augmentation, chatbot)
Personalization	No	No	No	Yes (diet, exercise, location)
Doctor recommendation	No	No	No	Yes (specialist recommendation)
Explainability	Not reported	Partial: Grad-CAM	Partial: qualitative overlays	Not reported
Clinical deployment	Not clinically validated (stated)	Not reported	Not reported	Not deployed; simulated-user testing

TABLE IV
CHALLENGES AND SOLUTIONS

Challenge	Evidence in Literature	Existing Approach	Remaining Gap
Scarce annotations	WSIs are very large and need multi-magnification expert annotation [3]	Graph-refined pseudo-labels [3]	Only two cancer types; simple features; unstable at 1% labels
Ambiguous single-source diagnosis	COVID-19 symptoms overlap; RT-PCR slow; imaging varies [1]	Fusion of X-ray, CT, RT-PCR [1]	No clinical validation; inconsistent reporting; small data
Privacy of cloud training	Semi-honest servers; HIPAA constraints [2]	Learnable encryption with smoothing [2]	Accuracy loss; no quantitative security metric; key attacks evolve
Privacy of patient records in apps	QR-gated EHR access, IMR verification [4]	Access restriction and verification [4]	Authors state advanced and end-to-end encryption are absent
Deployment and real-time use	Sync lags, map delays, one language [4]	Cloud scaling proposed as future work [4]	No public deployment or outcome data
Evaluation comparability	Accuracy [1], F1 [2], AUC [3], precision [4]	Each study uses its own protocol	No shared benchmark across the layers
Interpretability	Grad-CAM and overlays only [2], [3]	Qualitative visualization	No clinician-validated explanations

VII. RESEARCH GAPS AND OPEN CHALLENGES

Each gap below is tied to evidence in the reviewed studies; items marked *inferred* are conclusions drawn by this survey from the comparison.

- 1) *Limited clinical validation.* The COVID-19 fusion study states it has not been validated clinically [1], and the application was tested only with simulated data [4]. The imaging and pathology studies use public datasets and report no clinical deployment [2], [3].
- 2) *Dataset limitations.* The fusion study uses 1,000 samples [1]; the pathology study covers breast and kidney cancer only [3]; the application relies on one symptom dataset and notes that related systems were limited by small data [4].
- 3) *Generalization across hospitals and populations.* The fusion study calls for testing on larger, more diverse populations [1]. Camelyon16 includes two centers with different scanners [3], but no study reports cross-hospital generalization (*inferred*).
- 4) *Data privacy and security.* Learnable encryption addresses image privacy [2], while the application's authors acknowledge missing advanced encryption [4]. The fusion and pathology studies do not discuss privacy.
- 5) *Dependence on labeled data.* Pathology classification requires costly annotation [3]; the other studies train supervised models without addressing label cost (*inferred*).

- 6) *Interpretability and explainability.* Only Grad-CAM visualization [2] and qualitative overlays [3] are provided; explanation quality is not clinically assessed.
- 7) *Multimodal integration.* Fusion is demonstrated for one task [1], with internal inconsistencies about which modalities enter the fusion; its combination with privacy or label-efficient learning is not examined (*inferred*).
- 8) *Computational cost.* Complex models motivate cloud training [2], and a related multimodal model is reported to have very high computational cost [25], yet none of the four studies reports training time or resource usage.
- 9) *Real-time deployment.* RT-PCR delays motivate [1], and the application reports synchronization lags [4]; real-time performance is not quantified.
- 10) *Interoperability.* The application proposes hospital integration as future work and cites HIPAA [4]; no study evaluates integration with existing hospital systems.
- 11) *Standardization of evaluation.* Metrics, splits, and datasets differ (Table II), preventing valid cross-study comparison.
- 12) *Reliability and robustness.* Pseudo-labeling is unstable at 1% labels [3]; learnable encryption is vulnerable to dedicated attacks [18]; the fusion study reports a single split.
- 13) *Privacy-performance trade-offs.* Larger smoothing blocks improve attack resistance but lower F1 [2]. The trade-off is quantified for images only; the survey found no analogous analysis for multimodal or semi-supervised settings (*inferred*).

VIII. CROSS-STUDY SYNTHESIS

The four studies address different layers of one chain, from imaging through deep learning, multimodal fusion, privacy preservation, label-efficient learning, and personalized services toward clinical decision support (Fig. 2). Medical imaging supplies the data and deep learning supplies the predictive engine [1], [2], [3]. Multimodal fusion reduces dependence on a single imperfect source [1]. Privacy preservation determines whether sensitive data can leave the institution for training [2]. Semi-supervised learning reduces the annotation burden that limits access to large labeled sets [3]. Personalized applications translate predictions and records into services for patients and clinicians [4]. These are complementary directions, not independent technologies.




Fig. 2. Conceptual Framework Derived From the Reviewed Literature. It is a survey figure, not an implemented system or experimental result.

