STUDENT NAME: RACHNA SETHI

SUPERVISOR NAME: Dr. MONICA MEHROTRA

TOPIC: Recommender Systems for Healthcare Applications

DEPARTMENT: Department of Computer Science, Faculty of Sciences

FINDINGS

This research expands the traditional scope of recommender systems (RS), which are predominantly used in domains like e-commerce, entertainment, and social media, into the healthcare sector. By doing so, it highlights the transformative potential of RS in delivering personalised healthcare services. The study identifies and analyses two critical sub-domains within healthcare where RS can have a meaningful impact:

- **Image-based medical diagnosis**, where timely and accurate disease detection is crucial.
- **Physical activity recommendation**, which promotes preventive care and lifestyle management through personalised guidance.

Application of Deep Learning for Medical Image-Based Diagnosis

A key contribution in this area is leveraging deep learning and transfer learning techniques for building diagnostic recommender systems using medical images. Four prominent CNN architectures were applied to classify COVID-19 cases from chest X-ray images. These models, initially trained on the large-scale ImageNet dataset, were fine-tuned using smaller, domain-specific datasets, thus demonstrating the power of transfer learning in medical imaging tasks. InceptionV3 emerged as the most accurate and reliable model, outperforming existing diagnostic models in terms of classification accuracy. This contribution underscores the feasibility of using pre-trained deep networks to build highly accurate diagnostic tools even when labelled medical data is limited—a common scenario in real-world healthcare settings.

Another challenge addressed is data imbalance in medical datasets, particularly where some disease classes are underrepresented. To counter this a Generative Adversarial Network (GAN)-

based model was developed to synthetically generate COVID-19 positive chest X-ray images. These synthetic images were used to augment three different imbalanced datasets, improving the diversity and representativeness of the training data. Extensive experimentation demonstrated that CNNs trained on these augmented datasets showed enhanced sensitivity and robustness, even under varying degrees of class imbalance.

Physical activity recommendation

Recognising the lack of structured physical activity guidelines for disease-specific management, this study introduces a novel approach for developing a Personalised Physical Activity Recommender System using Natural Language Processing (NLP) techniques. It utilises yoga as a therapeutic modality, capitalising on its well-documented health benefits and widespread acceptance in both preventive and curative health contexts. A comprehensive textual database of yoga asanas and their associated health benefits was curated. State-of-the-art Natural Language Processing (NLP) models were trained and fine-tuned to predict suitable yoga asanas based on textual descriptions of health conditions. The models achieved high prediction accuracy, effectively mapping health conditions to appropriate therapeutic yoga practices. This contribution demonstrates how NLP can be used to empower users with personalised physical activity recommendations based on health-specific textual inputs.

Integration of AI Techniques for Healthcare Solutions

A unique aspect of this research is the integration of diverse AI paradigms—deep learning (CNNs, GANs) and NLP models—within a unified healthcare recommender framework. This research contributes to the growing body of work at the intersection of AI and healthcare by enhancing the accuracy and reliability of diagnosis through intelligent image-based systems. It also contributes to promoting preventive care via personalised activity recommendations.