

The Role of Computer Vision and Convolutional Neural Networks in Healthcare Recommender Systems: A Comprehensive Review

Hikmat Z. Neima
Department of Computer Science ,
College of CSIT , University of Basrah , Iraq.

Rana M. Ghadban
Department of Intelligent Medical Systems,
College of CSIT, University of Basrah, Iraq.

Hussein A. Jassim
Department of Computer Science ,
College of CSIT , University of Basrah , Iraq.

Ahmed R. Hmood
Department of Computer Science ,
College of CSIT , University of Basrah , Iraq.

Abstract - The integration of computer vision (CV) and deep learning (DL) has revolutionized the healthcare landscape by enabling intelligent and personalized clinical decision-making. Among the many innovations in this field, Convolutional Neural Networks (CNNs) have emerged as the cornerstone of computer vision technologies, particularly in medical image analysis. Simultaneously, Healthcare Recommender Systems (HRS) have evolved to assist patients and clinicians in selecting appropriate treatments, diagnostic pathways, or lifestyle interventions tailored to individual needs. However, the intersection between CNNs and HRS remains an under-explored but promising research frontier. This paper presents a comprehensive review of recent developments at this intersection, examining how CNNs are being embedded into HRS frameworks to enhance patient-specific recommendations. The foundational applications of CNNs were explored in medical imaging and personalized healthcare, followed by an analysis of how CNNs have been adapted within healthcare recommender systems. Drawing on an extensive analysis of 56 recent and high-impact academic papers, an organized synthesis of system architectures, datasets, performance evaluations, and clinical use cases were provided. The current review highlights the emerging potential of hybrid CNN-based recommendation models and identifies key challenges, such as data sparsity, interpretability, privacy, and real-time applicability. By offering a clear roadmap of current trends, limitations, and opportunities, this paper aims to guide future research in building robust, explainable, and scalable CNN-powered HRS for next-generation healthcare solutions.

Keywords— *Convolutional Neural Networks (CNN); Computer Vision; Healthcare Recommender Systems (HRS); Deep Learning; Medical Imaging; Personalized Medicine.*

I. INTRODUCTION

With the explosion of medical data comes a revolution within the healthcare industry, Digital Transformation. This system is based on a variety of health data streams , such as electronic health records (EHRs) , diagnostic images, as well as consumer-output data via wearable technologies. Because healthcare organizations cannot provide effective intervention suggestions in personalized care , the multimodal nature of multimodal data demands new ways of analytical (and not just the traditional systems) solutions in order to maximize the

benefit provided. Diversified deep learning technology, particularly a new generation of neural networks [1] , is the promising solution to this problem.

The implementation of computer vision technology (CV) assists clinicians in reading the X-ray and MRIs and pathology slides for diagnostic analysis. Convolutional Neural Networks (CNNs) perform significantly better than other CV tools for tasks in image classification, object detection, segmentation, and feature extraction. The promising performance of CNNs in cancer detection and ophthalmology has been demonstrated [2]. Healthcare Recommender Systems (HRS) have emerged as a key way for delivering personalized healthcare. Not only does medical personnel, but patients are also used to receive personalized treatment recommendations, information (including prescription information services) and lifestyle alterations through the systems [3]. Nonetheless, the old solutions of collaborative and content-oriented filtering in HRS are not sufficient to handle complex health data at high dimensions [4].

A. Motivation and Background

Healthcare artificial intelligence integration prompted the emergence of new avenues to facilitate automated decision making and individualized patient care. Computer vision and convolutional neural networks (CNNs) have been proposed for complex analysis of medical images to support doctors for diagnostics and treatment application. They were found to be particularly powerful in medical fields with a lot of imaging data and difficulty of interpretation that requires expert knowledge, for instance in radiology and ophthalmology [5]. HRS are popularly utilized to assist with treatment recommendations, lifestyle modifications, and patient monitoring. These systems take clinical data, patient's history, and preferences to create evidence-based recommendations. Through the use of deep learning-driven techniques , especially using CNNs , HRS is able to use improved feature extraction and produce more informative and context-sensitive results [6].

B. Scope and Objectives

This paper reviews the computer vision and CNN-based models to boost recommender systems performance in healthcare. It shows how recent developments in deep learning technology like CNNs make breakthroughs in healthcare recommendation tasks by enabling systems to leverage visual and contextual data from diverse sources.

The main objectives of this review are:

- To present a detailed overview of CNN fundamentals and their applications in medical imaging.
- To examine the core principles of healthcare recommender systems and how CNNs are integrated into their architecture.
- To survey recent advances across emerging subfields, including fairness, cold-start solutions, interpretability, reinforcement learning, causal modeling, and ethics in HRS.
- To identify current research challenges and suggest promising future directions.

C. Organization of the Review

The remainder of this review is organized as follows:

- Section 2 presents the fundamental background of medical imaging, CNN architectures, and the principles of healthcare recommender systems.
- Section 3 explains the integration of CNNs with HRS.
- Section 4 details discussion of recent advances across several subfields.
- Section 5 outlines challenges and open problems in this research domain.
- Section 6 concludes the paper and reports directions for future exploration.

II. FUNDAMENTALS AND BACKGROUND

This subsection is dedicated to the exploration of the basics of medical imaging, deep learning and recommendation algorithms, as they are the fundamental understanding in order to understand the convergence of CNN-based computer vision with healthcare recommender systems (HRS). Those principles formed the background for the subsequent interdisciplinary progress discussed.

A. Healthcare Recommender Systems (HRS)

Although CNNs have impacted the field of image analysis in the context of healthcare, the full application of CNNs in personalized decision-making systems only becomes more interesting when they are combined with other systems. Consequently, we have Healthcare Recommender Systems (HRS) that are progressively driven by deep learning to offer data-driven patient-centered decisions. HRS aim to deliver personalized medical content, including treatment plans, dietary advice, or clinical resources tailored to the patient's preferences [7]. These systems leverage enormous amounts of structured and unstructured health data for better clinical decision-making.

Categories of HRS:

- **Content-Based Filtering:** It recommends items based on the similarity between item features and user profiles. Similarity can be computed using cosine similarity:

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

This measures how closely a user's feature vector aligns with an item's attributes, which is especially important when explicit item metadata is available.

- **Collaborative Filtering:** It relies on historical interactions of users with items. A common approach is latent factor modeling through matrix factorization:

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

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

- **Hybrid Systems:** They combine both approaches to balance personalization and generalization, reducing issues like cold-start and sparsity [9]. Hybrid recommenders are increasingly used to improve robustness when input data is incomplete or inconsistent.

B. Medical Imaging in Healthcare

Now days in the healthcare, medical imaging not only underpins the health diagnosis but also patients' health monitoring. Doctors use medical imaging to examine the body from within and take a look at disease states without need for invasive diagnostic techniques. Diagnostic Imaging. In most cases, medical professionals will depend on X-ray imaging, computed tomography (CT), magnetic resonance imaging (MRI) scanner, ultrasound examination, and positron emission tomography (PET) imaging. Various imaging techniques also have their pros in the clinical application; as in case of MRI they help the quality of the MRI to assess tissue, while CT is better for bone examination, and evaluation of acute trauma [10].

Images are very high degrees of complexity and a big volume of data. — Examiners require extensive time and expose the possibility of marginal differences between manual interpretation in manual results of an MRI scan having its various image slices. Hence the automated diagnostics using machine learning and computer vision systems has increased the accuracy of diagnosis and efficiencies [11].

These systems treat medical images as multi-dimensional arrays:

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

The formula defines a medical image dimension structure by representing image height (H), width (W), and depth (D) for 3D modalities like CT or MRI, as well as channel count (C) which represents the number of channels, so that C equals 1 in a grayscale image and 3 in the case of RGB image. Many techniques, such as segmentation, detection, and classification, are often applied to these inputs to obtain a precise diagnosis.

C. Convolutional Neural Networks

Medical imaging provides detailed visual data. This data contains rich information ; however , processing and then extracting beneficial information from this extensive amount of data requires high computational complexity. Deep learning , specifically convolutional neural networks, serves as the next step in processing this data due to its high capabilities. Nowadays , modern image understanding tasks in healthcare rely heavily on the field of deep learning.

