Iraqi Journal for Computers and Informatics
Available online: https://ijci.uoitc.edu.iq/index.php/ijci
Volume 52, issue 1. 2026

Print ISSN: 2313-190X, Online ISSN: 2520-4912

Multimodal Deep Learning in Healthcare Recommender Systems: A Review of CNN-Based Architectures

Hikmat Z. Neima[a], Maha S. Abdulridha[b], Rana M. Ghadban[c], Ghaihab H. Adday[a]

[a] *Department of Computer Science, College of CSIT, University of Basrah, Iraq.*
[b] *Department of Computer Information Systems, College of CSIT, University of Basrah, Iraq.*
[c] *Department of Intelligent Medical Systems, College of CSIT, University of Basrah, Iraq.*

ARTICLEINFO

Keywords:

Multimodal Deep Learning,
Healthcare Recommender Systems,
Convolutional Neural Networks,
Explainable AI,
Federated Learning

ABSTRACT

Recent multimodal deep learning approaches are changing how healthcare systems process and integrate widely different types of patient information. Based on this, in this review we investigate the Convolutional Neural Network CNN-based architectures in Healthcare Recommender Systems HRS, which incorporate a variety of data from different sources, such as medical imaging, Electronic Health Records EHRs, wearable sensor streams, clinical narratives. At the core of this ecosystem, CNNs learn hierarchical representations from high-dimensional visual and other unstructured medical data and are increasingly the feature-extraction backbone of modern healthcare recommendation pipelines. We first discuss CNN-based architectures for HRS and then examine their combination with traditional recommendation schemes such as collaborative filtering, hybrid recommendation techniques and other multimodal fusion methods. The review furthermore covers new research developments to leverage CNN-based HRS, including vision–language models, reinforcement learning, causal inference, fairness-aware optimization, federated learning and edge/IoMT deployment. This review reports methodological advances and recent challenges. These challenges include generalization, interpretation , privacy preserving, and clinical integration. Taken together, these analyses summarizes the progress of recent research and also provide open challenges. The goal is to guide the design of transparent and scalable systems. These systems should be also ethically sound. They focus on CNN-based healthcare recommender systems for next-generation personalized medicine.

1. INTRODUCTION

The volume of medical data is rising very rapidly. As a result, healthcare is quickly moving toward digital systems. The change is driven by the many types of health data now available. These types include Electronic Health Records EHRs, medical images, and data collected from wearable devices. Healthcare providers seek cost-effective and tailored interventions. Diverse data sources provide many opportunities for improving medical decision-making and require new analytical methods. Advanced neural network architectures and other deep learning methods have been designed to tackle this issue. Computer Vision CV technology is used to interpret X-rays, MRI scans, and pathological slides for clinical diagnosis. Computer vision techniques such as Convolutional Neural Networks CNNs have been reported to enhance image classification, object detection, segmentation, and feature extraction. These successful approaches are particularly recognized for cancer detection and ophthalmology [2]. At the same time, Healthcare Recommender Systems HRS have emerged as a major form of personalized

E-mail address:
hikmat.taher@uobasrah.edu.iq [a]
ghaihab.aday@uobasrah.edu.iq [a]
maha.abdulradha@uobasrah.edu.iq [b]
rana.ghadban@uobasrah.edu.iq [c]

Corresponding* : Hikmat Z. Neima

Received 13 February 2026, Accepted 21 April 2026

DOI: 10.25195/ijci.v52i1744.

health delivery services. They are now widely used by healthcare professionals and patients for treatment recommendation, informational support, and lifestyle advice [3]. However, existing collaborative and content-based filtering techniques in HRS are inadequate for high-dimensional and complex healthcare data [4].

1.1. Motivation and Background

The integration of Artificial Intelligence (AI) into healthcare has opened new opportunities for automated decision-making and individualized patient management. Computer vision techniques and CNN-based models are introduced as medical images analyzer and as clinician's assistant in diagnosis and treatment planning. These approaches have proven their effectiveness in medical domains characterized. This characterization is done by large imaging volumes and high interpretative complexity such as radiology and ophthalmology [5].

Healthcare recommender systems are widely used to support treatment planning , lifestyle adjustments, and continuously monitoring patients. By incorporating clinical records , patient history, and individual preferences, HRS aim to generate evidence-informed recommendations. The adoption of deep learning techniques , especially CNN-based models , enhances feature extraction capabilities , enabling HRS to generate more informative , context-aware , and adaptive recommendations [6].

1.2. Scope and Objectives

In this review , we report the role of computer vision and CNN-based architectures in enhancing the performance of healthcare recommender systems. We also highlight how recent progress in deep learning, particularly CNNs, has advanced healthcare recommendation tasks by enabling the utilization of visual and contextual information obtained from heterogeneous data sources.

The primary objectives of this review are:

- To provide a comprehensive overview of CNN fundamentals and their applications in medical imaging.
- To analyze the foundational concepts of healthcare recommender systems and explore how CNNs are incorporated within their architectural design.
- To review recent developments across key emerging areas, including fairness , cold-start strategies , interpretability , reinforcement learning , causal modeling , and ethical considerations in HRS.
- To identify existing research gaps and propose potential future research directions.

1.3. Organization of the Review

The remainder of this review is structured as follows:

- Section 2 introduces the fundamental concepts of medical imaging, CNN architectures, and HRS principles.
- Section 3 discusses the integration of CNNs within HRS frameworks.
- Section 4 outlines recent advances in CNN-based HRS.
- Section 5 reviews representative CNN-based HRS solutions, datasets, evaluation methodologies, and limitations.
- Section 6 discusses system-level enhancements such as federated learning and edge/IoMT deployment aspects.
- Section 7 outlines remaining challenges and open research issues.
- Section 8 concludes the review and outlines directions for future investigation.

Unlike the previous surveys that give general overviews of deep learning applications in healthcare domains, this review focuses specifically on CNN-based multimodal healthcare recommender systems. It emphasizes the convergence of heterogeneous data sources with CNN architectures and also highlights emerging trends like explainable AI, federated learning, and fairness-aware modeling, providing a more applied perspective and concentrated approach.

1.4. Review Methodology

In this literature review, we summarize recent advances in multimodal deep learning, and CNN architectures for HRS , and analyzed reported contributions. Literature analysis sources covered Scopus, Web of Science, IEEE Xplore, PubMed, and ScienceDirect. The terms for the review included "healthcare recommender systems," "convolutional neural networks," "multimodal deep learning," "federated learning in healthcare," "vision–language models," and "AI-driven personalization."
The review includes the literature available up to 2025. Inclusion criteria comprised peer-reviewed journal articles or conference papers published between 2015 and 2025, studies that combine CNNs with healthcare recommendation frameworks or multimodal architectures, and articles reporting empirical or experimental evaluations or architectural insights. Articles that covered only generic recommender systems without relevance to the field of healthcare, or those involving diagnostic imaging without recommendation components were excluded. All abstracts were screened for methodological relevance and contribution coverage before analyzing full texts. Finally, representative and high-impact studies were chosen in order to be able to cover architectural-level and system-level perspectives. No formal meta-analytic statistical method was employed; instead, this review emphasizes conceptual synthesis, comparative analysis, and the mapping of future research directions in relevant sub-areas such

as fairness, explainability, reinforcement learning, causal inference, federated learning, and edge deployment. Approximately 200 studies were detected across the selected databases. After removing duplicates and screening titles and abstracts for relevance, nearly 110 studies remained for full-text review. Finally, 95 representative studies were chosen based on the explicit inclusion and exclusion criteria. Through a structured screening approach by focusing on relevance, quality, and contribution to the field, this selection ensured coverage of influential and methodologically relevant contributions for CNN-based healthcare recommender systems. The screening process followed multiple stages, including duplicate removal, title and abstract screening, and full-text review. This review gives a focus on CNN-based architecture of multimodal healthcare recommender systems. It emphasizes on the fusion of visual, clinical, and behavioral data. It also considers new emerging researches such as fairness-aware modeling, federated learning, and multimodal fusion techniques, providing a structured view of forthcoming challenges and opportunities.

2. FUNDAMENTALS AND BACKGROUND

This Section presents the basic principles of medical imaging, deep learning, and recommender algorithms that support the integration of CNN-based computer vision within HRS. These fundamental principles establish the conceptual framework necessary to understand the interdisciplinary developments discussed in the following Sections.

2.1. Healthcare Recommender Systems (HRS)

Although CNNs have made a large impact on medical image analysis, their importance for personalized decision-making becomes more significant when embedded in a recommendation-based system. HRSs are adopting deep learning techniques for providing data-driven, patient-centered clinical decision support. HRSs focus on the development of personalized medical advice for treatment plans, dietary guidance, and clinical resources, by considering patient characteristics and preferences. Such systems utilize structured and unstructured health data to improve decisions and personalization.

Categories of HRS**:**

- Content-Based Filtering**:** In content-based filtering, it has been the feature that its suggestions are made for when the item attributes are similar to user profiles. Common factors can be calculated with cosine similarity:

$$cos(u,v) = \frac{\sum_i^n u_i v_i}{\sqrt{\sum_i^n u_i^2}\sqrt{\sum_i^n v_i^2}} \quad (1)$$

It measures the matching with item properties; and this criterion works well in cases where explicit metadata is accessible.

- Collaborative Filtering**:** It is built by using historical user interactions with items. A usual approach is the latent factor modeling via matrix factorization:

$$R \approx UV^T \quad (2)$$

Where R denotes the matrix of user-item interaction, while U and V are the matrices of latent factor that represent the preferences of user and item learned during training [8].

- Hybrid Systems: Hybrid recommenders combine content-based and collaborative approaches to balance personalization and generalization, thereby mitigating challenges such as cold-start and data sparsity [9]. These systems enhance robustness, particularly when input data are incomplete or inconsistent.

2.2. Medical Imaging in Healthcare

Medical imaging is an essential component not only for disease diagnosis but also for long-term patient follow-up. Imaging is employed by physicians to visualize internal anatomical structures and pathological conditions using a non-invasive approach. Imaging methods in general are available as follows: X-ray , Magnetic Resonance Imaging MRI , Computed Tomography CT, ultrasound , and Positron Emission Tomography PET. Both modalities have special clinical merits; for example, MRI produces a high degree of soft tissue contrast, while CT is optimized in the evaluation of bone and acute trauma [10]. Medical images are inherently complex and high-dimensional, by nature. Manual coding may be time consuming and inter-observer variability is expected, particularly between slices of 3D imaging modalities. As a result, machine learning and computer vision based automatic diagnostic systems are increasingly being utilized to improve diagnostic precision and efficiency [11]. Medical images can be mathematically represented as multidimensional arrays:

$$I(x,y,z,c) \epsilon\ R^{H\times W\times D\times C} \quad (3)$$

Where a medical image is represented by its three spatial dimensions, height (H), width (W), and depth (D) for three-dimensional modalities such as CT or MRI, along with the channel dimension (C), which denotes the number of image channels. For example, C=1 for grayscale images and C=3 for RGB images. Various analytical techniques, including segmentation, detection, and classification, are applied to these multidimensional inputs to support accurate clinical diagnosis.

2.3. Convolutional Neural Networks

Medical imaging generates high-resolution visual data that contain complex spatial patterns. Advanced computational techniques are required to extract meaningful representations from such data. Deep learning, particularly CNNs , became a dominant approach to deal with medical imaging data processing. This is because of its robust representational capabilities. Modern healthcare image analysis heavily relies on deep learning methodologies.

Deep learning, a subset of machine learning, enables hierarchical feature learning through multiple processing layers. CNNs are especially effective for image analysis because they capture multi-scale spatial features while preserving local structural relationships [12].

A standard CNN architecture typically consists of:

1) Convolutional layers.

2) Activation functions (e.g. ReLU).

3) Pooling layers.

4) Fully connected layers.

It is possible to express the convolutional operation as:

$$Y_{i,j}^{(K)} = (X * W^{(K)})_{i,j} + b^{(K)} \qquad (4)$$

Where X denotes the input, $W^{(K)}$ is the kth filter, and $b^{(K)}$ is the bias term. Non-linearities are introduced through activation functions, such as:

$$f(x) = max\,(0, x) \qquad (5)$$

To decrease computational complexity and enhance spatial invariance, pooling layers are utilized. Softmax function is typically applied in the final classification layer:

$$\hat{y}_i = \frac{e^{z_i}}{\sum_j e^{z_j}} \qquad (6)$$

CNNs have demonstrated strong performance in medical imaging tasks, including tumor detection, organ segmentation, and disease classification [13]. Prominent architectures include AlexNet, VGG, ResNet, and U-Net, particularly for segmentation applications [14].

Due to the limited availability of annotated medical datasets, transfer learning is commonly adopted. Pre-trained CNN models—often initialized on large-scale datasets such as ImageNet—are fine-tuned on domain-specific medical data to improve generalization and reduce overfitting [12].