The synthesis also reveals structural gaps. The privacy study evaluates image classification only, while the application evaluates records and recommendation only, so protection of images and records is not combined. The label-efficient method has not been combined with encryption or multimodal inputs, and the fusion model does not use label-efficient learning. Clinical decision support is shown by none of the studies as an end-to-end outcome. Possible combinations, such as encrypted multimodal learning, are therefore a research direction suggested by this survey, not a result of any reviewed paper.

IX. FUTURE RESEARCH DIRECTIONS

A. Directions Reported by the Original Studies

The fusion study calls for larger, more varied datasets and tests across populations and settings [1]. The encryption study proposes additional smoothing filters, other datasets, attention modules, GAN-generated encrypted images, and super-resolution to handle low-resolution images [2]. The pathology study proposes nuclear shape and morphology features and evaluation on further cancers [3]. The application study proposes personalized AI-based plans, scalable cloud infrastructure, end-to-end encryption and HIPAA compliance, accessibility features, hospital integration, and RFID [4].

B. Directions Derived From This Survey

The following directions follow from the gaps in Section VII and are not claimed by the original authors: (1) privacy-preserving multimodal healthcare AI, combining fusion with secure learning; (2) reduced reliance on labeled pathology data together with privacy protection; (3) federated or distributed medical AI as a possible alternative to centralized cloud training, which none of the seed studies evaluates; (4) explainable and robust AI with clinician-assessed explanations; (5) integration of imaging with EHR and clinical information; (6) real-time personalized healthcare with measured latency; (7) larger and more diverse datasets and cross-hospital studies; (8) shared evaluation protocols; and (9) clinically validated AI systems.

X. LIMITATIONS OF THE SURVEY

This survey rests on four seed papers and the references they cite, so it is not exhaustive and may reflect source selection bias. No systematic database screening was performed. The studies differ in healthcare tasks, datasets, metrics, and splits, which makes direct performance comparison invalid. Some assessments rely on what the seed papers report, including secondary work described only through the citing paper. Several statements about inconsistencies are observations from the text and figures of the papers and could reflect presentation errors rather than methodological ones. Finally, the conceptual figures summarize interpretation rather than evidence.

XI. CONCLUSION


AI is used across healthcare, from image interpretation to patient-facing management. The reviewed studies show that medical image analysis remains central, that fusing complementary modalities improved results in one setting [1], that privacy-preserving training is feasible but trades accuracy for protection [2], that graph-refined pseudo-labels can reduce annotation needs [3], and that personalized applications can integrate recommendation, scheduling, and record access [4]. Shared challenges include limited clinical validation, small or narrow datasets, unquantified generalization, privacy-performance trade-offs, limited explainability, and non-comparable evaluation. Future healthcare AI will need to be reliable, secure, explainable, and clinically validated, and combining these layers remains an open research problem.

REFERENCES

Items [1]–[4] are the seed studies. Items [5]–[28] are secondary references cited within the seed papers; their details are reproduced as given there and were not independently verified against a publisher database.

- [1] A. M. Arul Raj and R. Sugumar, "Multi-modal fusion of deep learning with CNN based COVID-19 detection and classification combining chest X-ray images," in Proc. 7th Int. Conf. Intell. Comput. Control Syst. (ICICCS), 2023, pp. 569–575, doi: 10.1109/ICICCS56967.2023.10142910.
- [2] Q.-X. Huang, W. L. Yap, M.-Y. Chiu, and H.-M. Sun, "Privacy-preserving deep learning with learnable image encryption on medical images," IEEE Access, vol. 10, pp. 66345–66355, 2022, doi: 10.1109/ACCESS.2022.3185206.
- [3] H.-K. Shin, K.-H. Uhm, K. Choi, Z. Xu, S.-W. Jung, and S.-J. Ko, "Graph segmentation-based pseudo-labeling for semi-supervised pathology image classification," IEEE Access, vol. 10, pp. 93960–93970, 2022, doi: 10.1109/ACCESS.2022.3204000.
- [4] A. N. Saritha, G. Kalnoor, P. R. Bhat, P. A. Rao, P. Prajwal, and P. T. Pradeep, "Alleviate: An AI-powered mobile application for personalized healthcare management and doctor recommendation," IEEE Access, vol. 13, pp. 215822–215832, 2025, doi: 10.1109/ACCESS.2025.3647089.
- [5] C. L. Srinidhi, O. Ciga, and A. L. Martel, "Deep neural network models for computational histopathology: A survey," Med. Image Anal., vol. 67, Jan. 2021, Art. no. 101813.
- [6] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), Jul. 2017, pp. 4700–4708.
- [7] F. Chollet, "Xception: Deep learning with depthwise separable convolutions," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), Jul. 2017, pp. 1251–1258.
- [8] F. Hao, X. Liu, M. Li, and W. Han, "Accurate kidney pathological image classification method based on deep learning and multi-modal fusion method with application to membranous nephropathy," Life, vol. 13, no. 2, Jan. 2023, doi: 10.3390/life13020399.
- [9] M. Tanaka, "Learnable image encryption," in Proc. IEEE Int. Conf. Consum. Electron.-Taiwan (ICCE-TW), May 2018, pp. 1–2.
- [10] W. Sirichotedumrong, Y. Kinoshita, and H. Kiya, "Pixel-based image encryption without key management for privacy-preserving deep neural networks," IEEE Access, vol. 7, pp. 177844–177855, 2019.