Deep learning, which is a subfield of machine learning, allows models to learn hierarchical data representations through multiple processing layers. As a deep learning architecture, the convolutional neural network (CNN) is particularly effective for image analysis, because it can significantly handle multi-scale features and preserve image spatial relationship [12].

A typical CNN includes:

- 1) Convolutional layers,
- 2) activation functions (e.g. ReLU),
- 3) Pooling layers, and
- 4) fully connected layers.

It is possible to express the convolutional operation as:

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

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

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

To reduce computational complexity and introduce spatial invariance, pooling layers are applied. The final classification layer usually employs the softmax function:

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

CNNs have proven exceptionally useful in medical imaging tasks, such as tumor detection, organ segmentation, and disease classification [13]. Popular architectures include AlexNet, VGG, ResNet, and U-Net for segmentation tasks [14].

Given the limited availability of labeled medical data, transfer learning is frequently used. Pre-trained CNNs on datasets like ImageNet are fine-tuned on medical datasets, which improves generalization and reduces overfitting [12].

III. INTEGRATION OF CNNs IN HEALTHCARE RECOMMENDER SYSTEMS

A. CNNs in Medical Imaging

CNNs have a great ability to extract the complex spatial features from high-dimensional images. Therefore, they are widely utilized in the field of medical imaging. Because of their hierarchical structure and ability to model non-linear patterns, CNNs are suitable for visual analysis in clinical settings.

Nowadays, CNNs are widely applied in a various modality of medical imaging including X-ray, computed tomography (CT), magnetic resonance imaging (MRI), positron emission tomography (PET), and ultrasonography [15].

Medical imaging tasks powered by CNNs can be broadly categorized as follows:

- Image classification: Determining the overall category of an input image (e.g., identifying whether an X-ray scan is normal or indicative of pneumonia).
- Object detection: Localizing specific features within an image, such as identifying the coordinates of a tumor.
- Semantic segmentation: Assigning a class label to each pixel in an image, which is critical in delineating tumors, organs, or lesions.

In order to prove their clinical effectiveness, CNNs have been integrated into numerous studies. For instance, Shen et al. applied deep CNNs to classify breast cancer. They achieved accuracy rate comparable to the accuracy rate achieved in experienced radiologists [15]. Jayawardena et al. developed a diagnostic system based on CheXNet. A CNN was trained on chest X-rays [16], which raises the performance of radiologist-level of the original CheXNet model for pneumonia detection. Deep learning models can serve as reliable assistants in diagnostic, decision-making, proving their robustness and reliability, and this is illustrated through these breakthroughs. Different architectures have been proposed to meet specific tasks. For example, recent work [17] demonstrates that U-Net architectures remain widely used for biomedical segmentation due to the symmetry of the encoder-decoder design and the skip connections that preserve spatial detail. Residual connections were introduced in ResNet [14] to facilitate deep networks learning, i.e., to tackle the vanishing gradients problem. Lately, attention-based mechanisms, such as the Convolutional Block Attention Module (CBAM) [1], have been developed to facilitate the interpretability of CNN models by concentrating on the most informative regions in biometric or medical scans. Together, these architectural innovations have pushed the boundaries of what automated systems can achieve in the field of clinical imaging.

The capabilities, however, present difficulties for the system to operate effectively. The success of CNNs for generalization relies heavily on obtaining large amounts of annotated datasets. Healthcare institutions face major challenges when using these networks in clinical settings due to privacy matters, inconsistent data quality, and limited availability of professional image labels. The strengths of CNNs in prediction tasks do not eliminate their mystery since the inner workings operate like a 'black box'. Professionals are less likely to accept AI technology when they cannot identify the basis of automated decisions. The techniques of Gradient-weighted Class Activation Mapping (Grad-CAM) [18] generate visual explanations to show which image elements guided the model in its decisions thus improving model transparency and trust. CNNs advance in the healthcare industry because of their strong performance alongside their scalability and flexibility features despite facing technical challenges. Decision-support systems have integrated these tools as essential elements that expanded their diagnostic functions. The following section discusses how CNNs normalize their applications in diagnostics and imaging by

integrating with personalized healthcare recommender systems built for specific patient characteristics.

B. Deep Learning for Personalized Healthcare

Research findings have demonstrated the benefits of deep learning for recommender systems, especially for healthcare. CNNs can be used to treat input unstructured data such as radiographs, patients' charts or wearable sensor data and provide image and context information to user in the recommendation process [7].

Techniques such as RBM-CNN architectures enhance health insights by integrating feature extraction with collaborative filtering [19]. Neural collaborative filtering with deep autoencoders, CNN-LSTM architectures, and hybrid (HCLNet), used for HRS tasks are among the architectures [20]. These hybrid models have the potential to improve classification/recommendations of healthcare products by learning from both spatial and sequential data. They are capable of predicting future events to health, providing preventive strategies, and dynamically adjusting according to the changes in the patient data [21].

The medical field achieves patient-centered care innovations by using deep learning techniques for both treatment approaches developed by machines and automating the selection of the types of medicine used as well as methods to predict diseases. The sector's transfer to computational model called the recommender system that provides personalized healthcare suggestions becomes apparent as you interpret data to develop personalized health recommendation systems is introduced. Current medical recommender systems (HRS), for example, have adopted deep learning technologies to make automatic adjustments to patients' needs on an ongoing basis of clinical relevance [22], [23]. Deep learning allows HRS systems to learn from different datasets. These could be EHRs, genomic sequences, wearable sensor outputs, lifestyle files. These systems capitalize on relationships identified in multimodal data to recommend personalizing interventions for chronic diseases such as diabetes, cardiovascular disease, and mental health conditions [7].

CNNs were initially developed for visual data; however, they have been applied to personalized healthcare via hybrid and multimodal architectures. CNNs have been used in dietary recommendation applications or to analyze patterns of physical activity by processing food images or activity visuals derived from the sensor. Integrating these visual or physiological cues with structured health data creates an enhanced context awareness in HRS [24]. As Deep Learning owns a high ability to analyze large, multidimensional, and heterogeneous patient data, it has become a key technology enabling the development of personalized health technologies. Accordingly, the use of DL leads to providing accurate and context-sensitive recommendations. Deep learning systems with advanced architecture including CNNs and RNNs assist in combining patients' medical data from multiple sources (clinical records, sensor data, diagnostic images, and genomic information) to generate personalized recommendations that adjust to individual health situations [9].

Due to their ability to extract deep representations from unstructured inputs, such as radiological images, medical reports, and wearable sensor streams, CNNs are particularly suited for this domain. Integration of these abilities with that of longitudinal data analysis of the RNN or temporal convolutional

networks (TCNs) results in systems that could be able to model patient trajectory and predict future health events [25]. A relevant obvious example is recommending preventive care or early screening if there are minor signals noticed from past imaging or the trend of symptoms. For example, deep hybrid models have been applied to suggest follow-up screenings for high-risk patients, as well as to suggest treatment plans based on patient similarity analysis in a large healthcare formatted dataset [11]. These systems can also help with population-level health management by identifying developing health patterns in particular populations.

These smart systems improve patient engagement by customizing lifestyle and medication suggestions in real time. Mobile health (mHealth) platforms driven by deep learning can track adherence, alert to possible complications, and adjust interventions on an ongoing basis. This is conducive to moving towards proactive and preventive care [19]. The utilization of real-time feedback tools in these systems also enhances the communication between patients and healthcare providers.

C. Hybrid Systems: CNN and Recommender Models

Another trend is using combined CNNs and collaborative filtering methods instead. But when used together the system reveals correspondences between patients with identical medical profiles. Such a technique can prove to be very useful in the clinical arena, for instance, if different symptoms occur in patients who may receive the same treatment [26].

In recent works, several authors attempted to use Explainable AI (XAI) frameworks, the foundation behind deep learning-based HRS, to improve trust and interpretability. Such models promote user and clinician acceptance and system transparency, as well as justifying an explanation of the recommendation in a clear manner [3]. They also implement attention mechanisms and saliency maps in CNN based HRS to highlight specific features (e.g., image regions, clinical markers, and so on) that have the most influence on the recommendation.

CNN-enhanced HRS have also been tested on a variety of tasks including: individualized mental wellness coaching, individual adaptive exercise training, medication adherence reminders and pandemic-related health monitoring [27]. This further increases the participation of patients in health care and offers some useful knowledge about how to prevent disease.