3. INTEGRATION OF CNN IN HEALTHCARE RECOMMENDER SYSTEMS

3.1. CNNs in Medical Imaging

Deep Convolutional Neural Networks CNNs have a strong performance in extracting complex spatial features from high-dimensional medical images. As a result, they have achieved extensive use in medical imaging. CNN-based visual applications in clinical fields are highly effective due to their hierarchical structure and ability to represent non-linear relationships. They are now employed in other imaging modalities such as X-ray, Computed Tomography CT, Magnetic Resonance Imaging MRI, Positron Emission Tomography PET and ultrasonography [15]. Medical imaging tasks conducted by CNN can be categorized such as:

- Image classification: The classification of those input images, e.g., whether a chest X-ray is normal or is associated with pneumonia.
- Object detection: Detection and localization of interest areas in an image such as tumor boundary or the position of a tumor.
- Semantic segmentation: Marking each pixel on an image with the marking of a class, used to correctly name organs, lesions and tumors on the image.

Several studies have utilized CNNs in numerous research literature to show their clinical utility. For example, Shen et al. employed deep CNN architectures to classify breast cancer with a similar performance to expert radiologists [15]. Likewise, Jayawardena et al. reported on a diagnostic framework and applied CheXNet to the chest X-ray images training a CNN [16], showing that a radiologist's detection of pneumonia was greatly improved. And these papers illustrate that deep learning models may become smart and consistent assistants in diagnosing an ailment and improving decision-making during clinical care. Some CNN architectures have addressed individual imaging problems. Due to their symmetric encoder–decoder layout and skip connections, U-Net architectures continue to be recognized and widely deployed in biomedical segmentation in a way that preserves spatial information [17]. Residual connections introduced in ResNet enable effective training of deep networks by alleviating vanishing-gradient issues. Attention-based mechanisms such as the Convolutional Block Attention Module CBAM have been recently proposed to improve interpretability on medical or biometric images by emphasizing more complex or clinically salient regions [1]. This architectural upgrade is incremental and represents an increasing aspect of the domains in automating medical image analysis. To illustrate, Gradient-weighted Class Activation Mapping (Grad-CAM) provides visual explanations that highlight the image regions most influential in generating model predictions, thereby improving transparency and interpretability. Finally, we summarize these imaging capabilities in the next Section, embedded in a personalized health recommender system.

3.2. Deep Learning for Personalized Healthcare

Recent research has shown that deep learning can be used for recommender systems, especially in healthcare. CNNs can process unstructured data such as radiographs, patient charts, and wearable sensor data, making visual and contextual information available for recommendation systems [7]. For example, feature extraction combined with collaborative filtering methods can be combined to make health-related insights more powerful in RBM–CNN architectures [19]. Other methods such as neural collaborative filtering based on deep autoencoders, CNN-LSTM models, and hybrid approaches (e.g., HCLNet for healthcare recommendation tasks) have also been described [20]. These hybrid architectures make use of spatial and sequential information from the learning environment to increase classification and recommendation potential. They support prediction of future health events, help preventive strategies, and also adapt themselves to the ever-changing patient data [21]. By automated treatment planning, medication selection, and disease prediction, advanced patient-centered care has benefited immensely from deep learning. Nowadays, the trend on these topics changes dramatically. With a healthcare system turning to computational recommender frameworks, deep learning models are increasingly employed to interpret increasingly large, intricate data files to create personalized recommendations. Currently, in patient-centric clinical conditions the recommender systems of modern medicine increasingly integrate new deep learning models to continually adapt to their individualized care [22], [23]. The same is true for clinical medicine. These systems recognize patterns within multimodal data by learning from heterogeneous sources such as EHRs, genomic data, wearable sensor outputs, and lifestyle records. They then generate personalized recommendations for chronic illness, including diabetes, cardiovascular disease, or mental health disorders [7]. CNNs were initially built for visual data processing. However, they are now used in hybrid and multimodal architectures for designing personalized healthcare applications. For instance, CNNs have been deployed in dietary recommender systems by processing food images and in tracking physical activity by visual or sensor-based inputs. As a result, the integration of these signals with structured clinical data improves contextual awareness in HRS frameworks [24]. Deep learning models can analyze large-scale, multidimensional, and heterogeneous data. This makes them essential technology in personalized healthcare development [9]. Such architectures embed CNNs with RNNs to fuse medical records, diagnostic images, sensor streams, and genomic data. They also enable real-time generation of customized advice on medicine to give personalized patient advice with the support of RNNs-based architectures. CNNs are also used to generate representations from unstructured inputs such as images, text reports, and wearable sensor data. They are employed to generate models for patient trajectory and possible health incidents. These models are often combined with temporal models such as RNNs or TCNs [25]. Moderate patterns, for instance, emerging in some visual history or symptom trends might be a basis for preventing screening or early intervention. Hybrid deep learning models have been employed in particular to recommend suggested new screening for high-risk individuals and to recognize and analyze similarity in large healthcare data to inform treatment options for the treatment of the high-risk patient [11]. Such approaches also provide population-level solutions for health by detecting trends that are emerging in specific demographic groups. Such intelligent solutions can improve patient engagement. They also support personalized lifestyle recommendations and medication reminders in a timely manner. These systems, powered by deep learning in mobile health (mHealth), can monitor treatment adherence and detect complications. They can also support real-time intervention. This capability supports proactive and preventive care [19]. These immediate feedback outlets also facilitate open channels of communication between the healthcare providers and patients.

3.3. Hybrid Systems: CNN and Recommender Models

One of the most recent hybrid systems is to integrate CNN-based feature extraction into collaborative filtering. They have strong results in the clinical field, where patients with differing symptoms tend to respond similarly to particular treatments. Explainable AI (XAI) methods have recently been presented for CNN-based healthcare recommender systems through the concepts of transparency and trust. This explains the meaning of these recommendations and may, in turn, lead to the recommendations being more likely to be accepted by both professionals and patients [3]. In CNN-based HRS methods, attention mechanisms and saliency maps are often used to prioritize information that may render the useful recommendation (e.g., the area in an image or a particular clinical marker). CNN recommender systems have been explored in applications ranging from mental health coaching, to adapted exercise programs, medication adherence monitoring, to pandemic-related health surveillance applications [27]. These applications foster patients' involvement and offer organization decision-making tools for disease prevention and management. These hybrid CNN–HRS architectures also highlight key issues with data privacy, model bias, and responsible usage that we discuss more rigorously further on (e.g., fairness, explainability, ethical aspects). Based on current research findings, hybrid architectures that incorporate CNNs into healthcare recommender systems are transforming intelligent healthcare service delivery. In next-generation health support systems, multimodal data processing, real-time personalization, and interpretability mechanisms are critical elements. For multimodal healthcare recommender systems, CNNs are primarily used to extract hierarchical feature representations from medical images. These features are further combined with structured and sequential data through fusion mechanisms, which provide complementary information across modalities. The achievement of such systems relies on the alignment of heterogeneous data sources together and the capability to capture complex interactions.

Hybrid models based on CNNs, together with recurrent or attention-based models, also promote temporal learning and contextual learning. Representative models are summarized in Table 1 in terms of data modality, architectural design, key contributions, and limitations.

Table 1: *Comparative Analysis of CNN-Based Healthcare Recommender Systems*

Model / Ref.	Data Modality	Method / Architecture	Key Contribution	Limitations
Shen et al. [15]	Medical Imaging (Breast Cancer)	CNN-based classification	Achieved radiologist-level performance in cancer detection	Limited generalization across datasets
Jayawardena et al. [16]	Chest X-ray	CNN (CheXNet)	Improved pneumonia detection accuracy	Requires large labeled datasets
HCLNet [20]	Multimodal (EHR + Imaging)	CNN + LSTM hybrid	Combines spatial and temporal features	Increased model complexity
RBM–CNN [19]	Clinical + Behavioral	CNN + RBM	Enhanced recommendation quality using feature fusion	Limited interpretability
DeepReco [81]	Mobile Health Data	CNN + Collaborative Filtering	Improved personalized recommendation in mHealth	Latency and scalability issues
DGCF [82]	User–Item Graph Data	CNN + GNN	Better interaction modeling using graph structure	Computationally expensive
Al-Qazzaz et al. [80]	Clinical + Behavioral	CNN + KNN	Personalized ASD treatment support	Limited dataset diversity
Iwendi et al. [84]	IoMT + Nutrition Data	CNN-based recommender	Multimodal dietary recommendation	Data heterogeneity challenges
Ponselvakumar et al. [85]	Multimodal Medical Data	CNN-based multimodal system	Clinical trial matching and precision medicine	Requires multimodal alignment
Subramanian et al. [88]	Emotion + Sensor Data	CNN-based emotion recognition	Mental health monitoring and recommendation	Real-time deployment constraints

A comparative analysis of the reviewed studies has reflected notable differences in architectural design and learning strategies. Hybrid models integrating CNNs with collaborative filtering or sequential learning methods tend to improve personalization performance, while multimodal and graph-based models improve contextual understanding and relational modeling. All of these strategies, however, come with increased computational burden and limited generalization across diverse clinical datasets. These results emphasize the crucial compromises among performance, interpretability, and scalability. Figure 1 presents the overall architecture of a CNN-based multimodal healthcare recommender system, including data integration, feature extraction, fusion, and recommendation stages.

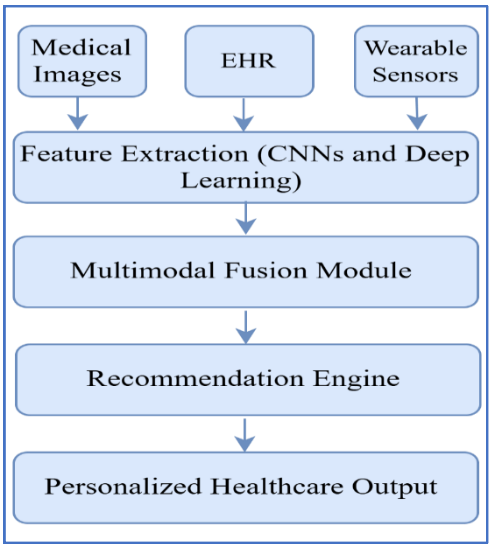

Figure 1: *Conceptual architecture of a CNN-based multimodal HRS integrating heterogeneous data sources, feature extraction, multimodal fusion, and recommendation generation.*

4. RECENT ADVANCES IN HRS

4.1. Cold-Start Solutions and Cross-User Generalization in HRS

HRS often struggle with cold-start due to insufficient historical data for new patients, making it difficult to generate accurate and personalized recommendations. This issue becomes particularly critical in time-sensitive clinical environments where decisions must be made for first-time users. To address this limitation, recent research has proposed meta-learning, zero-shot learning, and hybrid schemes to develop appropriate predictions based on the information from the patient, in other words adaptive systems with small and accurate input, to overcome this limitation. In [28], the authors described a meta-learning-based healthcare recommender framework for learning and integrating few-shot learning, which improves generalization to new users and clinical tasks. Their model learns transferable interaction patterns from previous users which translates to cross-user generalization. Likewise, W. Sarah [29] used a dual-attention meta-learning approach with the aim of increasing adaptation to unseen patients in chronic care environment by using task-relevant and user-specific data. By contrastive pretraining, L. Wang and E.-P. Lim [30] constructed a zero-shot clinical recommendation pipeline which enabled the model to generate recommendations without requiring labeled examples from the target domain. Z. Kuang et al. [31], on the other hand, to extend generalization, proposed a contrastive user-embedding framework with a focus on how shared latent representations can be exploited in order to enhance recommendation abilities, especially for incomplete, or partially observed patient profiles. S. Shetty et al. [32] explored cross-domain recommendation by alignment of the latent space in response to domain sparsity, facilitating knowledge exchange from data-rich to data-poor medical domains. In a hybrid cold-start solution, S. Gupta et al. [33] presented a graph-based model with collaborative filtering and content-based methods, using meta-path analysis and propagation mechanisms to generate accurate recommendations with limited user data. These approaches, taken together, represent promising directions in overcoming cold-start and cross-user generalization limitations in HRS. However, more questions remain about how to adapt to new and dynamic patient feedback, temporal model recalibration, and performance stability in rare disease settings.