- [11] J. E. Van Engelen and H. H. Hoos, "A survey on semi-supervised learning," *Mach. Learn.*, vol. 109, no. 2, pp. 373–440, 2020.
- [12] D.-H. Lee, "Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks," in *Proc. Workshop Challenges Represent. Learn.*, 2013, p. 896.
- [13] E. Arazo, D. Ortego, P. Albert, N. E. O'Connor, and K. McGuinness, "Pseudo-labeling and confirmation bias in deep semi-supervised learning," in *Proc. Int. Joint Conf. Neural Netw. (IJCNN)*, Jul. 2020, pp. 1–8.
- [14] M. N. Rizve, K. Duarte, Y. S. Rawat, and M. Shah, "In defense of pseudo-labeling: An uncertainty-aware pseudo-label selection framework for semi-supervised learning," in *Proc. Int. Conf. Learn. Represent.*, 2020.
- [15] Y. Boykov and G. Funka-Lea, "Graph cuts and efficient N-D image segmentation," *Int. J. Comput. Vis.*, vol. 70, no. 2, pp. 109–131, Nov. 2006.
- [16] S.-K. Tan, S.-C. Chong, K.-K. Wee, and L.-Y. Chong, "Personalized healthcare: A comprehensive approach for symptom diagnosis and hospital recommendations using AI and location services," *J. Informat. Web Eng.*, vol. 3, no. 1, pp. 117–135, Feb. 2024.
- [17] J. Tseng et al., "An interactive healthcare system with personalized diet and exercise guideline recommendation," in *Proc. Conf. Technol. Appl. Artif. Intell. (TAAI)*, 2015, pp. 525–532.
- [18] A. H. Chang and B. M. Case, "Attacks on image encryption schemes for privacy-preserving deep neural networks," 2020, arXiv:2004.13263.
- [19] W. Sirichotedumrong and H. Kiya, "A GAN-based image transformation scheme for privacy-preserving deep neural networks," in *Proc. 28th Eur. Signal Process. Conf. (EUSIPCO)*, Jan. 2021, pp. 745–749.
- [20] Y. Ding et al., "DeepEDN: A deep-learning-based image encryption and decryption network for internet of medical things," *IEEE Internet Things J.*, vol. 8, no. 3, pp. 1504–1518, Feb. 2021, doi: 10.1109/JIOT.2020.3012452.
- [21] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 618–626.
- [22] B. E. Bejnordi et al., "Diagnostic assessment of deep learning algorithms for detection of lymph node metastases in women with breast cancer," *JAMA*, vol. 318, no. 22, pp. 2199–2210, Dec. 2017.
- [23] N. A. Koohbanani, B. Unnikrishnan, S. A. Khurram, P. Krishnaswamy, and N. Rajpoot, "Self-Path: Self-supervision for classification of pathology images with limited annotations," *IEEE Trans. Med. Imag.*, vol. 40, no. 10, pp. 2845–2856, Oct. 2021.
- [24] J. Silva-Rodriguez, A. Colomer, J. Dolz, and V. Naranjo, "Self-learning for weakly supervised Gleason grading of local patterns," *IEEE J. Biomed. Health Informat.*, vol. 25, no. 8, pp. 3094–3104, Aug. 2021.
- [25] L. Björkdahl, O. Pauli, J. Östman, C. Ceccobello, S. Lundell, and M. Kjellberg, "Towards holistic disease risk prediction using small language models," 2024, arXiv:2408.06943.
- [26] M. M. Rahman et al., "Web-based multiple disease prediction system with transfer learning and healthcare chatbot," in *Proc. IEEE 3rd Int. Conf. Comput. Mach. Intell. (ICMI)*, Jan. 2024, pp. 1–5.
- [27] R. Surve, T. Purohit, R. Joseph, and P. Shaikh, "HealthCare chatbot using machine learning and NLP," in *Proc. 6th Int. Conf. Adv. Sci. Technol. (ICAST)*, Dec. 2023, pp. 411–416.
- [28] R. Umbare, R. Patil, T. Mukund, A. Rathod, and P. Rajurkar, "A smart healthcare application using artificial intelligence (AI) and machine learning (ML)," in *Proc. 5th Int. Conf. Image Process. Capsule Netw. (ICIPCN)*, Jul. 2024, pp. 204–210.



10.22214/IJRASET



45.98



IMPACT FACTOR:
7.129



IMPACT FACTOR:
7.429



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24*7 Support on Whatsapp)