For all of this promise, though, there are still concerns for data privacy, model bias and deployment under limited conditions. The handling of personal health data needs to be done so data can be handled in a manner that complies with regulations, such as HIPAA and GDPR. Moreover, models trained on imbalanced or unrepresentative datasets do not necessarily produce unbiased results, so models have to be fair-minded and be proven on strong validation.

To sum it up, architectures that add to CNN-based architectures contribute in recreating deep learning towards healthcare recommender systems in the new shape. Being able to process multimodal data, provide real-time personalization, as well as incorporate interpretability tools makes it a highly essential part of the intelligent delivery of healthcare in the future.

IV. RECENT ADVANCES IN HRS

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

HRS often suffer in the scenarios of cold-start. This is because of insufficient patients' historical data that can be used to generate accurate and personalized recommendations. Such problem is very critical especially in clinical settings where the decision-making is required in a case of time-sensitive for first-time users. For the sake of this limitation overcoming, many researchers have introduced meta-learning, zero-shot learning, and hybrid techniques to build adaptive systems capable of making informed predictions even with limited user input.

Authors of [28] have proposed a meta-learning-based health recommender framework. Their framework leverages few-shot learning to quickly adapt to new users and clinical tasks. In their model, transferable patterns are extracted from prior user interaction leading to achieve cross-user generalization. Similarly, W. Sarah [29] introduced a dual-attention meta-learning strategy that makes models to be adaptive to unseen patients in chronic care scenarios. This is done by focusing on both task-relevant and user-specific information. Complementing these, L. Wang and E.-P. Lim [30] created a zero-shot clinical recommendation pipeline by pretraining models using contrastive learning. This allowed the system to help people make decisions even without labeled examples from the target area.

To make it even more general, Z. Kuang et al. [31] showed a user embedding approach that was trained with contrastive goals. This method learned shared latent representations that make it easier to make good recommendations for patients whose profiles are similar but not complete. S. Shetty et al. [32] looked into cross-domain suggestion using latent space alignment in the context of domain sparsity. This lets knowledge flow from medical domains with lots of data to those with few, which makes personalization better for underserved user groups. In a hybrid approach to cold-start, S. Gupta et al. [33] suggested a graph-based solution that combines collaborative filtering with content-based methods. This solution uses meta-path analysis and propagation techniques to make reliable suggestions even when there is no user past.

All together, these methods show promising ways to deal with the two problems that healthcare recommender systems face: cold-start and generalization. Still, there are some things we don't know about how these frameworks will work with dynamic patient feedback, how they will change over time, and how they will be calibrated for rare conditions.

B. Fairness and Bias Mitigation in Healthcare Recommender Systems

With the increasing use of healthcare recommender systems (HRS) in real-life hospital settings, it is increasingly important to ensure that HRS are fair toward all groups of patients. However, when healthcare data is skewed or biased, such as from members of under-represented groups, unbalanced clinical records, or systemic differences, those biases can create suggestions that serve to disproportionately harm vulnerable groups over other groups. So, a good deal of recent work has been done on making algorithms aware of fairness so they can detect, measure and minimize these types of biases, and be helpful in the clinical environment.

As an example of this, L. H. Nazer et al. [34] designed an adaptive debiasing framework that exploits collaborative filtering to fix exposure and selection biases of HRS. Their approach dynamically updates the weights of training examples on the fly to account for distorted past data. This ensures that advice results are equitable for all user groups. On the other hand, M. Sasseville et al. [35] indicated fairness-aware matrix factorization application for personalized health information search. That was done to make recommendations fairer for less-represented users. Their formulation embeds group fairness constraints into the model objective to ensure that people from all backgrounds are equally treated. M. Haroon et al. [36] expanded the concept of fairness in clinical settings and introduced a collaborative filtering architecture to introduce fairness constraints at optimization layers to minimize the differences in the exposure of healthcare items. Similarly, P. Rouzrokh et al. [37] developed an oncology decision support system bias-aware patient clustering approach. Their model aligns patients' subgroups with more easily discernible and evenly distributed characteristics across demographics, to make this fairer. Researchers in [38] presented an adversarial learning model to eliminate the biases due to age and gender in the latent space. This helped to align demographic-sensitive features and enabled easy generalization across different types of people. Lastly, J. Yi et al. [39] developed DebiasRec, an RS framework sensitive to fairness using representations that consider causes and interventions to support transparent and equitable recommendations, and particularly in sensitive domains (e.g., mental health).

These studies demonstrate the growing dedication to fairness and accountability in health care recommender systems. Nonetheless, several issues arise that must be thoroughly explored. These are the problems of the balance of justice and prediction accuracy, what constitutes a universal fairness measure in medicine, and maintaining clarity in complex RS pipelines. It remains important to incorporate fairness-aware features as HRS continues to grow. This makes AI powered healthcare solutions more ethical and socially responsible.

C. Explainable AI (XAI) and Interpretability in HRS

In healthcare recommender systems (HRS), demand for explainability has become significant given the importance of clinical decision-making and its high-stakes aspect. Many black-box models, more particularly deep CNN-based recommenders, are prone to concerns regarding transparency and reliability. XAI is one of the methods developed to overcome these limitations to provide human-interpretable explanations for recommendations. A fundamental work for visual interpretability is the Grad-CAM method by A. S. Narasimha Raju et al. [40], that emphasizes spatial zones in medical skin images (i.e. regions within the skin) which influence most the CNN decisions and has helped clinicians to validate an automated output for the patient's condition. This visual-analysis approach has been modified for a variety of clinical recommender problems, most notably when focusing on patient stratification based on images.

Beyond visualizations, rule-oriented interpretability is also getting a good deal of attention. For example, Alzahrani et al. devised a neuro-fuzzy staging recommender system combining CNN-facilitated patterns with fuzzy reasoning to enhance the interpretability of risk classifications for chronic diseases [41]. This hybrid approach combines quantitative learning and

qualitative reasoning, thereby enabling a bridge between the clinician and the trainer and connecting the outputs to transparent rule sets. At the systemic level, we have seen the advent of explainable recommender frameworks, such as DECAF (Deep Context-Aware Framework), for optimizing patient referral. DECAF combines clinical context with user history for accurate textual and feature explanations of each referral recommendation, enhancing acceptance by healthcare providers [42].

On the other hand, bias-aware frameworks like DeBiasRec take fairness into account in explainable RS. Deliberately proposed by Lei et al., DeBiasRec provides transparency using intervention-aware embeddings that are traceable to the user interactions and fairness adjustments with explainable explanations for click prediction tasks in sensitive domains like healthcare [43].

It is worth mentioning that M. Benleulmi et al. [44] presents the state-of-the-art review of the available literature on XAI and its adoption in deep learning-based recommender systems that identify the relevant deficiencies, including the absence of standardized evaluation metrics from which explanations can be evaluated, as well as the limited incorporation of causal reasoning frameworks. Finally, Wang et al. proposed a deep knowledge graph-based approach to recommender systems where explainability is derived from semantic relationships across multi-hop paths in the knowledge graph, allowing users to understand recommendations through interpretable symbolic chains (e.g., symptom \rightarrow disease \rightarrow treatment) [45].

These developments reflect a paradigm shift from opaque clinical RS models to transparent, multi-modal, and rule-grounded systems, enhancing clinician trust and aligning more closely with evidence-based medical practice.

D. Reinforcement Learning in HRS

Healthcare recommender systems (HRS) have shown success with Reinforcement Learning (RL) as a method for sequential decision-making, allowing for more tailored and adaptive treatment suggestions. A dynamic, long-term care setting is ideal for RL because, unlike conventional supervised learning approaches, it learns optimal policies by interactions with the environment, taking into account the fact that patient circumstances change over time. The use of RL in clinical decision pathway modeling, drug regimen optimization, and behavioral health intervention customization has been the subject of multiple investigations.

M. Mehdi [23] put forward a reinforcement learning model that relies on actor-critic feedback to provide tailored treatment suggestions. In order to optimize glycemic control policies over several stages, their approach connects patient trajectories with long-term treatment outcomes. An analogous deep RL framework for precision oncology was presented by Zhang et al. [46]. In this framework, therapy selection is represented as a sequential policy optimization job that makes use of genetic and clinical characteristics that are distinct to each patient. In comparison to static classifiers, the system's survival prediction capabilities are enhanced by its ability to dynamically adjust treatment sequences in response to feedback. Similarly, Y. Yang and Y. Zhao [47] used Q-learning to create a mental health intervention recommender that takes into account the patient's contextual information, such as their reported outcomes and

their level of participation over time, and uses these to generate reinforcement signals.