4.2. Fairness and Bias Mitigation in Healthcare Recommender Systems

As HRS become more widely implemented in clinical settings, there is an urgent need to ensure fairness across patient populations. Bias can be reflected in recommendations that disadvantage vulnerable groups when healthcare datasets are biased due to underrepresented groups, imbalanced clinical records, or systemic disparities. Consequently, significant attention has focused on fairness-aware algorithmic design. These approaches aim to identify, measure, and correct such biases in clinical recommender systems. L. H. Nazer et al. [34] applied an adaptive debiasing framework (i.e. collaborative filtering) to address the exposure and selection bias issues associated with HRS. Their approach also dynamically re-weights training samples to counteract their reliance on past data distortion to ensure fairness across user groups in recommending data. M. Sasseville et al. [35] proposed a fairness-aware matrix factorization of the personalized health information retrieval algorithm by including the bounds of group fairness in the optimization objective. Similarly, M. Haroon et al. [36] used fairness constraints on collaborative filtering networks to mitigate the differential exposure to healthcare items. In the field of oncology decision support, P. Rouzrokh and others [37] proposed a bias-aware clustering model to bring together patient sub-groups in an even-paced distribution with balanced demographic characteristics. The authors of [38] mitigated the age/gender biases by using adversarial learning techniques, which were proven to increase demographic generalization of latent representations. Furthermore, J. Yi et al. [39] have proposed DebiasRec, a fairness-sensitive recommender architecture that integrates causal and intervention representations for transparent and fair recommendation performance, particularly on sensitive datasets such as mental health. These contributions indicate that HRS is now being increasingly concerned with fairness and accountability. However, there are still several unresolved themes to address: reconciling predictive performance and the fairness constraints, the specification of standardized metrics for fairness relevant to medical application, and maintaining interpretability in increasingly complex recommendation pipelines. Fairness-aware mechanisms, however, are essential as the uptake of HRS in distinct clinical settings continues to unfold (e.g. where equitable and responsible decision support is essential to effective and trustworthy healthcare delivery).

4.3. Explainable AI (XAI) and Interpretability in HRS

For most healthcare recommender systems, explainability has become an important requirement due to the high-stakes nature of clinical decision-making. Deep CNN-based recommenders are often black-box models, which raises questions of transparency and reliability. Thus, explainable AI (XAI) approaches have emerged that can offer interpretable insights into recommendation processes. Grad-CAM is a visualization method that highlights spatial regions in medical images. These regions most influence CNN predictions. This approach helps clinicians validate automated diagnostic predictions and has been applied to image-driven aspects of patient classification. In addition to simple visual explanations, rule-based interpretability mechanisms have been under focus. Alzahrani et al. [41] proposed a neuro-fuzzy staging recommender that combines CNN-crafted features and fuzzy logic reasoning, to improve the transparency in chronic disease risk classification.

At the system level, frameworks like DECAF [42] integrate contextualized and historical patient data into clinical recommendations to produce textual and feature-level explanations for referral and ultimately achieve clinician acceptance. Bias-aware systems like DebiasRec [43] broaden explainability even more by embedding intervention-aware models that identify recommendations in both user behavior and fairness changes. M. Benleulmi et al. [44] conducted a comprehensive review on the application of XAI in deep learning based recommender systems, pointing out drawbacks, including no standards of model evaluation with regard to explanation analysis and the lack of causal reasoning integration. Additionally, Wang et al. [45] developed a knowledge-graph-based recommender paradigm which accounts for the explainability of the results by semantic multi-hop relations (e.g., symptom → disease → treatment), and made transparent the interpretable symbolic reasoning chains. These advances indicate a shift from opaque clinical recommender systems to transparent, multimodal, and semantically grounded models that closely follow evidence-based patterns in healthcare.

4.4. Reinforcement Learning in HRS

Reinforcement Learning RL has emerged as an effective option to facilitate sequential decisions in healthcare recommender systems. RL learns optimal policies through interaction with dynamic environments which is a significant advantage that makes it suitable for long-term patient management scenarios where health conditions change over time, unlike traditional supervised models. RL is frequently used for clinical pathway modeling, medication optimization, and also personalized behavioral interventions in HRS. M. Mehdi [23] developed an actor–critic-based RL model for personalized treatment recommendation on long-term glycemic control optimization guided by patient trajectories. Similarly, Zhang et al. [46] proposed a deep RL architecture for precision oncology, treating the selection of therapy as a sequential policy optimization problem driven by elements of genetic and clinically relevant patient profile. This novel approach of adaptive classification outperformed static methods in prediction of survival. Y. Yang and Y. Zhao [47] took advantage of contextual patient data and engagement signals as reinforcement feedback to apply Q-learning to mental health intervention recommendations. For chronic disease management, Y. Zhao et al. [48] included domain knowledge in reward design to improve clinical safety. D. J. Tan et al. [49] suggested a multi-agent RL model to describe clinical departments as agents interworking in shared environments to achieve optimized coordinated care output. Z. Guo et al. [50] integrated RL with causal inference to reduce treatment selection bias and increase policy generalization through counterfactual simulation. Moreover, M. B. Tariq and H. A. Habib [51] studied trust-aware reinforcement learning on mobile healthcare platforms, and integrated user trust dynamics with policy learning in order to improve recommendation adherence. Taken together, RL-based HRS illustrate a transition to adaptive, context-aware systems that can model long-term clinical impacts.

4.5. Causal Inference in Healthcare Recommender Systems

The use of causal inference is increasingly important in healthcare recommender systems for controlling for confounding variables, estimating treatment effects, and generalization of system across patient populations. Contrary to purely associative models, causal-aware recommenders try to approximate counterfactual scenarios, which enables more accurate and personalized clinical treatment. H. Hosseinmardi et al. [52] proposed the treatment effect-aware recommender that uses counterfactual estimators to model user-level heterogeneity in electronic health records and reduce recommendation bias. X. Wang [53] proposes a dual-branch neural counterfactual learning framework, predicting both factual and counterfactual outcomes based on patient trajectories, which is robust to sparse and imbalanced data. X. Ma et al. [54] developed domain-adaptive causal embeddings for aligning patient data across institutions through invariant causal embeddings. Z. Zhang et al. [55] applied instrumental variable-based causal modeling to distinguish treatment pathways as well as latent confounders in observational clinical environments. R. Qiu et al. [56] reported the recurrent counterfactual network integrating temporal attention and propensity scores to mitigate time-dependent confounders in longitudinal recommendations. K. Krauth et al. [57] addressed fairness in causal recommender systems by learning balanced representations to reduce demographic bias. S. Xu et al. [58] presented a hybrid causal reinforcement learning approach that connects short-term motivation with long-term efficacy of treatment. E. Cavenaghi et al. [59] presented a graph-based causal discovery technique for modeling clinically relevant connections between symptoms and therapies. Y. Ding et al. [60] combined the structure of causal models to the attention-based neural architectures to make the causal graphs more understandable. Finally, J. Liu et al. [61] presented an adversarial counterfactual inference architecture for balancing factual and counterfactual distributions to enhance personalization and robustness. These studies illustrate the increasing sophistication of causal methodologies in HRS, paving the way for safer, more accountable, and clinically grounded AI-driven recommendations.

4.6. Ethics in Healthcare Recommender Systems

HRSs introduce crucial ethical issues because they can affect care decisions and influence patient outcomes. Unlike general-purpose recommenders, HRSs must prioritize transparency, informed consent, responsible data governance, and patient autonomy [4], [62], [63], [64], [65]. Algorithmic opacity might erode physician–patient trust when recommendations lack interpretability within clinical workflows [62], while insufficient regulatory safeguards risk reinforcing structural biases or reducing patients to abstract statistical profiles [4], [63]. Fairness, accountability, privacy protection, and equitable access must therefore be systematically incorporated throughout system design and deployment [64], [65]. Ethical incorporation should cover everything within the HRS lifecycle, from data acquisition and model training to validation, deployment, and post-deployment monitoring. Human-in-the-loop oversight and explainable recommendation mechanisms further enhance clinical reliability and regulatory compliance.

4.7. Vision-Language Models VLM and Multimodal Transformers in Healthcare Recommender Systems

Recently developed VLM pretraining and multimodal transformers have improved the performance of healthcare recommender systems by facilitating the integration of visual medical data with textual clinical information. S. Raminedi et al. [13] presented a transformer-based encoder–decoder architecture to generate radiology reports based on imaging inputs, facilitating diagnostic explanation and case-based recommendations. In the same way, Q. Li et al. [66] found that general-purpose VLM models like BLIP-2 and OFA can also scale up to medical domains for robust zero-shot performance and report-grounded recommendations without task-specific retraining. MedViT [67], a medical-domain vision transformer, revealed a good combination of domain-specific priors and unified pretraining to enhance classification and captioning performance across medical imaging tasks. W. Huang et al. [68] proposed a key-semantic refinement in a vision-language model toward more meaningful information extracts from extended radiology reports, enabling downstream tasks like disease classification and phrase grounding. Schmidgall et al. in [69], proposed a Multimodal Vision-Language Architecture MAVL with medical-knowledge embeddings for patient-specific recommendation generation. MedUnifier [70] moved a step further into these realms to fuse bidirectional language modeling and visual generation as a unified framework, giving a new generation of narrative directly from imaging inputs. While this is a domain in the works, these multimodal architectures provide a promising next step in the scalable, semantically grounded architecture of healthcare recommender systems.

4.8. Evaluation on Real-World User Studies in Healthcare Recommender Systems

Evaluation of healthcare recommender systems in the real world, not only with offline benchmarks, is important to confirm their validation. Since a real-world deployment does not rely on retrospective evaluation alone, it draws upon contextual feedback from patients, clinicians, and caregivers, enabling a better understanding of usability, acceptance, and clinical impact. In order to explore the efficacy of HRS in real-world clinical practice, mHealth platforms have been designed in the current research. Additionally, [71] also explored a smartphone-based nutritional recommendation system through an initial pilot deployment and user input to identify major drivers to dietary adherence. Similarly, [72] combined wearable sensing with mobile HRS for lifestyle recommendation by using system performance, usability, and contextual responsiveness during natural use. P. Alves et al. [73] proposed a Dockerized microservices-based recommender, evaluated in a cluster randomized trial focused on prehypertension monitoring of long-term behavioral adherence and blood pressure. Similarly, [74] studied a recommendation for reducing obesity risk and the long-term engagement tracking of high-risk women. Interactive and human-in-the-loop approaches are examples of models that have been researched by others. D. Diyasena et al. [75] introduced a clinician-interactive recommender allowing experts to refine recommendations through clinicians' interaction (with higher trust and effective perception [75]). F. Gräßer et al. [76] examined a therapy recommender in mental health treatment sessions and observed the role of emotional context and trust. The elderly patients' medication adherence study has also been studied. In [77], clinical and elder user feedback on a mobile medication support system was explored, focusing on key points regarding interpretability, accessibility, and continued use. Randomized trials and longitudinal studies require considerable resource investment, but they also represent vital evidence for ecological validity. Emerging evaluation paradigms are increasingly bringing in passive sensing, clinician insights, and patient assessments to measure system-level impact. While this is a complex area of study, research must ensure that co-design principles, trust calibration, and adaptive feedback remain clinically relevant for HRS—combined with close alignment to ethical tenets for the future.

5. REVIEW OF STATE-OF-THE-ART APPROACHES

In this Section , we explore state-of-the-art systems that include CNNs for recommender systems in healthcare systems. Architectural design, dataset utilization, evaluation metrics, and recent trends in clinical recommendation research can all be captured to achieve a complete picture of the situation in clinical recommendation. The discussion highlights both technological discoveries and ongoing shortcomings, offering recommendations that could inform future advancements.

5.1. Research Trends and Frameworks

This Section describes specific lines of research and representative projects in which CNNs have been adopted in healthcare recommender applications for classification and recommendations. Rather than addressing common goals, we focus instead on advances in CNN architectures, multimodal integration approaches, and hybrid learning and data-driven models to achieve scalable and personalized health advice recommendation. Integration of deep learning and HRS has been growing rapidly since the era of data in the last few decades and more recently, as CNN architectures tend to capture structured representations in diverse medical data. Hybrid and data-rich models linking imaging to clinical databases and behavioral cues have increasingly defined user-centric healthcare delivery approaches. M. C. Comes et al. [78] used CNN with transfer learning on breast DCE-MRI imaging for early prediction of neoadjuvant chemotherapy response. Although largely diagnostic, their embedding-based process demonstrates how pathology-informed representations can contribute to personalized recommendation pipelines. B. Ihnaini et al. [9] focused on chronic disease management and used deep ensemble methods with CNNs to further enhance the accuracy and applicability of treatment recommendations for diabetes. Mahyari et al. [27] utilized a dynamic CNN-based recommendation system to provide personalized exercise coaching to users in a mobile health (mHealth) environment. Their model incorporates continuous learning that involves expertise feedback and provides adaptive customization. A. Nilla and E. B. Setiawan [79] pioneered a mixed CNN-based healthcare recommender integrating collaborative filtering and content processing techniques which can effectively integrate multiple modalities, enabling personalized health content targeting. N. K. Al-Qazzaz et al. [80] applied machine learning-based hybrid system for diagnosing autism spectrum disorder with pre-trained CNN for symptom severity and KNN classifiers for prediction for treating severe mental disturbance and how the recognition of patterns in a clinical scenario could be implemented to facilitate the individualized treatment planning of this condition. M. Arakeri et al. [81] introduced a CNN-centric collaborative filtering pipeline named DeepReco, which tackles mobile healthcare recommendation challenges. S. Bourhim et al. [82] introduced DGCF, an architecture based on GNNs with a large batch of convolutional operations in the health recommender system to enhance interaction modeling between user and item. Fallahi and

Mohammadzadeh [8] published CNN methods which were used on collaborative filtering and showed their remarkable ability to process sparsity and high-dimensional data problems. K. Karre [83] addressed the treatment prioritization techniques for children suffering from autism and demonstrated how to apply CNN-based models to different age groups and cognitive features. Iwendi et al. [84] proposed a CNN-based recommender for dietary personalization enabled by an Internet of Medical Things IoMT device, that showed how multimodal patient information and nutrition data can offer enhanced management for chronic disease. Ponselvakumar et al. [85] proposed a precision medical recommender system for the matching of patients with clinical trials that was proposed by multi-modal CNN inputs. Y. S. Cho and P. C. Hong [86] demonstrated the scalability of CNN-based diagnostic systems for malaria screening in low-resource settings, demonstrating potential wider impact of AI as a tool on social issues to promote fairness. Cai et al. [87] performed a scoping study on HRS from 2010 down to 2022 showing an advancement and integration between CNN and deep learning. Most recent work has been done by B. Subramanian et al. In [88], CNN features maps for the emotion recognition framework also on mobile mental health monitoring to detect mood states to intervene with personalized treatment. For that reason, within the context of CNN frameworks, recommendation mechanisms are being integrated more and more into other modalities such as visual analytics, behavioral inference, sensor streams, and structured clinical record. This has represented a transition toward broad, adaptive, and context-aware healthcare recommender ecosystems. Recent research trends and the representative frames in CNN-based medical recommender systems are summarized in Table 2.