In RL-based HRSs used for chronic disease, Y. Zhao et al. [48] used domain knowledge to shape rewards in a way that improved clinical safety and reward modeling. By combining policy learning with clinical best practices, this hybrid approach lessens the likelihood of harmful recommendations. In order to optimize recommendations within shared limitations, D. J. Tan et al. [49] presented a multi-agent reinforcement learning system for collaborative recommendation in hospital networks. Each agent represents a clinical department. Their method improves patient care coordination by balancing the trade-offs between departments. Z. Guo et al. [50] used the combination of reinforcement learning and causal inference to mitigate treatment selection bias, enhance policy generalization by simulating real-world counterfactual outcomes, and improve the system.

M. B. Tariq and H. A. Habib [51] devised a reinforcement learning model considering trust during healthcare services recommendation on mobile platforms. To ensure the system evolves not only in response to incentives but also based on user dependability and feedback consistency, the model uses user trust dynamics to inform choice of actions. In terms of promoting the effectiveness of the recommender, this is especially important in telehealth and mobile settings, because the extent to which patients take part and conform to this recommendation directly factors into a recommender's success. Taken together, RL-based methods represent an emerging trend towards intelligent, context-aware, morally grounded HRS capable of achieving long-term clinical impacts.

E. Causal Inference in Healthcare Recommender Systems

Causal inference has acquired a vital role in healthcare recommender systems (HRS), particularly with respect to confounding, treatment effect estimation, and generalizability across patient populations. In contrast to classical associative models, causal-aware recommenders seek to approximate counterfactual scenarios, allowing for more reliable and personalized decision support. More recent efforts have implemented causal inference for bias management of data, robustness, and explanation generation that is grounded in real-world medical reasoning.

By using counterfactual estimators, H. Hosseinmardi et al. [52] modeled user-level heterogeneity and established a treatment effect-aware RS for tailored treatments. Through the process of learning individual causal effects from past interventions, the system is able to significantly decrease suggestion bias in electronic health records. Based on this, X. Wang [53] introduced a neural counterfactual learning framework with two branches that uses patient trajectory embeddings to estimate factual and counterfactual outcomes. This improves stability when dealing with sparse data and imbalanced treatments. In order to improve generalizability across unexplored domains, X. Ma et al. [54] investigated a domain-adaptive causal embedding method that aligns patient data from various institutions using invariant causal representations. At the same time, to make more reliable recommendations in observational clinical settings, Z. Zhang et al. [55] used instrumental variable-based causal modeling to separate treatment pathways from latent confounders. A recurrent counterfactual network was suggested by R. Qiu et al.

[56] to handle time-dependent confounding. This network combines temporal attention with propensity scores, allowing the model to be used for longitudinal recommendation tasks.

To ensure fair recommendations, K. Krauth et al. [57] concentrated on fairness-aware causal RS, which uses balanced representations to remove bias caused by demographic differences. A hybrid causal reinforcement learning paradigm was also proposed by S. Xu et al. [58]; this approach uses counterfactual returns to link immediate incentives with the long-term efficacy of a treatment. A knowledge-aware RS can be guided by a graph-based causal discovery technique, which E. Cavenaghi et al. [59] suggested for building clinically significant links between symptoms and therapies. By combining attention-based neural architectures with structural causal models, Y. Ding et al. [60] improved interpretability and made it possible for users to follow recommendations through causal graphs that are easy to understand. In conclusion, a counterfactual inference method utilizing adversarial training was suggested by J. Liu et al. [61] to improve the personalization and robustness of healthcare recommendation tasks by balancing the distributions of facts and counterfactuals.

Causal inference in HRS is becoming more sophisticated and impactful, as shown in these ten works. Safer and more accountable AI-driven healthcare recommendations are possible thanks to these methodologies, which tackle issues including confounding bias, data imbalance, domain change, and fairness.

F. Ethics in Healthcare Recommender Systems

Given the growing penetration of healthcare recommender systems (HRS) in clinical decision-making, ethical issues related to their design, deployment and usage have become salient. Unlike generic recommender systems, HRS drive high-stakes decisions regarding patient well-being, autonomy, and access to care. And so, developers need to address transparency, informed consent, data ownership, and social responsibility of AI models. A. Chotrani et al. [62] highlight the distinct ethical landscape of HRS. Key challenges include opacity of algorithmic logic and the risk of undermining physician-patient trust — as well as the demand for explainable and contestable recommendations in clinical workflows.

Healthcare recommender systems (HRS) raise ethical implications in their development, implementation, and use as they infiltrate clinical decision-making. HRSs do much more than general-purpose recommender systems: they shape decisions that affect a patient's freedom, health, and access to treatment. Hence, there are issues that developers have to address—privacy, consent after being fully informed, who owns the data, and the social responsibility of AI models. The main points raised by A. Chotrani et al. [62] mention are the complexity of the algorithmic logic, the potential for trust to decrease among doctor and patient, and the necessity for clear and arguable recommendations within clinical workflows. Such problems are but a few of the ethical problems unique to HRS.

Moreover, because such AI was utilized in clinical practice, it creates weak ethical frameworks with risk of biased or unexpected results. S. Tiribelli et al. [4] note that many current AI-based systems, like the ones used in the field of HRS, lack sufficient legal and moral protections built in. As a group, they support models based on the biomedical ethics principles of non-maleficence, liberty, justice, and beneficence. These rules need to be adhered to for years, if not decades, of an HRS, from

collecting data and developing algorithms to use them and keeping them running as intended. Without this foundation, HRS may inadvertently strengthen structural biases, or might view patients as statistical profiles that ignore the nuance of their individual circumstances.

Several attempts have attempted incorporating ethical reasoning by designing the mechanism of recommendation for clinical systems. S. Tiribelli and D. Calvaresi [63] do, for instance, examine the relationship between AI ethics and clinical decision support systems and recommend improvements to the practice. Such features include human-in-the-loop frameworks, differential access rules, and support for the ability to document the reasoning behind decisions. This not only establishes trust, but it also complies with the user standards and regulatory requirements easily. Similarly, Authors of [64] also confirm the need for responsible AI on personalization for patients. Fairness, transparency and empowerment of patients were named by them their most valuable factors.

Finally, there is an introduction of more formal and useful ethical models for the growth of ethical HRS. M. Kwiatkowska [65] presented a framework with multi-layered ethical design for medical recommendation systems. That means auditing processes after they are deployed, privacy controls to handle new information and consent protocols centered on the patient. Their blueprint unifies progress in technology with clinical integrity and social aspirations. Taken together, these efforts illustrate that HRS's morality is something not simply to tick off a list under regulators, but rather an important aspect of their long-term legitimacy and public trust.

G. Vision-Language Models (VLP) and Multimodal Transformers in Healthcare Recommender Systems

Implementation of vision-language models (VLPs) with multimodal transformers has recently expanded the potential of health recommender systems (HRS) towards visual content (medical images) and textual knowledge (clinical notes) integration. S. Ramedini et al. [13] developed a transformer-based encoder-decoder architecture for medical report generation based on radiology images which provides a dual-pronged treatment in the form of diagnosis explanation and case-based recommendations. Similarly, Q. Li et al. [66] also investigated that general-use VLPs such as BLIP-2 and OFA were applicable in medical areas after being trained. They realized high zero-shot performance and found ways to make report-grounded recommendations without retraining the VLPs for the new job.

A medical-specific vision transformer named MedViT [67] demonstrated that medical-specific priors and unified pretraining strategies improve both classification and captioning performance in multiple medical image-based problems. This builds a strong bridge between learning the representation of health information and directing a patient's clinical journey. More sophisticated ones extend on this idea by making the semantic understanding over time by reasoning together text and images in a meaningful manner and using relevant reasoning that is useful there. W. Huang et al. [68] proposed a key-semantic report refinement approach in a vision-language base model. This would enhance learning by extraction of progressively relevant information from long radiology stories. This helps for the remaining jobs (disease classification, phrase

grounding) and helps annotate well. This is a step in the right direction, to make personalized HRS a scalable process.

In addition, S. Schmidgall et al. [69] developed MAVL, an architecture of multimodal vision-language, and incorporate knowledge embeddings for the specific diagnosis of a specific patient. This helps ensure that AI can make recommendations. These systems are early in their development, but a promising area for future recommender systems involved in clinical decision support is their functionality to integrate semantic text processing and contextual imaging.

On the other hand, MedUnifier leverages this concept by introducing both language learning and visual generation into one architecture. Letting it learn in both directions [70], it makes the model better at writing clear clinical narratives straight from imaging inputs.

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

Real-world user studies are key for the validation of healthcare recommender systems (HRS) beyond algorithmic benchmarks into real world relevant metrics such as usability, acceptance and clinical impact. Unlike offline evaluation where retrospective datasets are used, real-world deployment includes the contextual feedback of the patients or caregivers and healthcare providers by observing the system under environmental variation and human behavior.

More recent applications have focused on creating mHealth test beds for validation of HRS in the real world. For instance, Authors in [71] developed and evaluated a smartphone application for individualized nutritional advice. They collected patient feedback on recommendation relevance and interaction flow based on a usability study and pilot deployment, which uncover the fundamental design space for dietary compliance support. Similarly, Researchers in [72] incorporated in-the-moment sensing from wearables within a mobile HRS for providing lifestyle suggestions. Their assessment included technical performance, usability, and context-dependency during natural use.

P. Alves et al. [73] employed a Dockerized microservices-based mobile health recommender, evaluated in a cluster randomized trial targeting prehypertension. The system provided protocol specific behavioral change messages and emphasized optimal health behavior adherence, which was longitudinally monitored to access BP outcomes and adherence improvement. Likewise, Authors of [74] studied an obesity risk reduction recommender for high-risk women, with long term engagement tracking and user reported outcomes in real deployment.

Some researches have been also conducted related to interactive/human-in-the-loop recommender paradigms, especially in eHealth context where dynamic contextual changes are necessary. D. Diyasena et al. [75] presented a human-in-the loop design that allows medical domain experts to interactively refine their recommendations while using the system. Their quantitative study demonstrated enhancements in system trust and perceived decision support effectiveness when clinicians were co-designers of the logic. In mental health use, F. Gräßer et al. [76] implemented a therapy recommender system based on patient-therapist interactions. They evaluated user satisfaction, perceived personalization and explainability in

real treatment sessions, suggesting that emotional context as well as trust are crucial.

Senior medication adherence brings its own unique issues, for which recommender systems are designed to minimize. Researchers in [77] studied a mobile medication support application through feedback from the elderly and clinicians within a clinical care environment. System recommendations were evaluated in terms of interpretability, accessibility and engagement in accordance to nursing guidelines revealing the importance of these factors for an elderly population. Although RCTs and longitudinal deployments are resource-heavy, they offer unparalleled grounding in the ecological validity and human-centered performance of HRS. New evaluation approaches are integrating passive sensing, clinician feedback and patient-reported outcomes to provide a full picture of the effect on the system. Future work should also retain an emphasis on co-design principles, trust calibration and ongoing feedback loops to ensure that HRS deployments do not lose their ethical mooring but continue to be clinically useful.

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

This section thoroughly examines the leading advancements in building Convolutional Neural Networks (CNNs) into healthcare recommender systems. A comprehensive view of the current study of medical settings can be produced by considering research trends combined with system architecture analysis, dataset, and metric evaluation. The analysis shows important obstacles present in existing research that help researchers develop future innovations in this field.

A. Research Trends and Frameworks

In line with the summary section 5, in this sub-section we describe concrete research trends and illustrative paradigms for the application of CNNs in healthcare recommender systems. Instead of returning to the general motivations, this section shows how CNN-based architectures, multimodal fusion, and hybrid learning architectures have already undergone to evolve in recent works to provide personalized, scalable and data-oriented healthcare recommendation. The convergence of deep learning and healthcare recommender systems has gained momentum in recent years with the emergence of Convolutional Neural Networks (CNNs) and their capacity for data extraction. Indeed, interest on hybrid architectures and data-rich approaches integrating medical imaging, clinical records as well as behavioral data has begun to grow rapidly, which will enable us to provide user-oriented and scalable healthcare assistance that is personalized for the consumer.

M. C. Comes et al. [78] used CNNs with transfer learning to analyze breast DCE-MRI imaging for early prediction of neoadjuvant chemotherapy response. The embeddings, although diagnostic-oriented, also demonstrate that pathology-driven understanding can be included in personalized healthcare recommendation processes. B. Ihnaini et al. [9] concentrated on chronic illness treatment and proposed to enhance the appropriateness and the efficiency of the recommendations based on deep ensemble techniques including CNNs for the diabetic patient.

Mahyari et al. [27] proposed a dynamic CNN-fueled recommendation system within mHealth for exercise coaching. Their model is on continuous learning with expert feedback which allows for dynamic customization. A. Nilla and E. B.

Setiawan [79] investigated a CNN-based hybrid health recommendation system based on collaborative filtering and content-based methods. Their universal structure enables the integration of multimodal information for tailored health content recommendations.

N. K. Al-Qazzaz et al. [80] devised hybrid autism spectrum disorder diagnosing system which integrates pre-trained CNN and K-nearest neighbor (KNN) classifiers for symptom severity classification of the patients and presented how pattern recognition helps personalized therapy planning. M. Arakeri et al. [81] introduced DeepReco a CNN-integrated collaborative filtering approach specially designed for mobile health recommender systems.

S. Bourhim et al. [82] presented the DGCF model, which is a hybrid of graph neural networks (GNNs) and convolutional filters that is proposed to enhance user-item interactions in health recommender pipelines. Fallahi and Mohammadzadeh [8] showed up and summarized CNN-based deep learning methods for collaborative filtering systems that has been introduced in the literature and also pointed out their capability of handling sparsity and data complexity-related behaviors in such systems.

K. Karre [83] return to some literature with their approach to treatment prioritization with severity analysis for children diagnosed with autism, and demonstrated the generalizability of CNNs across age groups and cognition problems. Iwendi et al. [84] created an IoMT-based CNN-based recommender system

suitable for diet personalization. Their work also demonstrated that patients and nutritional data with CNNs were utilized for the purpose of better care of individuals with chronic diseases. A. P. Ponselvakumar et al. [85] have presented a precision medicine recommender system to match patients to clinical trials with multimodal CNN inputs.

Y. S. Cho and P. C. Hong [86] illustrated that CNN based diagnostic models are capable to scale up malaria screening in low resource settings, indicating what a potential role it offers for AI tools in strengthening public health efforts and promoting equity across low-resource settings. Y. Cai et al. [87] also undertook a scoping review for HRS from 2010 to 2022, and observed a gradual trend in the integration of CNN and deep learning. B. Subramanian et al. [88] recently applied their emotion recognition framework in a mobile domain for the detection of mental wellness issues and applied CNN feature maps to identify the patient mood states.

Overall, more and more, we are seeing work in which recommendation mechanisms are meshing across domains, including visual analytics, behavioral inference/machination, sensor data, and structured medical records into CNNs. These approaches represent a shift towards more holistic, interpretable, and responsive healthcare recommender ecosystems. Table I presents an overview of recent research trends and frameworks in CNN-based healthcare recommender systems.

TABLE I. Research Trends and Frameworks in CNN-Based Healthcare Recommender Systems.