Table 2: *Research Trends and Frameworks in CNN-Based Healthcare Recommender Systems*

Research Domain / Direction	**Fundamental Idea**	**Key References**	**Primary Contributions and Relevance to HRS**
CNN-Driven Imaging-Based Recommendation	Applying transfer learning with CNN architectures to extract clinically meaningful representations that inform treatment recommendations	Comes et al. [78]	Enabled early forecasting of chemotherapy response and introduced pathology-aware embeddings to enhance personalized therapeutic decision-making
Ensemble CNN Frameworks for Chronic Care	Designing multi-model CNN ensembles to strengthen recommendation reliability in chronic disease management	Ihnaini et al. [9]	Improved accuracy and robustness of diabetes-related recommendation systems through model fusion strategies
Adaptive and Continual CNN Recommenders	Embedding real-time feedback and incremental learning within CNN-based recommendation pipelines	Mahyari et al. [27]	Delivered dynamically personalized exercise guidance via expert-informed updates and continuous system adaptation
Hybrid CNN with Collaborative and Content Filtering	Merging CNN-derived representations with collaborative and content-based recommendation strategies	Nilla & Setiawan [79]	Achieved multimodal personalization of healthcare content through an extensible and hybrid recommendation framework
CNN–KNN Integrated Severity Modeling	Combining deep CNN feature extraction with KNN classification for symptom intensity evaluation	Al-Qazzaz et al. [80]	Facilitated individualized ASD therapy planning by integrating deep pattern learning with traditional classification techniques
CNN-Augmented Collaborative Filtering in mHealth	Embedding CNN-extracted features into collaborative filtering architectures for mobile health applications	Arakeri et al. [81] (DeepReco)	Developed a mobile-centric recommender leveraging CNN-enhanced user–item representations
CNN–GNN Fusion for Interaction Representation	Integrating convolutional mechanisms within graph neural network structures to refine user–item modeling	Bourhim et al. [82] (DGCF)	Strengthened structural learning and relational modeling in healthcare recommender systems through graph–CNN synergy
CNN-Based Collaborative Filtering under Sparsity	Utilizing convolutional architectures to address sparsity and behavioral variability in healthcare recommendation data	Fallahi & Mohammadzadeh [8]	Demonstrated both empirical and theoretical robustness of CNN-based CF models in sparse clinical datasets
CNN Applications in ASD Treatment Prioritization	Employing CNN-driven cognitive and severity assessment to guide therapy ranking decisions	Karre [83]	Showed cross-demographic adaptability of CNN-based systems for personalized autism treatment prioritization
IoMT-Enabled CNN Dietary Recommendation	Integrating IoMT data streams with CNN-based decision-support models for nutritional planning	Iwendi et al. [84]	Provided tailored dietary recommendations for chronic disease patients by using multimodal IoMT inputs

Multimodal CNN Architectures for Precision Medicine	Exploiting heterogeneous CNN inputs to support therapy selection and clinical trial alignment	Ponselvakumar et al. [85]	Enabled personalized clinical trial matching through multimodal medical data integration
CNN-Based Diagnostic Support in Low-Resource Contexts	Deploying CNN-powered diagnostic systems in underserved and resource-limited regions	Cho & Hong [86]	Supported scalable malaria screening and promoted equitable healthcare recommendations in constrained environments
Evolutionary Trends of CNN in HRS	Analyzing the progression of CNN integration within healthcare recommender research	Cai et al. [87]	Documented sustained expansion and methodological advancement of CNN-based HRS approaches (2010–2022)
Emotion-Aware CNN Recommenders for Mental Health	Leveraging CNN feature representations to infer emotional states for adaptive mental health recommendations	Subramanian et al. [88]	Developed a mood-sensitive mobile recommender supporting psychological monitoring and personalized intervention

5.2. Representative Dataset

Diverse datasets are essential for efficient CNN-based healthcare recommender systems. These datasets are necessary to validate model generalizability. They also ensure stable performance across diverse populations and clinical environments. CNNs trained on medical imaging datasets for disease detection and feature extraction have predominantly been used in healthcare applications. Well-known datasets include ChestX-ray14 and CheXpert, which are widely used for thoracic disease screening. MIMIC-CXR combines radiographic images with detailed clinical reports to improve diagnostic interpretability and improve the understanding of disease context [10]. The COVIDx dataset [89], which was widely used because of the COVID-19 pandemic, was essential for CNN model training to detect viral pneumonia from chest radiographs. Aside from imaging, such clinical datasets offer some extra added value to health recommender development. MIMIC-III can provide extensive patient attributes, such as laboratory measurements, medications, ICU admission, and physician details, to optimize training the patient-centered recommendation models [90]. Nutritional, behavioral and demographic characteristics are supplied by the National Health and Nutrition Examination Survey NHANES to fit as basis for lifestyle focused health recommender systems [91]. In the same vein, the UCI Diabetic dataset has provided a good basis for prediction and optimization of chronic disease prediction and patient care through CNN-based algorithms [92]. Wearable and physiological signal datasets are extending the application of CNN-based recommender systems, such as WESAD that presents also unique features. For instance, WESAD enables the modeling of personalized mental health intervention for patients based on multimodal physiological signals of stress and affect state, enabling personalization of clinical data. Time-series based data such as those collected in MIT-BIH have developed into ECG recordings, including a good source of data for cardiovascular monitoring and adaptive health feedback systems, thus enhancing the understanding of monitoring system. In emotion-based health applications, AffectNet and custom webcam-based emotion datasets are examples of datasets used to achieve real-time mood recognition and therapeutic support by using CNN in-situ feature extraction [94]. While these datasets are helpful, they have some limitations. Most of these datasets have limited demographic heterogeneity and are subjected to bias predictions that might not be generalized to other patient populations. Others are unimodal, offering only a single modality of the data (i.e. images without the addition of corresponding textual data or physiological information), limiting the potential to evolve a hybrid model [95]. Furthermore, privacy laws limit user access to detailed and longitudinal patient records in combination with data inpatient reports, which results in the gap between experimental studies and real-world clinical implementation [96]. Further construction of CNN-based healthcare recommender systems will require further developing and ethical dissemination of multimodal, demographically representative, longitudinal data to better represent the real-world clinical complexity.

5.3. Performance Metrics

The evaluation of CNN recommender systems in healthcare contexts is typically guided by measurement techniques for the identified quality goals, namely classification accuracy, prediction reliability, ranking effectiveness, and user-centered evaluation. They are utilized to assess the robustness of deep learning architectures and the practical usefulness of resulting recommender systems in clinical settings [45]. For classification-based tasks such as medical imaging and disease detection, the evaluation metrics used include accuracy, precision, recall (sensitivity), specificity, F1-score, and area under the receiver operating characteristic curve (AUC-ROC) [89], [97]:

Accuracy represents the proportion of correctly classified instances among all predictions. The computation of accuracy is given in equation (7):

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \qquad (7)$$

Precision measures the proportion of true positive predictions among all predicted positive cases. The formulation is presented in equation (8):

$$Precision = \frac{TP}{TP + FP} \quad (8)$$

Recall (Sensitivity) quantifies the proportion of actual positive cases that are correctly identified by the model. It is calculated as shown in equation (9):

$$Recall = \frac{TP}{TP + FN} \quad (9)$$

The F1-Score represents the harmonic mean of precision and recall, providing a balanced measure when class distributions are uneven. The calculation is given in equation (10):

$$F1 - Score = \frac{2 * Precision * Recall}{Precision + Recall} \quad (10)$$

Specificity SP reflects the proportion of true negative instances correctly identified. It is computed as shown in equation (11):

$$Specificity = \frac{TN}{TN + FP} \quad (11)$$

Area Under the Receiver Operating Characteristic Curve ROC-AUC evaluates the model's ability to distinguish between classes across varying decision thresholds [33].

In recommendation-oriented tasks, where personalized ranking and prediction accuracy are central, additional evaluation metrics are applied.

Mean Absolute Error MAE measures the average absolute difference between predicted and actual values. The formulation is provided in equation (12):

$$MAE = \frac{1}{n} \sum |y_i - \hat{y}_i| \quad (12)$$

Root Mean Squared Error RMSE calculates the square root of the average squared prediction errors, emphasizing larger deviations. It is defined in equation (13):

$$RMSE = \sqrt{\frac{1}{n} \sum (y_i - \hat{y}_i)^2} \quad (13)$$

Mean Reciprocal Rank MRR assesses ranking quality by evaluating the position of the first relevant item in a recommendation list. The calculation is presented in equation (14):

$$MRR = \frac{1}{|Q|} \sum \frac{1}{rank_i} \quad (14)$$

In practical deployments, systems like DeepReco [45] not only focus on predictive quality in real-world deployments, but also on functional performance like low latency, interpretability, and adapting to the need of its users in a real-time application. Given that health recommender systems now more and more tend to deal with the processing of sensitive patient data, fairness-aware evaluation frameworks like those defined by L. Bruni et al. [98]—and transparency mechanisms such as Grad-CAM [18] and SHAP [99], are increasingly important for CNN data flow-based performance evaluation.

5.4. Limitations in Current Studies

Although there are positive outcomes for healthcare recommender systems with CNN-based models, a number of limitations can be identified in the literature. Large parts of the literature rely on small samples and restricted population and geographic boundaries. These drawbacks result in models that perform well in controlled environments but do not generalize well to real-world settings. Skewed representation across demographics, e.g. age groups, race, and socio-demographic diversity introduces bias in recommendations that are particularly challenging within clinical decision-support spaces, and is often detrimental [21]. Second and most frequently, this issue relates to the interpretability of CNN architectures. While CNNs are capable of learning complex and potentially very informative feature representations from medical data sources, the underlying decision rules for the algorithms are generally opaque. In a domain such as healthcare, where accountability and transparency are key features, such 'black-box' nature of quality makes clinical reliability problematic. Methods such as Grad-CAM and SHAP have been introduced from an explainability perspective. However, it has not been widely applied in real healthcare recommender systems.

For nearly everything in the deployed or empirical implementations today, most still do not have fully interpretable outputs which reduces trust among clinicians and end-users, and limits their practical use [40], [42].
Furthermore, a number of the proposed systems have not undergone real-time validation in operational clinical settings. Most evaluation studies are retrospective. They are usually conducted in controlled laboratory settings. Only a few are tested in real-world and diverse environments [85]. Because of this, it is difficult to move from successful lab results to real clinical applications. For example, the integration of Electronic Health Record EHR systems is often not well planned. Compliance with regulatory requirements is also weak in many cases. The implementation of EHR in daily clinical practice lacks extensive studies and clear documentation [81]. Data privacy and security is the major challenge appearing in applying CNN-based recommendation systems on a large scale in healthcare. These systems often use sensitive patient data; therefore, they ought to follow strict regulatory frameworks such as HIPAA and GDPR. Only a very small number of studies take the privacy-preserving strategies into account, such as federated learning or differential privacy. In certain instances, the techniques used for anonymization did not meet clinical standards [100]. Performance evaluation did not generally consider fairness and ethical considerations. In reality, models may inadvertently benefit majority groups while oversimplifying minority subgroups, resulting in healthcare inequities and disparities in access to and attainment of healthcare [4], [62]. There has been minimal progress on the formal fairness-aware evaluation criteria, bias auditing mechanisms, and ethical design principles in the development and evaluation of the system. Addressing these issues is essential as a foundational step for future research endeavors.

5.5. Clinical Translation, Regulatory Considerations, and Future Benchmarks

Despite significant progress in CNN-based healthcare recommender systems, several challenges still limit their use in dynamic clinical settings. Such challenges include weak integration with existing hospital systems, the quality of the data that is available from different hospitals, and the lack of standard clinical validation protocols. Furthermore, the reliance upon retrospective and experimental datasets limits the reliability of these systems in clinical practice settings. To successfully implement such systems into clinical practice, the development of user-friendly interfaces that match the decision-making processes of clinicians is essential - an area that has been insufficiently explored so far in existing studies.

Considering the regulatory and ethical implications of deploying AI-based healthcare recommender systems, there exist a variety of regulations that must be considered prior to the deployment of such systems. Regulations related to patient data protection, such as HIPAA and GDPR, must be considered, as must regulations related to the requirement such as transparency, explainability, and auditability of AI systems prior to their approval for use in healthcare environments. The lack of standards related to regulating AI systems that can learn from and adapt to new information over time presents challenges to their deployment and approval for use.

Therefore, future research in this area should focus upon developing standards for evaluating CNN-based healthcare recommender systems. Such evaluations can utilize publicly available datasets of multimodal health information, metrics beyond accuracy, such as fairness, robustness, and interpretability in measuring the performance of the systems, and can perform large prospective studies to confirm the performance of these systems. These efforts will be critical to the safe and effective deployment of AI-based healthcare recommender systems.

6. SYSTEM-LEVEL ENHANCEMENTS AND DEPLOYMENT CONSIDERATIONS

6.1. Federated Learning and Privacy-Preserving Healthcare Recommender Systems

Keeping patient privacy and efficient predictive accuracy is a strong challenge for HRS because of distributed data of heterogeneous systems. Whereas Federated Learning (FL) has emerged as a key approach to avoid this problem, with model training being done in collaboration rather than across institutions where raw data are shared. With decentralized learning that has no need to communicate with data centers of centralized architectures, FL avoids the need for expensive communications and can still provide best privacy. It enables generalization in tightly controlled domains such as medical diagnosis, treatment planning and behavioral health monitoring. M. Adnan et al. [101], a federated multitask deep learning approach based on federated data, has been proposed for early breast cancer prediction in a federated world. The approach treats patients from multiple hospitals, by optimizing the local model to one condition from a certain institution, and maintaining it in a one condition globally consistent setting. It reduces prediction error between patients from different population groups and protects the privacy of users. Because of data diversity, S. Che et al. [102], proposed FedHealthRec as a personalized federated collaborative filtering method. It adapted local recommendation parameters to user behavior, including the privacy-preserving knowledge distillation mechanism that also resolves cold-start concerns not to publish any private data. M. H. Nguyen et al. [103] presented a federated medical dialogue application for using a medical dialogue system to manage chronic diseases. Transformer-based language models are combined with on-device models for real-time, privacy-aware medical provision via centralized policy optimization. However, model distillation methods are essential in order to support performance under both low computing and unstable network environments like on the mobile healthcare domain. S. D. N et al. [104] proposed FedCure, which can perform a federated recommender of globally shared knowledge based on locally-personalized embeddings. Using device heterogeneity as a parameter, they provide recommendations based on customer performance and participation frequency. In an appropriate direction, E. Mantey et al. [105] introduced a federated recommender framework based on Hyperledger Fabric by using a blockchain, and this has enabled secure data-based storage for consumers in case their mobile health platform, on-demand, or fully involved users also provided explainable choices from the mobile health systems. Building on privacy-prescribing designs to further explore this, S. Gupta et al. [33] introduced FedGraphRS, a federated recommender network based on graph representations which applies homomorphic encryption and graph neural networks that exploits homomorphic encryption algorithms to simulate user-item interaction between decentralized individual institutions. E. A. Mantey et al. [106] built from this paradigm integrating LSTM and GRU-based temporal modeling into a blockchain-governed infrastructure for chronic disease management. They have incorporated EHRs into a distributed ledger and write post-predictive follow-up recommendations of the longitudinal health trajectory. Altogether, these papers show how federated learning is redefining HRS as a decentralized, privacy-preserving ecosystem. Nonetheless, non-IID data distribution issues, model drift and fairness across all participating clients are still unresolved. Balancing personalization with global data generalization is an important research question for federated healthcare recommender systems.

6.2. Edge Deployment and IoMT in Healthcare Recommender Systems

The development of the Internet of Medical Things IoMT and edge computing has changed how recommender systems, particularly healthcare recommender systems, are developed to provide real-time decision support which is context-aware and privacy-preserving. Legacy cloud-based algorithms have suffered from limitations and are risky with data lags and potential of transmission when employed in mission-critical medical applications. To these limitations, techniques such as federated learning and heterogeneous edge optimization have been integrated into the recommender pipelines and are applied in practice. A. S. Fathima et al. [107] introduced a federated edge–cloud framework to provide dynamic offloading of medical resource recommendations to 5G IoMT devices. They leverage the edge preprocessing together with blockchain priority to scale the scheduling. M. H. Nguyen et al. [103] proposed a heterogeneous BlockNet-based federated architecture for diagnostic recommendation to develop new neural networks with varying computational capabilities for edge devices. This enhances scalability, efficiency, and privacy in distributed healthcare environments. Sensor-based integration further reinforces the need for emotion-aware healthcare processes. B. Subramanian et al. [88] proposed a digital twin-based deep learning framework to process real-time emotion recognition with webcam inputs at commodity edge devices. This scalable design allows measuring psychological state and providing timely recommendation delivery in mobile settings. On-device federated learning continues to grow the ability of personalization at the edge. A. Alqhatani and S. B. Khan [108] have introduced an IoT-based Half-and-Half Attention Enhanced Deep Collaborative Transformer integrated with federated learning and show increased recommendation accuracy and latency for extensive IoMT applications. Previous research on lightweight residual networks has investigated efficient representation learning methods that can be used for constrained edge-deployed vision modules used in healthcare recommender systems [20]. Other disease-specific IoMT recommender solutions have also been developed. The model of chronic disease management proposed in [21] exemplifies the application of similarity modeling by using EHR data in relation to IoT-enriched healthcare environments for facilitating proper diagnosis and treatment planning. To further extend recommendation logic, structural modeling approaches have been proposed. S. Bourhim et al. [82] proposed a graph neural network-based deep collaborative recommender that captures complex relationships between users and items, and shows how structural learning improves robustness for healthcare recommender architectures. Personalized treatment recommendation is yet another layer of critical deployment. In the course of the aforementioned recommendations, Alam [22] presented the Personalized Multimodal Treatment Response System PMTRS to ingest imaging, clinical and demographic data as intermediaries for prescription-level recommendations. Model partitioning methods, like [109], target distributed-optimized deep learning architectures in a wide range of edge devices, and may be used, like their predecessors, as a base for hierarchical recommender pipeline designs in healthcare. S. D. N et al. [104] discovered that heterogeneous-aware federated learning improves the accuracy and latency of IoMT and it could also be used in diabetes monitoring, retinopathy screening, maternal health assessment, and managing patients remotely. Taken together, these results underscore a paradigm shift from centralized healthcare recommendation environments to distributed, scalable, and resilient edge-AI landscapes. Despite that, we must overcome some obstacles, like energy constraints, federated model drift, and cross-device variability.

7. CHALLENGES AND OPEN RESEARCH ISSUES

While CNNs have been gradually used in HRSs , many unsolved problems hinder their widespread clinical use. These issues consist of technical limitations, privacy and fairness, and operational restrictions. All these are still in an active state of research.

7.1. Data Heterogeneity and Non-IID Distributions

Healthcare data are inherently heterogeneous, given that patient demographics, imaging devices, clinical practices, and institutional protocols vary greatly. In federated settings such as non–Independent and Identically Distributed (non-IID) data distribution means unstable gradient updates and slow convergence. Recent federated recommender systems attempt to overcome these issues. Personalized FL-based recommender systems using meta-learning frameworks, such as REPTILE-style first-order adaptation, have been proven to enhance personalization over decentralized data (such as PrivRec and DP-PrivRec applications). Graph-based techniques like FedGraphRS [33] utilize structural relationships between users and items to improve generalization to cross-client contexts. Hybrid models, such as FedCure [104], which combine design specialization features at the device level and adaptive aggregation, can achieve greater personalization, but are susceptible to unreliable clients and inconsistent participation. In conclusion, extreme data imbalance, interaction sparsity, and cross-institution variability are ongoing challenges with federated healthcare recommendation.

7.2. Personalization vs. Generalization Trade-offs

Effective healthcare recommendations require personalization that accounts for comorbidities, historical behavior, lifestyle factors, and long-term patient trends. However, increasing personalization can lead to reduced global model generalizability.

High-heterogeneity federated architectures like BlockNets [103] show that modular specialization will increase local relevance but may damage broad-based robustness. Likewise, FedCure [104] is able to enhance localized performance at the individual level by local tuning, yet is overly sensitive to frequency variances in participation. Dynamic contextual federated recommenders with temporal and behavioral modeling, such as those described in [110], need ways to adapt personalized experiences to the patient. Future system design needs to balance patient-specific relevance and stable generalization across heterogeneous clinical populations.

7.3. Multimodal Fusion and Temporal Misalignment

Current modern healthcare recommender techniques are incorporating multimodal data sources such as medical images, clinical notes, physiological signals, and demographic variables. The image–text alignment with high-resolution image and text is illustrated from transformer-based fusion architectures like ViGPT2 [111] by co-attention. Cross-modal transformers [112], [113] show the effective fusion between textual and visual features. However , multimodal CNN-based HRS suffer from multiple operational issues such as:

- Vitals and imaging with mismatched temporal sampling.
- Missing, or noisy modalities.
- Late-arriving sensor data.
- Latency-dependent cross-attention calculations.

These challenges are accentuated by real-time systems or longitudinal patient monitoring scenarios. Designing more robust architectures to cope with multimodal uncertainty and temporal variations is still an open question in research.

7.4. Explainability and Clinical Trust

Interpretability is very important in clinical decision-making. To provide insights and textual rationales, alternative applications such as Grad-CAM [18] as well as attention-based visualization in vision–language models [66] shed light on important segments. Graph-based explanation techniques, such as node and edge attribution techniques [45], increase the interpretability of graph-structured recommender systems. These mechanisms still may, however, not be up to clinical workflow standards. Co-design frameworks [41] proposed that clinicians should plan to develop nested explanations for model-level evidence (e.g., emphasized imaging features, or the presence of symptom clusters) and the explicit statements that can explain treatment recommendations. Systematic human-centered validation of explanation quality and calibration of trust is the less explored area and is pivotal to its broader utilization.

7.5. Fairness, Bias, and Data Imbalance

These biases may result from demographics, infrequent diseases, and resource-poor areas. While several approaches, including LDW-CNN [20], attempt to address class imbalance by using optimization through feature discrimination approaches, demographic fairness and cross-institutional equity have been comparatively less studied. Despite many efforts in developing domain adaptation strategies [114], existing methods to address distributional effects in medical CNN-processing have not yet included established fairness metrics for clinical risk prediction. Preventive bias auditing, calibrated prediction frameworks, and demographically-aware training methods are needed in healthcare recommender systems for fairness.

7.6. Edge Deployment and Resource Constraints

Deployment of CNN-based healthcare recommender models to the edge brings the burden of energy consumption, memory, and computation. In addition to an efficient attention mechanism, MedViT [67] illustrated various practical limits related to transformer inference on bandwidth-constrained devices. [115] described partitioned and distributed inference techniques that can distribute components of models onto varied edge devices. Federated IoMT systems including PrivRec [33], FedCure [104], and IoMT-aware architectures [116] are attractive for these decentralized learning systems. Blockchain coordination frameworks, particularly MedShare [100], improve the resilience of distributed scheduling and secure data exchange. However, ensuring continuity of operation through poor connectivity, hardware limitations, and variable participation of clients remains a challenge in this area.

7.7. Evaluation Gaps and Real-World Validation

Although quite a few similar frameworks also show performance enhancement, much of this work relies on synthetic/simulated federated learning situations. For instance, PrivRec [33] and FedCure [104] validated the performance of FL simulations under controlled conditions without long-term deployment and clinician-in-the-loop validation. Overall, deployment-oriented studies [27] emphasized the importance of institutional integration, real-time constraint evaluation, and end-user evaluation. Currently, there are user studies, multi-center trials, and comparisons with what is recommended by physicians for HRS in a limited number of cases. The development and standardization of clinical evaluation protocols combined with cross-institutional collaboration can lead to the advancement from research prototypes to clinical application.

In summary, both the CNN-based and multimodal-based healthcare recommender systems have undergone tremendous levels of innovation, but major challenges are present when dealing with heterogeneous data, reconciling personalization with generalization, robust multimodal fusion between models, fairness and interpretability, edge deployment, and real validation in practicalities. Addressing these challenges is going to require coordination among machine learning, medical informatics, human–computer interaction, and clinical workflow engineering. The general scope of some of the key challenges highlighted in this Section is summarized in Table 3.