Research Trend / Area	Core Idea	Representative Works	Key Contributions / Implications for HRS
CNNs for Medical Imaging-Driven Recommendation	Using CNN + transfer learning to extract diagnostic features for treatment-oriented recommendation	Comes et al. [78]	Early chemo-response prediction; pathology-driven embeddings that support personalized treatment recommendations
Deep Ensembles for Chronic Disease Management	CNN-based ensemble frameworks to enhance recommendation accuracy for chronic patients	Ihnaini et al. [9]	Improved appropriateness and efficiency of diabetes-related recommendations using multi-model fusion
Dynamic & Continually Learning CNN Recommenders	Incorporating continuous feedback loops into CNN-based RS for adaptive health coaching	Mahyari et al. [27]	Real-time customization of exercise recommendations through expert feedback and continuous learning
Hybrid CNN + CF/Content-Based Architectures	Combining CNN feature extraction with collaborative filtering and content-based reasoning	Nilla & Setiawan [79]	Multimodal integration for tailored health content; versatile architecture for generalized health RS
CNN-KNN Hybrids for Disorder Severity Analysis	Leveraging CNNs with classical ML for symptom severity classification	Al-Qazzaz et al. [80]	Personalized ASD therapy planning using CNN-driven pattern recognition and KNN-based classification
CNN-Enhanced Collaborative Filtering for mHealth	Fusion of CNN feature extractors with CF pipelines for mobile environments	Arakeri et al. [81] (DeepReco)	Mobile-focused RS with CNN-augmented user-item representations
CNN + GNN Hybrids for Structural Interaction Modeling	Applying convolutional filters to graph neural networks to strengthen user-item modeling	Bourhim et al. [82] (DGCF)	Enhanced interaction modeling in health RS through deep graph + CNN fusion
CNN-Based Methods for CF Robustness	Addressing sparsity and behavioral complexity with CNN-based collaborative filtering	Fallahi & Mohammadzadeh [8]	Theoretical foundation showing CNN effectiveness for sparse, complex CF health datasets
CNNs for Autism Treatment Prioritization	CNN-driven severity and cognitive-level analysis to guide therapeutic interventions	Karre [83]	Demonstrates CNN generalizability across demographics for personalized ASD treatment ranking
IoMT-Driven CNN Recommenders for Diet Personalization	Integrating IoT medical devices with CNN-powered dietary decision support	Iwendi et al. [84]	Personalized nutritional recommendations for chronic illness via IoMT data + CNN processing
Precision Medicine with Multimodal CNN Inputs	Using CNNs to match patients to individualized clinical trials or therapies	Ponselvakumar et al. [85]	Personalized clinical trial matching using multimodal medical inputs
CNNs for Low-Resource Public Health Support	Deploying CNN-based diagnostics in rural or resource-limited environments	Cho & Hong [86]	Scalable malaria screening and equitable public-health recommendation support
Long-Term Trend Analysis of CNN Adoption in HRS	Reviewing evolution of CNN/HRS integration over a decade	Cai et al. [87]	Observed acceleration of CNN-based approaches in HRS (2010–2022)

CNN-Based Emotion Recognition for Mental Health RS	Using CNN feature maps to detect emotional states for wellness recommendations	Subramanian et al. [88]	Mood-aware RS for mobile mental health monitoring
--	--	-------------------------	---

B. Representative Datasets

Effective robust and diverse datasets form the basis for all successful CNN-based healthcare recommender systems. These datasets validate and verify the generalizability and robustness across different populations and clinical contexts of recommendations. CNNs have primarily been trained for clinical tasks including disease detection and visual feature extraction, mainly based on medical imaging datasets. These include ChestX-ray14 and CheXpert that have gained interest as good tools for screening thoracic disease and MIMIC-CXR that integrates chest radiographs with detailed clinical reports, increasing the diagnostic efficiency and diagnostic interpretability [10]. COVIDx was an invaluable tool in the pandemic campaign for training CNNs to spot viral pneumonia from X-rays [89].

Aside from imaging, structured clinical datasets, such as MIMIC-III, possess rich information about laboratory values, medication history, ICU admissions, and physician notes to be used in the training of patient-relevant CNN recommender systems [90]. NHANES (National Health and Nutrition Examination Survey) has brought nutritional, behavioral, and demographic information, useful in lifestyle recommendation models [91]. The UCI Diabetic dataset has been a standard reference for CNNs in predicting chronic diseases and improving care [92].

Wearable sensor datasets such as WESAD, which includes physiological signals related to stress and affective states, have enabled CNNs to personalize mental health interventions. Similarly, ECG-based datasets like MIT-BIH provide valuable time-series data for cardiovascular monitoring and tailored health feedback [93]. For emotion-driven applications, datasets such as AffectNet and other custom webcam-based emotional datasets support the development of CNNs in real-time mood recognition and therapy support systems [94].

Current datasets appear to possess the necessary utility, though at the same time they have several issues. They have inadequate demographic diversity, which leads to biased model output that may generalize to new patient populations. Others are unimodal, offering only a single type of input (e.g., images without corresponding textual or biometric data), thus limiting their effectiveness in training hybrid models [95]. Privacy regulations also hinder access to longitudinal and high-resolution patient data, creating a gap between model development and real-world deployment [96].

Future progress in CNN-augmented healthcare recommenders depends on the creation and sharing of multimodal, ethically curated datasets that reflect the complexity, diversity, and longitudinal nature of real clinical environments.

C. Performance Metrics

Evaluating the performance of CNN-empowered healthcare recommender systems involves a diverse array of quantitative metrics tailored to different tasks, including classification accuracy, prediction error, ranking effectiveness, and user satisfaction. These metrics help ensure both the robustness of

underlying models and the utility of generated recommendations in practical clinical scenarios [45].

For classification tasks, which are central in medical image analysis and disease diagnosis, the most common metrics include [89], [97]:

1. Accuracy measures the proportion of correctly classified instances. Equation 7 demonstrates accuracy calculation.

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

2. Precision indicates the proportion of true positives among all predicted positives. Equation 8 shows the precision calculation.

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

3. Recall (Sensitivity) reflects the proportion of actual positives correctly identified. Equation 9 shows the recall calculation.

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

4. F1-Score is the harmonic mean of precision and recall. Equation 10 shows the F1-score calculation.

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

5. Specificity (SP) is the proportion of true negatives correctly identified. Equation 11 shows the specificity calculation.

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

6. ROC-AUC: The area under the receiver operating characteristic curve, reflecting the model's ability to discriminate between classes across thresholds [33].

In the context of recommendation systems, where personalized ranking is key, performance is measured using:

7. Mean Absolute Error (MAE): Measures the average magnitude of errors between predicted and actual values. Equation 12 shows the MAE calculation.

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

8. Root Mean Squared Error (RMSE): Computes the square root of the average squared differences. Equation 13 shows the RMSE calculation.

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

9. Mean Reciprocal Rank (MRR): Evaluates ranking by measuring the position of the first relevant item. Equation 14 shows the MRR calculation.

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

Real-world systems, such as DeepReco [45], emphasize the importance of low latency, high interpretability, and adaptability to dynamic user needs. As recommender systems increasingly integrate sensitive health data, fairness-aware evaluation (as outlined by L. Bruni et al. [98] and model transparency tools such as Grad-CAM [18] and SHAP [99] are gaining traction in CNN performance assessments.

D. Limitations in Current Studies

Although CNNs make a positive impact on integration of healthcare recommender systems, there are still a number of limits that need to be overcome in most of the studies to date. That large majority are using very small datasets, very small samples of data not all with even slightly differing demographic representation in the samples or also geographic diversity. This frequently produces models that work for controlled settings but generally don't generalize to larger, real-world populations. Poor representativeness of data from differing age-groups, ethnicities, and socio-economic backgrounds can introduce bias in recommendations which can lead to accidental bias in making recommendations, especially concerning in clinical application [21].

Yet another general issue is the poor interpretability of deep CNN architectures. CNNs are able to learn sophisticated models for medical data (for example, complex representations), but the mechanisms being used for decision making are complex and thus very opaque. In a clinical setting where accountability and transparency are key, this black-box nature leads to concerns. Though Grad-CAM, SHAP and others have been added to increase model explainability, their use in such settings has been limited, and other solutions to explainability issues are still in limited use; many still operate without interpretable outputs [40], [42] which limits general trust between clinicians and end-users.

And besides, many existing systems have not yet been tested in real time in clinical settings. Evaluations are usually retrospective and performed under optimal laboratory conditions, with scarce validation in noisy and changing healthcare settings [85]. The gap between trial success and clinical implementation remains a significant obstacle to system adoption. EHR integration, regulatory harmonization, and seamless workflow embedding have been unexplored or underdeveloped [81].

For more data privacy and security aspects also hinder the application of CNN-driven healthcare recommenders. Models built on patient information are required to comply with stringent patient privacy regulations such as HIPAA and GDPR. Nevertheless, only a handful of research includes privacy and security principles such as federated learning or differential privacy, and the data anonymization approach is frequently insufficiently performed by clinical benchmarks [100].

Finally, the fairness and ethics are generally underestimated in performance assessments. Models may inadvertently bias majority groups or oversimplify minority populations,

contributing to health disparities [4], [62]. Measures of fairness-aware evaluation, bias correction, and fairness-focused design principles to date have yet to be discovered in this area and should be highlighted in future.

VI. SYSTEM-LEVEL ENHANCEMENTS AND DEPLOYMENT CONSIDERATIONS

A. Federated Learning and Privacy-Preserving Healthcare Recommender Systems

Maintaining patient privacy in healthcare recommender systems (HRS) while still delivering good performance over heterogeneous datasets is increasingly relevant. Federated Learning (FL) is a key method that has emerged to address this need, allowing multiple actors to train models in unison without relaying raw data. FL takes advantage of local devices or medical centers to decentralize learning, in contrast to traditional centralized architectures. The private nature of this means it decreases the communication costs and is crucial especially in heavily regulated areas like medical diagnosis, treatment recommendation and behavior monitoring.

M. Adnan et al. [101] developed a simultaneous multitask deep learning model for the early prediction of breast cancer. This model separates patient data from various hospitals. The method adapts local models to each hospital's unique patterns while developing a globally-standardizable model. It provides a greater consistency of predictions for groups, while ensuring data privacy. To mitigate data heterogeneity, S. Che et al. [102] introduced FedHealthRec, a federated health recommender framework which is based on personalized collaborative filtering. The model adapts local recommendation weights to specific user behaviors and preferences. Its contribution is to integrate a privacy-preserving knowledge distillation module to address the cold-start challenge and to provide fast adaptation process for new users without exposing sensitive information.

M. H. Nguyen et al. [103] developed a federated medical dialogue system which was designed for the management of chronic diseases. Their architecture combines on-device transformer-based language models with server-side policy learning, allowing real-time privacy-preserving medical guidance. Transformer distillation retains performance with a lack of computational resource and an inconsistent connection which are usual constraints in mobile healthcare. S. D. N et al. [104] found FedCure, a federated recommendation system that balances the sharing of global knowledge with local user embeddings to provide tailor-made health advice. Their model considers device heterogeneity and personalized recommendations based on the client's computing power and frequency of participation.

In an aligned line, E. Mantey et al. [105] proposed a blockchain-assisted federated recommender (using Hyperledger Fabric). This hybrid framework is suitable for secure storage of data as well as explainable recommendations in mobile health platforms, even if data inconsistency and/or clients who are missing. S. Gupta et al. [33] furthermore introduced FedGraphRS, a graph-based federated recommender that protects data through homomorphic encryption and inscribes user-item interactions in a GNN. This is an extension to privacy-centric designs. In this system, people from different institutions learn from each other in the same graph structure, but while decentralized training occurs. E. A. Mantey et al. [106] further

developed this idea including LSTM and GRU-based temporal models to a blockchain to help govern chronic diseases. They employ a distributed ledger to hold electronic health records (EHRs) and recommend follow-up actions based on accurate predictions of how health will evolve across time.

These inputs illustrate how FL is turning HRS into decentralized systems concerned with user privacy. Problems of this kind persist, however. However, there are several remaining ones to be solved. These include non-IID data distribution, model drift, and client fairness. The tension between personalization and generalization is further explored in this article. These are all active fields of study in that area of change for which there are still many more stages.

B. Edge Deployment and IoMT in Healthcare Recommender Systems

With the exponential growth of the Internet of Medical Things (IoMT) and edge computing, there has been a paradigm shift in healthcare recommender systems that supports real-time, privacy-preserving and context-aware decision making. In conventional cloud-models, delays and risk of data transmission especially in the mission-critical healthcare applications are some issues. To combat this, federated learning, on-device inference, and heterogeneous edge optimization have been employed to tailor recommender pipelines for real-world scenarios.

For instance, A. S. Fathima et al. [107] introduced a federated edge-cloud model to dynamically offload medical-resource recommendations among 5G-based IoMT devices. It serves to facilitate scheduling efficiency; their model integrates edge level pre-processing and blockchain-enhanced prioritization. M. H. Nguyen et al. [103] designed a heterogeneous BlockNet based federated learning structure for diagnostic recommendation which enables various edge devices with different computational abilities to run customized deep neural networks. This method achieves significant improvement with respect to efficiency, data privacy and scalability. Sensor integration is a key for emotion-aware healthcare services. B. Subramanian et al. [88] introduced a digital-twin based deep learning framework for real-time emotion recognition, which could monitor psychological state and provide timely suggestions. This system runs on the webcam input on commodity edge devices and form a scalable foundation for personalized mental health support.

Edge intelligence can benefit even further from on-device federated learning. An IoT-based Half-and-Half-Attention Enhanced Deep Collaborative Transformer with federated learning was proposed by A. Alqhatani and S. B. Khan [108] which enhanced personalization accuracy and reduced latency for large scale recommendation tasks—the playground in whose evaluation can potentially guide the design of future IoT driven healthcare recommenders systems. Recent work on lightweight deep residual learning shows in addition some efficient representation learning techniques that can be used to help compact edge-deployed vision modules for healthcare recommender-using systems [20]. Disease-specific IoMT recommender technologies have been introduced as well. One possible path towards understanding concept space-based recommendations is the [21] management framework for chronic diseases, which illustrates how an IoT-enhanced healthcare environment may rely on electronic health records in

order to enhance diagnosis and management of chronic disease conditions through similarity-driven predictions. Moreover, we just only consider the structural modeling methods have broader recommendation logic. S. Bourhim et al.[82] proposed a deep collaborative recommender system that is community driven, models complex user–item interactions based on the principles of graph-neural-network, and provided rationales for how structural relationships can augment robustness in healthcare RS architectures.

Tailored treatment support is another crucial level. Alam [22], an algorithm named Personalized Multimodal Treatment Response System (PMTRS) was proposed for a personalized prescription by fusing the imaging, clinical, or demographic information at a multimodal feature fusion level. These people are different from those at the edge of learning, making it possible to deploy even further: scattered edge inference techniques. The work in [109] discusses how deep-learning models can be partitioned and optimized over heterogeneous edge devices, highlighting architectural design principles applicable for future hierarchical healthcare-recommender pipelines. S. D. N et al. [104] presented FedCure, a heterogeneous-aware personalized federated learning technique that gains in terms of accuracy and latency in IoMT settings considering applications for diabetes monitoring, retinopathy screening, maternal-health assessment, and remote patient monitoring.

Cumulatively, these developments highlight progression from static centralized healthcare recommendation models to edge-AI ecosystems that are distributed, scalable, and resilient. However, there are still some challenges in the federated learning context, such as energy constrained devices, drift model under the federated settings and cross device variation that require further exploration under future research.

VII. CHALLENGES AND OPEN RESEARCH ISSUES

Despite the growing integration of Convolutional Neural Networks (CNNs) into healthcare recommender systems (HRS), several unresolved challenges limit their widespread adoption in clinical settings. These challenges span technical limitations, privacy and fairness considerations, and operational constraints—each representing an active area of ongoing research.

A. Data Heterogeneity and Non-IID Distributions

Healthcare data are naturally heterogeneous as a result of differences in demographics, imaging devices, clinical procedures and institutional protocols. In federated systems, such non-IID behavior leads to poor quality gradients and unstable convergence. Recent federated recommender systems, e.g., FL-based personalized recommenders which employ meta-learning methods (e.g., REPTILE-style first-order adaptation), achieve better personalization on decentralized data such as PrivRec and DP-PrivRec. Similarly, graph-based models like FedGraphRS [33] exploit the structural relationship for cross-client generalization. Hybrid approaches like FedCure [104], which involve device-wise specialization and adaptive aggregation, yet still face unreliable clients and non-uniform participation. In general, addressing extremes of data imbalance, sparsity in interactions and cross-institution variability remains a major challenge.

B. Personalization vs. Generalization Trade-offs

When making personalized health recommendations, comorbidities, and the history of behavior, lifestyle or long-term tendencies should be taken into account. Nevertheless, commonly the personalization hurts the global generalization of models. Heterogeneous FL systems with block structure such as BlockNets [97. pdf] show that whereas a focus on local relevance can be enhanced by modular specialization, cross-population robustness can be potentially weakened. Similarly, FedCare [67. pdf], in that it adds local tuning to boost the accuracy of individual-level but becomes sensitive to participation frequency. Time and behavior-based personalization approaches (e.g., applied Context-Aware FL recommenders [110] illustrate the importance of agile adjustment. Future architectures should obtain an equilibrium between personalized relevancy and stable global generalization with different patient profiles.