Table 3: *Key Challenges in CNN-Based Healthcare Recommender Systems*

Challenge Domain	Underlying Issue	Illustrative Studies / Frameworks	Remaining Gaps and Research Challenges
Heterogeneous Data and Non-IID Learning	Variability across institutions, devices, and patient demographics leads to non-IID data distributions, causing unstable optimization and slow convergence in federated settings	PrivRec, DP-PrivRec, REPTILE-based meta-learning, FedGraphRS [33], FedCure [104]	Ongoing issues include severe data imbalance, sparse user–item interactions, unreliable or intermittently active clients, uneven participation rates, and cross-site variability
Personalization–Generalization Balance	Fine-grained personalization enhances patient-specific accuracy, but can weaken global model stability in distributed federated architectures	BlockNets [103], FedCure [104], Temporal/behavior-aware FL models, Context-Aware FL Recommenders [110]	Persistent tension between individualized optimization and cross-population robustness; vulnerability to participation irregularities; potential overfitting to specific patient profiles
Multimodal Integration and Temporal Desynchronization	Heterogeneous modalities (imaging, text, physiological signals, sensors) often arrive asynchronously, generating alignment challenges and computational overhead	ViGPT2 [111], Co-attention mechanisms, Cross-modal transformers [112], [113]	Temporal misalignment, incomplete or delayed modalities, latency-sensitive attention operations, and decreased reliability in real-time or longitudinal monitoring environments
Interpretability and Clinical Confidence	Current explanation mechanisms often lack clinically layered and actionable insights suitable for high-risk medical decisions	Grad-CAM [18], Vision–Language Model interpretability [66], Graph-based explainers [45], Co-design methodologies [41]	Limited human-centered evaluation, weak integration between visual evidence and clinical reasoning, and insufficient calibration of trust for deployment in clinical workflows
Fairness and Demographic Bias	Skewed representation of demographic groups and rare conditions contributes to inequitable prediction outcomes	LDW-CNN [20], Domain adaptation strategies [114]	Lack of standardized fairness metrics tailored to healthcare; continued demographic disparities; insufficient development of calibrated and equity-aware learning frameworks
Edge Implementation and Resource Limitations	Deployment of CNN / Transformer-based HRS on constrained edge devices introduces limitations in computation, memory, and energy consumption	MedViT [67], Distributed inference approaches [115], PrivRec [33], FedCure [104], IoMT-centric systems [116], MedShare [100]	Latency sensitivity, energy constraints, unstable connectivity, hardware heterogeneity, inconsistent client availability, and operational instability in remote or low-resource environments
Evaluation Deficiencies and Clinical Validation	Heavy dependence on simulated federated environments rather than operational clinical deployments	ZoKrates, Scale, PrivRec [33], FedCure [104], Deployment-focused analysis in [27]	Scarcity of clinician-in-the-loop studies, limited real-world multi-center trials, insufficient longitudinal validation, and absence of standardized cross-institution evaluation frameworks

8. CONCLUSIONS

In this paper, we studied the convergence between CNNs and HRSs , specifically, and recent advancements in both fields. Thus, the analysis showed that CNN-based approaches enable healthcare systems to obtain more meaningful representations based on complex visual and multimodal data, thus enabling them to contribute to decisions on individual data better. Finally, when deployed in recommendation architectures, these architectures shift patient-centered healthcare from the diagnostic perspective to holistic approaches. Despite substantial advances, there is a wide array of hurdles to overcome when considering large-scale clinical adoption. Data heterogeneity, interpretability, fairness, privacy preservation, and real-world validation are significant challenges to overcome, and they continue to be major challenges. Emerging promising areas like federated learning, explainable AI, causal modeling, and reinforcement learning are still to be developed and more interdisciplinary studies are required to incorporate these in a stable, clinically validated system. The aspect of ethical governance and regulatory compliance will continue to be of critical focus in the design, implementation, and evaluation stages in the area of health recommender systems. While this review of the literature has provided an overview of current state-of-the-art, architectural work, and open issues that have been made possible in the field of CNN-based multimodal learning for health care recommender systems, the future of this technology will benefit from the development of multimodal health datasets, quality criteria for their assessments, and a better alignment of technological advancements to the needs of patients in healthcare systems.

Beyond the development of improved methodologies, the current frameworks for health care recommender systems based upon CNN-based multimodal approaches will need to be tested within the real world. Many of the current studies of these frameworks have used retrospective datasets or other controlled studies to gather their results. These study designs do not guarantee that the findings can be accurately applied to the real world of healthcare systems and their patients. Therefore, future studies should focus upon verifying these findings over time, from different institutions, and across a variety of demographic groups to ensure their reproducibility and generalizability to other populations in healthcare.

The development of health recommender systems indicates the need for the integration of these systems into the various applications within healthcare systems. However, these technological developments are only one part of a much larger issue regarding the development and application of these technologies. Furthermore, the solutions to these recognized issues will require the collaboration of various groups within the healthcare systems, such as data scientists, healthcare providers, biomedical engineers, and policymakers. Through the appropriate implementation of health care recommender systems as well as the collaboration of those groups within healthcare systems, the technologies have the potential to play a vital role in the future of health care.

REFERENCES

[1] Z. Zhang and M. Wang, "Convolutional Neural Network with Convolutional Block Attention Module for Finger Vein Recognition," Feb. 14, 2022, *arXiv*: arXiv:2202.06673. doi: 10.48550 / arXiv.2202.06673.

[2] K. Aljohani and T. Turki, "Automatic Classification of Melanoma Skin Cancer with Deep Convolutional Neural Networks," *AI*, vol. 3, no. 2, Art. no. 2, June 2022, doi: 10.3390 / ai3020029.

[3] T. N. T. Tran, A. Felfernig, C. Trattner, and A. Holzinger, "Recommender systems in the healthcare domain: state-of-the-art and research issues," *J. Intell. Inf. Syst.*, vol. 57, no. 1, pp. 171–201, Aug. 2021, doi: 10.1007 / s10844-020-00633-6.

[4] S. Tiribelli, "The AI ethics principle of autonomy in health recommender systems," 2023, *Università degli studi di Sassari*. doi: 10.14275 / 2465-2334/20230.TIR.

[5] G. Alotaibi, M. Awawdeh, F. F. Farook, M. Aljohani, R. M. Aldhafiri, and M. Aldhoayan, "Artificial Intelligence AI diagnostic tools: utilizing a Convolutional Neural Network CNN to assess periodontal bone level radiographically—a retrospective study," *BMC Oral Health*, vol. 22, no. 1, p. 399, Sept. 2022, doi: 10.1186/s12903-022-02436-3.

[6] A. Del Rio, J. Jimenez, R. Medina-García, C. Lozano-Hernández, F. Alvarez, and J. Serrano, "Improving Quality of Life in Chronic Patients: A Pilot Study on the Effectiveness of a Health Recommender System and Its Usability," *Appl. Sci.*, vol. 13, no. 10, p. 5850, May 2023, doi: 10.3390 / app13105850.

[7] R. Alanazi, "Identification and Prediction of Chronic Diseases Using Machine Learning Approach," *J. Healthc. Eng.*, vol. 2022, no. 1, p. 2826127, 2022, doi: 10.1155 / 2022 / 2826127.

[8] A. F. RahmatAbadi and J. Mohammadzadeh, "Leveraging Deep Learning Techniques on Collaborative Filtering Recommender Systems," *J. Appl. Res. Ind. Eng.*, no. Online First, Nov. 2021, doi: 10.22105/jarie.2021.275620.1264.

[9] B. Ihnaini *et al.*, "A Smart Healthcare Recommendation System for Multidisciplinary Diabetes Patients with Data Fusion Based on Deep Ensemble Learning," *Comput. Intell. Neurosci.*, vol. 2021, no. 1, p. 4243700, Jan. 2021, doi: 10.1155 / 2021 / 4243700.

[10] J. Irvin *et al.*, "CheXpert: A Large Chest Radiograph Dataset with Uncertainty Labels and Expert Comparison," *Proc. AAAI Conf. Artif. Intell.*, vol. 33, no. 01, Art. no. 01, July 2019, doi: 10.1609 / aaai.v33i01.3301590.

[11] C. Suman, S. Saha, A. Gupta, S. K. Pandey, and P. Bhattacharyya, "A multi-modal personality prediction system," *Knowl.-Based Syst.*, vol. 236, p. 107715, Jan. 2022, doi: 10.1016 / j.knosys.2021.107715.

[12] H. Alshazly, C. Linse, E. Barth, and T. Martinetz, "Ensembles of Deep Learning Models and Transfer Learning for Ear Recognition," *Sensors*, vol. 19, no. 19, p. 4139, Sept. 2019, doi: 10.3390 / s19194139.

[13] S. Raminedi, S. Shridevi, and D. Won, “Multi-modal transformer architecture for medical image analysis and automated report generation,” *Sci. Rep.*, vol. 14, no. 1, p. 19281, Aug. 2024, doi: 10.1038 / s41598-024-69981-5.

[14] M. Shafiq and Z. Gu, “Deep Residual Learning for Image Recognition: A Survey,” *Appl. Sci.*, vol. 12, no. 18, Art. no. 18, Jan. 2022, doi: 10.3390 / app12188972.

[15] L. Shen, L. R. Margolies, J. H. Rothstein, E. Fluder, R. McBride, and W. Sieh, “Deep Learning to Improve Breast Cancer Detection on Screening Mammography,” *Sci. Rep.*, vol. 9, no. 1, p. 12495, Aug. 2019, doi: 10.1038/s41598-019-48995-4.

[16] C. A. Jayawardena, N. Wedasingha, N. Kolambage, and S. Perera, “Pneumonia Detection and Lung Disease Assessment from Chest X-rays: Developing A Diagnostic Support System,” Jan. 2025, doi: 10.54389/MKYM4075.

[17] U. Manor, V. Thiyagarajan, K. Harris, and A. Sheridan, “Sparse Annotation is Sufficient for Bootstrapping Dense Neuron Segmentation,” Nov. 14, 2024, *In Review*. doi: 10.21203 / rs.3.rs-5339143/v1.

[18] 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,” *Int. J. Comput. Vis.*, vol. 128, no. 2, pp. 336–359, Feb. 2020, doi: 10.1007 / s11263-019-01228-7.

[19] P. Chinnasamy, W.-K. Wong, A. A. Raja, O. I. Khalaf, A. Kiran, and J. C. Babu, “Health Recommendation System using Deep Learning-based Collaborative Filtering,” *Heliyon*, vol. 9, no. 12, Dec. 2023, doi: 10.1016/j.heliyon.2023.e22844.

[20] B. R. Reddy and R. L. Kumar, “HCLNet: A Hybrid CNN-LSTM Model for Enhanced Healthcare Product Classification for Recommendation,” *Eng. Technol. Appl. Sci. Res.*, vol. 14, no. 6, pp. 18935–18940, Dec. 2024, doi: 10.48084/etasr.8946.

[21] Y. A. Nanehkaran *et al.*, “Diagnosis of Chronic Diseases Based on Patients’ Health Records in IoT Healthcare Using the Recommender System,” *Wirel. Commun. Mob. Comput.*, vol. 2022, no. 1, p. 5663001, Jan. 2022, doi: 10.1155 / 2022 / 5663001.

[22] S. Alam, “Personalized Multimodal Treatment Response System PMTRS Framework for Personalized Healthcare”.

[23] M. Mehdi, “Personalized Recommendation Using Reinforcement Learning”.

[24] S. Zhang *et al.*, “Innovative Material-Based Wearable Non-Invasive Electrochemical Sweat Sensors towards Biomedical Applications,” *Nanomaterials*, vol. 14, no. 10, Art. no. 10, Jan. 2024, doi: 10.3390 / nano14100857.

[25] Y. Yan, K. Zhao, J. Cao, and H. Ma, “Prediction research of cervical cancer clinical events based on recurrent neural network,” *Procedia Comput. Sci.*, vol. 183, pp. 221–229, Jan. 2021, doi: 10.1016 / j.procs.2021.02.052.

[26] Q. Zhang, J. Lu, and Y. Jin, “Artificial intelligence in recommender systems,” *Complex Intell. Syst.*, vol. 7, no. 1, pp. 439–457, Feb. 2021, doi: 10.1007 / s40747-020-00212-w.

[27] A. Mahyari, P. Pirolli, and J. A. LeBlanc, “Real-Time Learning from an Expert in Deep Recommendation Systems with Application to mHealth for Physical Exercises,” *IEEE J. Biomed. Health Inform.*, vol. 26, no. 8, pp. 4281–4290, Aug. 2022, doi: 10.1109 / JBHI.2022.3167314.

[28] M. Shin, Y. Cho, H. Kang, and M. Kim, “Meta-learning-based lightweight learning framework for healthcare recommendation system,” in *2022 13th International Conference on Information and Communication Technology Convergence ICTC*, Oct. 2022, pp. 1107–1109. doi: 10.1109 / ICTC55196.2022.9952713.

[29] W. Sarah, “Cross-Cultural Sentiment Dynamics: Enhancing Product Recommendation Systems with Multilingual Emotion-Aware Review Analysis”.

[30] L. Wang and E.-P. Lim, “Zero-Shot Next-Item Recommendation using Large Pretrained Language Models,” Apr. 06, 2023, *arXiv*: arXiv:2304.03153. doi: 10.48550 / arXiv.2304.03153.

[31] Z. Kuang, X. Xiong, G. Wu, F. Wang, J. Zhao, and D. Sun, “A Recommendation System for Trigger–Action Programming Rules via Graph Contrastive Learning,” *Sensors*, vol. 24, no. 18, p. 6151, Sept. 2024, doi: 10.3390/s24186151.

[32] S. Shetty, V. S. Ananthanarayana, and A. Mahale, “Cross-modal Deep Learning-based Clinical Recommendation System for Radiology Report Generation from Chest X-rays,” *Int. J. Eng.*, vol. 36, no. 8, pp. 1569–1577, 2023, doi: 10.5829/ije.2023.36.08b.16.

[33] S. Gupta, A. Kumar, and D. Prasad, “Multimodal Graph-based Recommendation System using Hybrid Filtering Approach,” *Int. J. Comput. Digit. Syst.*, vol. 17, no. 1, pp. 1–15, Jan. 2025, doi: 10.12785 / ijcds/1571047857.

[34] L. H. Nazer *et al.*, “Bias in artificial intelligence algorithms and recommendations for mitigation,” *PLOS Digit. Health*, vol. 2, no. 6, p. e0000278, June 2023, doi: 10.1371 / journal.pdig.0000278.

[35] M. Sasseville *et al.*, “Bias Mitigation in Primary Health Care Artificial Intelligence Models: Scoping Review,” *J. Med. Internet Res.*, vol. 27, no. 1, p. e60269, Jan. 2025, doi: 10.2196 / 60269.

[36] M. Haroon, A. Chhabra, X. Liu, P. Mohapatra, Z. Shafiq, and M. Wojcieszak, “YouTube, The Great Radicalizer? Auditing and Mitigating Ideological Biases in YouTube Recommendations,” Mar. 25, 2022, *arXiv*: arXiv:2203.10666. doi: 10.48550 / arXiv.2203.10666.

[37] P. Rouzrokh *et al.*, “Mitigating Bias in Radiology Machine Learning: 1. Data Handling,” *Radiol. Artif. Intell.*, vol. 4, no. 5, p. e210290, Sept. 2022, doi: 10.1148 / ryai.210290.

[38] A. Mihan, A. Pandey, and H. G. Van Spall, “Mitigating the risk of artificial intelligence bias in cardiovascular care,” *Lancet Digit. Health*, vol. 6, no. 10, pp. e749–e754, Oct. 2024, doi: 10.1016 / S2589-7500 (24) 00155-9.

[39] J. Yi, F. Wu, C. Wu, Q. Li, G. Sun, and X. Xie, “DebiasedRec: Bias-aware User Modeling and Click Prediction for Personalized News Recommendation,” Apr. 15, 2021, *arXiv*: arXiv:2104.07360. doi: 10.48550 / arXiv.2104.07360.

[40] A. S. Narasimha Raju, S. Gurunathan, R. Kumar Gatla, S. Ankalaki, and G. S. Hukkeri, “XAI-SkinCADx: A Six-Stage Explainable Deep Ensemble Framework for Skin Cancer Diagnosis and Risk-Based Clinical Recommendations,” *IEEE Access*, vol. 13, pp. 176583–176621, 2025, doi: 10.1109 / ACCESS.2025.3616738.

[41] D. Sanghavi, A. Bhatt, D. Patel, F. Sanjanwala, and K. Bhowmick, “Transparency in Medical Recommendations: A Comprehensive Methodology of Explainable AI Techniques in Healthcare,” vol. 33, no. 4, 2024.

[42] “Explainable Recommender System for Patient Referral in the Healthcare Setting - ProQuest.” Accessed: July 11, 2025. [Online]. Available: https://www.proquest.com/openview/22fe667061d6a52c96ef6e376ce92d57/1?pq-origsite=gscholar&cbl=2026366&diss=y.

[43] T. Bhardwaj and K. Sumangali, “An explainable federated blockchain framework with privacy-preserving AI optimization for securing healthcare data,” *Sci. Rep.*, vol. 15, no. 1, p. 21799, July 2025, doi: 10.1038 / s41598-025-04083-4.

[44] M. Benleulmi, I. Gasmi, N. Azizi, and N. Dey, “Explainable AI and deep learning models for recommender & nbsp; systems: State of the art and challenges,” *JUCS - J. Univers. Comput. Sci.*, vol. 31, no. 4, pp. 383–421, Mar. 2025, doi: 10.3897 / jucs.122380.

[45] R. F. Doh, C. Zhou, J. K. Arthur, I. Tawiah, and B. Doh, “A Systematic Review of Deep Knowledge Graph-Based Recommender Systems, with Focus on Explainable Embeddings,” *Data*, vol. 7, no. 7, Art. no. 7, July 2022, doi: 10.3390/data7070094.

[46] M. Amiri, F. Sarani Rad, and J. Li, “Delighting Palates with AI: Reinforcement Learning’s Triumph in Crafting Personalized Meal Plans with High User Acceptance,” *Nutrients*, vol. 16, no. 3, p. 346, Jan. 2024, doi: 10.3390/nu16030346.

[47] Y. Yang and Y. Zhao, “Personalized Sports Health Recommendation System Assisted by Q-Learning Algorithm,” *Int. J. Human–Computer Interact.*, vol. 41, no. 4, pp. 1889–1901, Feb. 2025, doi: 10.1080/10447318.2023.2295693.

[48] Y. Zhao *et al.*, “Time-aware Path Reasoning on Knowledge Graph for Recommendation,” *ACM Trans Inf Syst*, vol. 41, no. 2, p. 26:1-26:26, Dec. 2022, doi: 10.1145 / 3531267.

[49] D. J. Tan, Q. Xu, K. C. See, D. Perera, and M. Feng, “Advancing Multi-Organ Disease Care: A Hierarchical Multi-Agent Reinforcement Learning Framework,” Aug. 07, 2025, *arXiv*: arXiv:2409.04224. doi: 10.48550/arXiv.2409.04224.

[50] Z. Guo, P. Song, C. Feng, K. Yao, and J. Liang, “Causal Inference for Multi-Criteria Rating Recommender Systems,” *ACM Trans Inf Syst*, vol. 43, no. 6, p. 158:1-158:30, Sept. 2025, doi: 10.1145 / 3757737.

[51] M. B. Tariq and H. A. Habib, “A Reinforcement Learning Based Recommendation System to Improve Performance of Students in Outcome Based Education Model,” *IEEE Access*, vol. 12, pp. 36586–36605, 2024, doi: 10.1109/ACCESS.2024.3370852.

[52] H. Hosseinmardi, A. Ghasemian, M. Rivera-Lanas, M. Horta Ribeiro, R. West, and D. J. Watts, “Causally estimating the effect of YouTube’s recommender system using counterfactual bots,” *Proc. Natl. Acad. Sci. U. S. A.*, vol. 121, no. 8, p. e2313377121, doi: 10.1073 / pnas.2313377121.

[53] X. Wang, “Towards Causality-centric Recommender Systems,” Thesis, 2025. Accessed: Nov. 08, 2025. [Online]. Available: https://opus.lib.uts.edu.au/handle/10453/187995.

[54] X. Ma, M. Li, and X. Liu, “Advancements in Recommender Systems: A Comprehensive Analysis Based on Data, Algorithms, and Evaluation,” vol. 5, no. 1, 2024.

[55] Z. Zhang, Q. Dai, X. Chen, Z. Dong, and R. Tang, “Robust Causal Inference for Recommender System to Overcome Noisy Confounders,” in *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, Taipei Taiwan: ACM, July 2023, pp. 2349–2353. doi: 10.1145/3539618.3592055.

[56] R. Qiu, S. Wang, Z. Chen, H. Yin, and Z. Huang, “CausalRec: Causal Inference for Visual Debiasing in Visually-Aware Recommendation,” in *Proceedings of the 29th ACM International Conference on Multimedia*, Oct. 2021, pp. 3844–3852. doi: 10.1145 / 3474085.3475266.

[57] K. Krauth, Y. Wang, and M. Jordan, “Breaking Feedback Loops in Recommender Systems with Causal Inference,” *ACM Trans. Recomm. Syst.*, vol. 4, no. 1, pp. 1–20, Mar. 2026, doi: 10.1145 / 3728372.

[58] S. Xu *et al.*, “Causal Structure Learning for Recommender System,” *ACM Trans. Recomm. Syst.*, vol. 3, no. 1, pp. 1–23, Mar. 2025, doi: 10.1145 / 3680296.

[59] E. Cavenaghi, A. Zanga, F. Stella, and M. Zanker, “Towards a Causal Decision-Making Framework for Recommender Systems,” *ACM Trans. Recomm. Syst.*, vol. 2, no. 2, pp. 1–34, June 2024, doi: 10.1145 / 3629169.

[60] Y. Ding, Y. Sun, and J. Feng, “The Application of Causal Inference Algorithms in Federated Recommender Systems,” *IEEE Access*, vol. 12, pp. 29748–29758, 2024, doi: 10.1109 / ACCESS.2023.3342861.

[61] J. Liu, Z. Zhang, X. Zhao, and J. Li, “Counterfactual Adversarial Learning for Recommendation,” in *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, in CIKM ’23. New York, NY, USA: Association for Computing Machinery, Oct. 2023, pp. 4115–4119. doi: 10.1145 / 3583780.3615152.

[62] A. Chotrani, “Ethical Considerations in Deploying Machine Learning Models in Healthcare,” vol. 10, no. 1, 2021.

[63] S. Tiribelli and D. Calvaresi, “Rethinking Health Recommender Systems for Active Aging: An Autonomy-Based Ethical Analysis,” *Sci. Eng. Ethics*, vol. 30, no. 3, p. 22, May 2024, doi: 10.1007 / s11948-024-00479-z.

[64] H. Javed, H. A. Muqeet, T. Javed, A. U. Rehman, and R. Sadiq, “Ethical Frameworks for Machine Learning in Sensitive Healthcare Applications,” *IEEE Access*, vol. 12, pp. 16233–16254, 2024, doi: 10.1109/ACCESS.2023.3340884.

[65] M. Kwiatkowska, “Recommender Systems in Healthcare: A Socio-Technical Systems Approach,” in *Recommender Systems for Medicine and Music*, Z. W. Ras, A. Wieczorkowska, and S. Tsumoto, Eds., Cham: Springer International Publishing, 2021, pp. 13–28. doi: 10.1007 / 978-3-030-66450-3_2.

[66] Q. Li, “Harnessing the Power of Pre-trained Vision-Language Models for Efficient Medical Report Generation,” in *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, Birmingham United Kingdom: ACM, Oct. 2023, pp. 1308–1317. doi: 10.1145 / 3583780.3614961.

[67] O. N. Manzari, H. Ahmadabadi, H. Kashiani, S. B. Shokouhi, and A. Ayatollahi, “MedViT: A robust vision transformer for generalized medical image classification,” *Comput. Biol. Med.*, vol. 157, p. 106791, May 2023, doi: 10.1016/j.compbiomed.2023.106791.

[68] W. Huang *et al.*, “Enhancing the vision-language foundation model with key semantic knowledge-emphasized report refinement,” *Med. Image Anal.*, vol. 97, p. 103299, Oct. 2024, doi: 10.1016 / j.media.2024.103299.

[69] S. Schmidgall, J. Cho, C. Zakka, and W. Hiesinger, “GP-VLS: A general-purpose vision language model for surgery,” Aug. 06, 2024, *arXiv*: arXiv:2407.19305. doi: 10.48550 / arXiv.2407.19305.

[70] Z. Zhang, Y. Yu, Y. Chen, X. Yang, and S. Y. Yeo, “MedUnifier: Unifying Vision-and-Language Pre-training on Medical Data with Vision Generation Task using Discrete Visual Representations”.

[71] “(PDF) Development and Evaluation of a Nutritional Smartphone Application for Making Smart and Healthy Choices in Grocery Shopping,” *ResearchGate*, Aug. 2025, doi: 10.4258 / hir.2017.23.1.16.

[72] “Integrating Wearable Devices and Recommendation System: Towards a Next Generation Healthcare Service Delivery Request PDF,” *ResearchGate*, Aug. 2025, Accessed: Nov. 19, 2025. [Online]. Available: https://www.researchgate.net/publication/330535536_Integrating_Wearable_Devices_and_Recommendation_System_Towards_a_Next_Generation_Healthcare_Service_Delivery.

[73] P. Alves, A. Martins, P. Novais, and G. Marreiros, “Improving Group Recommendations using Personality, Dynamic Clustering and Multi-Agent MicroServices,” in *Proceedings of the 17th ACM Conference on Recommender Systems*, in RecSys ’23. New York, NY, USA: Association for Computing Machinery, Sept. 2023, pp. 1165–1168. doi: 10.1145 / 3604915.3610653.