C. Multimodal Fusion and Temporal Misalignment

Multimodal Recent health care HRS are increasingly recording and managing multiple modalities, ranging from medical images, clinical notes to physiological signals and demographics. For example, ViGPT2 transformer-based fusion architectures [111], a strong image-text alignment and co-attention framework is presented for medical report generation. Co-attention and cross-modal transformers [112], [113] demonstrate how textual and visual features can be fused successfully. However, the multimodal CNN-based HRS have several issues:

- vitals and imaging with mismatched temporal sampling,
- missing or noisy modalities,
- late-arriving sensor data,
- and latency-sensitive cross-attention computation.

These difficulties are heightened when complex systems have to work in real time process or over extended patient pathways.

D. Explainability and Clinical Trust

Transparent and interpretable recommendations are required in clinical decision making. Interpretability tools such as Grad-CAM [18] and attention-based interpretability in vision-language models [66] is a mechanism to see evidence regions and their textual rationales. Graph explainers such as node/edge attribution methods [45] Additionally support graph-structured HRS transparency. However, these methods may be inadequate for workflows with high stakes. As demonstrated in the co-design frameworks [41], doctors need layered explanations that integrate evidence at the model-level (e.g., by referencing visual signs or symptom clusters) with a decision rationale (e.g., why a medication is appropriate). Human-centered, rigorous validation of explanations as a step towards engineering clinical trust is still outstanding.

E. Fairness, Bias, and Data Imbalance

Under-representation of specific populations, rare diseases or low-resource geographies gives rise to bias. Works such as LDW-CNN [20] deals with class imbalance and enhance predictive stability via optimization-based feature discriminants, very little has been done about demographic fairness and cross-

hospital equity. Some approaches involve reducing bias by learning with balanced training set and domain adaptation [114] covers distributional issues of medical CNN processing), there are still no standardized fairness metrics for clinical risk that nursing staff needs. There is an urgent need for methods of bias auditing, calibrated prediction and demographically-aware training.

F. Edge Deployment and Resource Constraints

Deploying CNN-based HRS at the edge introduces constraints in terms of energy, memory, and compute capability. The MedViT hybrid CNN-Transformer architecture [67] highlights the practical limitations of transformer inference on resource-constrained devices and demonstrates techniques for efficient attention. Distributed and partitioned inference strategies, reviewed in [115], further show how model segments can be offloaded across heterogeneous devices. FL-based IoMT systems—including PrivRec [33]), FedCure [104] , and IoMT-aware architectures [116]—demonstrate viable paths for efficient decentralized learning. Blockchain-enhanced decentralized coordination frameworks such as MedShare [100] offer additional robustness for distributed scheduling and secure data handling. Still, reliable edge operation under weak connectivity, hardware failures, and variable participation is an unsolved challenge.

G. Evaluation Gaps and Real-World Validation

Both ZoKrates and Scale outperform prior work by using synthetic data to perform FL. Uploaded methods such as PrivRec [33] and FedCure [104] focuses on performance in simulated FL settings and does not assess longitudinal integration or clinician-in-the-loop decision-making. The considerations for deployment testing in [27] importance is placed on integrating into the institutions, real-time constraint evaluation and end-user validation. Only a limited number of systems have been subjected to user studies, cross-hospital deployment trials, or are compared with physician recommendations. Standardized clinical assessment protocols and multi-center cooperation is needed to transfer research models into clinics.

While the CNN-based and multimodal HRS research is growing fast, a number of critical challenges are encountered in data heterogeneity, fair explainable personalization, multimodal real-time constraints managing, edge/IoMT deployment and robust validation in real clinical settings. Solving these problems will demand interdisciplinary breakthroughs in machine learning, medical informatics, HCI and clinical workflow design. Table II summarizes the aforementioned challenges.

TABLE II. Key Challenges in CNN-Based Healthcare Recommender Systems.

Challenge Area	Core Issue	Representative Works	Open Limitations
Data Heterogeneity & Non-IID Distributions	Highly heterogeneous healthcare data across demographics, devices, and institutions lead to unstable gradients and weak FL convergence	PrivRec, DP-PrivRec, REPTILE-style meta-learning, FedGraphRS [33], FedCure [104]	Persistent data imbalance, sparse interactions, unreliable clients, non-uniform participation, and cross-institution variability
Personalization vs. Generalization Trade-offs	Personalized tuning improves individual relevance but reduces population-wide robustness in heterogeneous FL systems	BlockNets (97), FedCare (67), Time/Behavior-based personalization, Context-Aware FL Recommenders [110]	Difficulty balancing global generalization with personalized relevance; sensitivity to participation frequency; risk of overfitting to patient-specific patterns
Multimodal Fusion & Temporal Misalignment	Medical images, clinical notes, vitals, and sensor data arrive asynchronously or with noise, creating misalignment and latency in multimodal HRS	ViGPT2 [111], Co-attention / cross-modal transformers [112], [113]	Temporal mismatch, missing/late modalities, latency-sensitive attention, and unreliability in real-time or long-term patient pathways
Explainability & Clinical Trust	Current interpretability methods fail to provide layered, clinically meaningful explanations required for high-stakes decisions	Grad-CAM [18], VLM interpretability [66], Graph explainers [45], Co-design frameworks [41]	Lack of human-centered explanations; inadequate validation; limited integration of visual evidence + rationale; insufficient trust for clinical deployment
Fairness, Bias & Data Imbalance	Under-representation of demographic groups and rare diseases leads to biased predictions and unequal clinical outcomes	LDW-CNN [20], Domain adaptation for fairness [114]	No standardized fairness metrics; demographic bias remains unaddressed; limited calibrated and equitable prediction frameworks
Edge Deployment & Resource Constraints	Running CNN/Transformer HRS models on low-power or intermittently connected devices imposes memory, compute, and energy constraints	MedViT [67], Distributed inference [115], PrivRec [33], FedCure [104], IoMT-aware RS [116], MedShare [100]	Latency, battery limits, connectivity issues, hardware failures, inconsistent client participation, and unreliable operation in remote settings
Evaluation Gaps & Real-World Validation	Heavy reliance on synthetic FL datasets and simulations instead of real clinical workflows or multi-center studies	ZoKrates, Scale, PrivRec [33], FedCure [104], Deployment analysis in [27]	Lack of clinician-in-the-loop testing; few real-world trials; weak longitudinal evaluation; no standardized assessment protocols across hospitals

VIII. CONCLUSION

This review has explored the intersection of computer vision, convolutional neural networks (CNNs), and healthcare recommender systems (HRS), highlighting their advancements and the emerging synergy among them. Through an in-depth analysis of recent literature, we observed that CNNs have significantly improved the ability of medical systems to extract meaningful patterns from complex visual data, making them an invaluable component in diagnostic and treatment support tools. These models could open new horizons when combined with recommendation systems to deliver promising personalized healthcare solutions.

Although researchers have made significant findings, they still face challenges in expanding use at work and making systems easier to understand within real-world clinical workflows. AI systems used in healthcare today create major substantial ethical concerns because decisions made by them need to be fully explained and developers need to face legal responsibilities. The healthcare systems need to protect sensitive patient information because they need access to medical data in their operations. Making sure new solutions follow data protection rules and build user trust must remain top priorities during their creation.

This review derives its significance from its harmonious integration of two rapidly evolving fields: CNN-based computer vision and healthcare recommender systems, enabling researchers and developers to gain a comprehensive view of their convergence. This paper provides a foundation for advancing interdisciplinary progress in intelligent healthcare by analyzing key trends, along with research

frontiers and new research possibilities. It also provides novel insights for researchers and developers by presenting important trends in combined fields, enabling them to bring about conceptual shifts that contribute to advancing healthcare personalization to new potentials. The main novelty of this work lies in linking disparate study fields to show how their common power can transform healthcare service delivery. This review functions beyond a modern state-of-the-art assessment by establishing a forward-thinking guide to creating AI-powered healthcare recommender systems which remain ethical and operationally effective.

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. Ramedini, 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. Thiagarajan, 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. Pirololi, 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, D. 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. Gasmli, N. Azizi, and N. Dey, "Explainable AI and deep learning models for recommender 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. Muqet, 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. Wiczkowska, 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.combiomed.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," *J. 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 programmes 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.