[74] “Knowledge-based dietary nutrition recommendation for obese management,” *ResearchGate*, Aug. 2025, doi: 10.1007/s10799-015-0218-4.

[75] D. Diyasena, N. Arambepola, and L. Munasinghe, “Effectiveness of Human-in-the-loop Design Concept for eHealth Systems,” 2022.

[76] F. Gräßer, F. Tesch, J. Schmitt, S. Abraham, H. Malberg, and S. Zaunseder, "A pharmaceutical therapy recommender system enabling shared decision-making," *User Model. User-Adapt. Interact.*, vol. 32, no. 5, pp. 1019–1062, Nov. 2022, doi: 10.1007 / s11257-021-09298-4.

[77] "Development of Smart Elderly Care Mobile Application for Health Management System." Accessed: Nov. 13, 2025. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/9936853.

[78] M. C. Comes *et al.*, "Early prediction of neoadjuvant chemotherapy response by exploiting a transfer learning approach on breast DCE-MRIs," *Sci. Rep.*, vol. 11, no. 1, p. 14123, July 2021, doi: 10.1038/s41598-021-93592-z.

[79] A. Nilla and E. B. Setiawan, "Film Recommendation System Using Content-Based Filtering and the Convolutional Neural Network CNN Classification Methods," *J. Ilm. Tek. Elektro Komput. Dan Inform.*, vol. 10, no. 1, p. 17, Feb. 2024, doi: 10.26555 / jiteki.v9i4.28113.

[80] N. K. Al-Qazzaz, A. A. Aldoori, A. K. Buniya, S. H. B. M. Ali, and S. A. Ahmad, "Transfer Learning and Hybrid Deep Convolutional Neural Networks Models for Autism Spectrum Disorder Classification From EEG Signals," *IEEE Access*, vol. 12, pp. 64510–64530, 2024, doi: 10.1109 / ACCESS.2024.3396869.

[81] M. Arakeri, D. M P, A. V. Kavan, K. S. Murthy, N. L. Nishitha, and N. Lakshmi, "Intelligent pesticide recommendation system for cocoa plant using computer vision and deep learning techniques," *Environ. Res. Commun.*, vol. 6, no. 7, p. 075003, July 2024, doi: 10.1088 / 2515-7620/ad58ae.

[82] S. Bourhim, L. Benhiba, and M. A. J. Idrissi, "A Community-Driven Deep Collaborative Approach for Recommender Systems," *IEEE Access*, vol. 10, pp. 131144–131152, 2022, doi: 10.1109 / ACCESS.2022.3230323.

[83] K. Karre, "MONOLITHIC HYBRID THERAPY RECOMMENDER SYSTEM FOR AUTISTIC CHILDREN USING MACHINE LEARNING," . *Vol.*, no. 12.

[84] C. Iwendi, S. Khan, J. H. Anajemba, A. K. Bashir, and F. Noor, "Realizing an Efficient IoMT-Assisted Patient Diet Recommendation System Through Machine Learning Model," *IEEE Access*, vol. 8, pp. 28462–28474, 2020, doi: 10.1109/ACCESS.2020.2968537.

[85] A. P. Ponselvakumar, S. Anandamurugan, K. Logeswaran, S. Nivashini, S. K. Showentharya, and S. Swetha Jayashree, "Advancement in Precision Medicine and Recommendation System for Clinical Trials Using Deep Learning Methods," *IOP Conf. Ser. Mater. Sci. Eng.*, vol. 1055, no. 1, p. 012110, Feb. 2021, doi: 10.1088 / 1757-899X / 1055 / 1 / 012110.

[86] Y. S. Cho and P. C. Hong, "Applying Machine Learning to Healthcare Operations Management: CNN-Based Model for Malaria Diagnosis," *Healthcare*, vol. 11, no. 12, p. 1779, June 2023, doi: 10.3390 / healthcare11121779.

[87] Y. Cai, F. Yu, M. Kumar, R. Gladney, and J. Mostafa, "Health Recommender Systems Development, Usage, and Evaluation from 2010 to 2022: A Scoping Review," *Int. J. Environ. Res. Public. Health*, vol. 19, no. 22, Art. no. 22, Jan. 2022, doi: 10.3390 / ijerph192215115.

[88] B. Subramanian, J. Kim, M. Maray, and A. Paul, "Digital Twin Model: A Real-Time Emotion Recognition System for Personalized Healthcare," *IEEE Access*, vol. 10, pp. 81155–81165, 2022, doi: 10.1109/ACCESS.2022.3193941.

[89] K. KC, Z. Yin, M. Wu, and Z. Wu, "Evaluation of deep learning-based approaches for COVID-19 classification based on chest X-ray images," *Signal Image Video Process.*, vol. 15, no. 5, pp. 959–966, July 2021, doi: 10.1007/s11760-020-01820-2.

[90] "Datasets | Center for Artificial Intelligence in Medicine & Imaging." Accessed: Nov. 14, 2025. [Online]. Available: https://aimidatasetindex.stanford.edu/datasets.

[91] CDC, "National Health and Nutrition Examination Survey," National Health and Nutrition Examination Survey. Accessed: Nov. 14, 2025. [Online]. Available: https://www.cdc.gov/nchs/nhanes/index.html.

[92] "Comparative Analysis of Machine Learning Models for Predictive Healthcare in Chronic Disease Management," *J. Intell. Syst. Internet Things*, vol. 17, no. 2, 2025, doi: 10.54216 / JISIoT.170204.

[93] G. Mzoughi, J. Fattahi, M. Mejri, R. Ghayoula, and S. Bahroun, "ECG-Based Stress Surveillance Using an Attention-Driven Hybrid CNN-RNN Model," *2025 17th Int. Conf. Electron. Comput. Artif. Intell. ECAI*, pp. 1–10, June 2025, doi: 10.1109 / ECAI65401.2025.11095605.

[94] "AffectNet+: A Database for Enhancing Facial Expression Recognition with Soft-Labels." Accessed: Nov. 14, 2025. [Online]. Available: https://arxiv.org/html/2410.22506v1.

[95] "Biometric Datasets," unidata.pro. Accessed: Nov. 14, 2025. [Online]. Available: https://unidata.pro/datasets/biometrics/.

[96] "Longitudinal Patient Data LPD - IQVIA." Accessed: Nov. 14, 2025. [Online]. Available: https://www.iqvia.com/locations/united-kingdom/solutions/life-sciences-industry-solutions/market-intelligence/

longitudinal-patient-data.

[97] D. R. Beddiar, M. Oussalah, U. Muhammad, and T. Seppänen, “A Deep learning based data augmentation method to improve COVID-19 detection from medical imaging,” *Knowl.-Based Syst.*, vol. 280, p. 110985, Nov. 2023, doi: 10.1016/j.knosys.2023.110985.

[98] L. Bruni *et al.*, “Cervical cancer screening programs and age-specific coverage estimates for 202 countries and territories worldwide: a review and synthetic analysis,” *Lancet Glob. Health*, vol. 10, no. 8, pp. e1115–e1127, Aug. 2022, doi: 10.1016 / S2214-109X (22) 00241-8.

[99] R. Ramamoorthy, “HME-NET: A HYBRID MULTIMODAL ENSEMBLE FRAMEWORK FOR EXPLAINABLE AND EFFICIENT BONE CANCER DIAGNOSIS,” *Int. J. Appl. Math.*, vol. 38, pp. 1095–1109, Sept. 2025, doi: 10.12732/ijam.v38i4s.289.

[100] M. Wang, Y. Guo, C. Zhang, C. Wang, H. Huang, and X. Jia, “MedShare: A Privacy-Preserving Medical Data Sharing System by Using Blockchain,” *IEEE Trans. Serv. Comput.*, pp. 1–1, 2021, doi: 10.1109/TSC.2021.3114719.

[101] M. Adnan, S. Kalra, J. C. Cresswell, G. W. Taylor, and H. R. Tizhoosh, “Federated learning and differential privacy for medical image analysis,” *Sci. Rep.*, vol. 12, no. 1, p. 1953, Feb. 2022, doi: 10.1038 / s41598-022-05539-7.

[102] S. Che *et al.*, “Federated Multi-view Learning for Private Medical Data Integration and Analysis,” *ACM Trans. Intell. Syst. Technol.*, vol. 13, no. 4, pp. 1–23, Aug. 2022, doi: 10.1145 / 3501816.

[103] M. H. Nguyen *et al.*, “On-device diagnostic recommendation with heterogeneous federated BlockNets,” *Sci. China Inf. Sci.*, vol. 68, no. 4, p. 140102, Apr. 2025, doi: 10.1007 / s11432-024-4162-2.

[104] S. D. N, A. B, S. Hegde, C. S. Abhijit, and S. Ambesange, “FedCure: A Heterogeneity-Aware Personalized Federated Learning Framework for Intelligent Healthcare Applications in IoMT Environments,” *IEEE Access*, vol. 12, pp. 15867–15883, 2024, doi: 10.1109 / ACCESS.2024.3357514.

[105] E. A. Mantey, C. Zhou, J. H. Anajemba, Y. Hamid, and J. K. Arthur, “Blockchain-Enabled Technique for Privacy-Preserved Medical Recommender System,” *IEEE Access*, vol. 11, pp. 40944–40953, 2023, doi: 10.1109/ACCESS.2023.3267431.

[106] E. A. Mantey, C. Zhou, V. Mani, J. K. Arthur, and E. Ibeke, “Maintaining privacy for a recommender system diagnosis using blockchain and deep learning.,” *Hum.-Centric Comput. Inf. Sci.*, vol. 13, Oct. 2023, Accessed: July 11, 2025. [Online]. Available: https://rgu-repository.worktribe.com/output/1645830.

[107] A. S. Fathima *et al.*, “Empowering Consumer Healthcare Through Sensor-Rich Devices Using Federated Learning for Secure Resource Recommendation,” *IEEE Trans. Consum. Electron.*, vol. 71, no. 1, pp. 1563–1570, Feb. 2025, doi: 10.1109 / TCE.2025.3541549.

[108] A. Alqhatani and S. B. Khan, “IoT-Driven Hybrid Deep Collaborative Transformer with Federated Learning for Personalized E-Commerce Recommendations: An Optimized Approach,” *Scalable Comput. Pract. Exp.*, vol. 25, no. 5, Art. no. 5, Aug. 2024, doi: 10.12694 / scpe.v25i5.3007.

[109] M. Souza, “Predictive Counterfactuals for Treatment Effect Heterogeneity in Event Studies with Staggered Adoption,” Nov. 11, 2019, *Social Science Research Network, Rochester, NY*: 3484635. doi: 10.2139 / ssrn.3484635.

[110] Q. Wang, H. Yin, T. Chen, J. Yu, A. Zhou, and X. Zhang, “Fast-adapting and privacy-preserving federated recommender system,” *VLDB J.*, vol. 31, no. 5, pp. 877–896, Sept. 2022, doi: 10.1007 / s00778-021-00700-6.

[111] T. Hussain *et al.*, “EFFResNet-ViT: A Fusion-Based Convolutional and Vision Transformer Model for Explainable Medical Image Classification,” *IEEE Access*, vol. 13, pp. 54040–54068, 2025, doi: 10.1109/ACCESS.2025.3554184.

[112] H. Jin *et al.*, “Evaluating the Effectiveness of a Generative Pretrained Transformer-Based Dietary Recommendation System in Managing Potassium Intake for Hemodialysis Patients,” *J. Ren. Nutr.*, vol. 34, no. 6, pp. 539–545, Nov. 2024, doi: 10.1053 / j.jrn.2024.04.001.

[113] Y. Liu, “Medical Image Classification Based on Transformer Model and Ordinal Loss:,” in *Proceedings of the 1st International Conference on Engineering Management, Information Technology and Intelligence*, Shanghai, China: SCITEPRESS - Science and Technology Publications, 2024, pp. 708–713. doi: 10.5220 / 0012969200004508.

[114] A. Gharahighehi, C. Vens, and K. Pliakos, “Fair Multi-Stakeholder News Recommender System with Hypergraph ranking,” Feb. 09, 2021, *arXiv*: arXiv:2012.00387. doi: 10.48550 / arXiv.2012.00387.

[115] Md. M. H. Shuvo, S. K. Islam, J. Cheng, and B. I. Morshed, “Efficient Acceleration of Deep Learning Inference on Resource-Constrained Edge Devices: A Review,” *Proc. IEEE*, vol. 111, no. 1, pp. 42–91, Jan. 2023, doi: 10.1109/JPROC.2022.3226481.

[116] T. Muazu, M. Yingchi, A. U. Muhammad, M. Ibrahim, O. Samuel, and P. Tiwari, “IoMT: A Medical Resource Management System Using Edge Empowered Blockchain Federated Learning,” *IEEE Trans. Netw. Serv. Manag.*, vol. 21, no. 1, pp. 517–534, Feb. 2024, doi: 10.1109 / TNSM.2023.3308331.