

# AI-Powered Predictive Modeling for Early Diagnosis of PCOS and Thyroid Disorders Using Electronic Health Record Data

Mohammad Fayazuddin<sup>1</sup>, Zeba Tanwar<sup>2</sup>, Deepak<sup>3</sup>, Maharajan. S<sup>4</sup>, Suruthi Suga. A<sup>5</sup>

<sup>1</sup>*M. Pharm, Kakatiya University, Warangal*

<sup>2</sup>*BSc Paramedical (OT & Anesthesia), Topiwala National Medical College and Nair Hospital- MUHS Mumbai, Maharashtra*

<sup>3</sup>*Pharm.D, MMCP, Maharishi Markandeshwar, Mullana Ambala*

<sup>4,5</sup>*B. Pharm, K.K. College of Pharmacy, TN MGR University, Chennai*

**Abstract**—Polycystic Ovary Syndrome (PCOS) and thyroid disorders are among the most prevalent endocrine conditions worldwide, yet both continue to be significantly underdiagnosed and often identified only after prolonged clinical progression. This delay in diagnosis is largely attributed to the non-specific and overlapping nature of symptoms, including fatigue, weight fluctuations, menstrual irregularities, and metabolic disturbances, which frequently lead to fragmented clinical evaluation and missed opportunities for early intervention. In recent years, the widespread adoption of Electronic Health Records (EHRs) has resulted in the accumulation of extensive longitudinal patient data, encompassing laboratory results, clinical observations, and treatment histories. Despite this, much of the data remains passively stored and underutilised for proactive disease detection. Advancements in artificial intelligence (AI), particularly in machine learning and deep learning, offer a promising pathway to transform this paradigm by enabling the identification of complex patterns and predictive signals embedded within EHR data. This paper explores the application of AI-powered predictive modelling for the early diagnosis of PCOS and thyroid disorders, focusing on the integration of multi-dimensional clinical data and temporal health trajectories. By leveraging structured and semi-structured EHR datasets, predictive models can detect subtle biochemical and physiological changes that precede overt clinical manifestation, thereby facilitating earlier diagnosis and timely clinical intervention. The study discusses various modelling approaches, including traditional statistical methods, ensemble learning techniques, and advanced deep learning architectures such as hybrid convolutional neural network–long short-term memory (CNN-LSTM) models, which are particularly effective in capturing both cross-sectional feature interactions and longitudinal trends. In addition, the importance of explainable artificial intelligence (XAI) is emphasised to ensure that model predictions are interpretable, clinically relevant,

and aligned with established medical knowledge, thereby enhancing clinician trust and supporting informed decision-making. Furthermore, the paper critically examines the challenges associated with implementing AI-driven predictive systems in healthcare settings, including issues related to data quality, interoperability, algorithmic bias, and ethical considerations such as patient privacy and regulatory compliance. It also highlights the potential of such systems to reduce diagnostic delays, improve disease management, and promote equitable access to healthcare, particularly in resource-constrained environments.

In conclusion, AI-powered predictive modelling using EHR data represents a significant advancement in the early detection of endocrine disorders. By shifting the focus from reactive diagnosis to proactive risk identification, such approaches have the potential to improve clinical outcomes, reduce long-term complications, and contribute to the broader goal of precision and preventive medicine.

**Index Terms**—Polycystic Ovary Syndrome, Thyroid Disorder, Electronic Health Records, Machine Learning, Deep Learning, Predictive Modelling, Clinical Decision Support, SHAP Explainability, Endocrine Disorders, Early Diagnosis

## I. INTRODUCTION

The global burden of endocrine disorders has grown steadily over the past few decades, reflecting not only changes in lifestyle and environmental exposures but also improvements in diagnostic awareness and reporting systems. Among these disorders, Polycystic Ovary Syndrome (PCOS) and thyroid dysfunction stand out as two of the most prevalent yet persistently

underdiagnosed conditions. Both disorders exert wide-ranging effects on metabolic, reproductive, and psychological health, often extending far beyond their primary physiological domains. Despite their clinical significance, early detection remains a major challenge, largely due to the subtle, heterogeneous, and often overlapping nature of their presentations.

PCOS is widely recognised as one of the most common endocrine disorders affecting women of reproductive age, with prevalence estimates ranging from 8% to 13% depending on diagnostic criteria and population studied. It is characterised by a constellation of features including hyperandrogenism, ovulatory dysfunction, and polycystic ovarian morphology. However, the clinical expression of PCOS is highly variable, with some individuals presenting with overt symptoms such as hirsutism and menstrual irregularities, while others may exhibit only mild metabolic disturbances or remain asymptomatic for extended periods. This heterogeneity complicates diagnosis and often leads to fragmented care, where individual symptoms are treated in isolation rather than recognised as part of a broader endocrine syndrome.

Thyroid disorders, encompassing hypothyroidism, hyperthyroidism, and subclinical variants, are similarly widespread and clinically diverse. The thyroid gland plays a central role in regulating metabolism, energy balance, and hormonal homeostasis, and even minor dysfunction can have significant systemic consequences. Hypothyroidism, for instance, is commonly associated with fatigue, weight gain, and depression, while hyperthyroidism may present with anxiety, weight loss, and cardiovascular abnormalities. Subclinical forms of thyroid disease further complicate the clinical landscape, as they often lack overt symptoms yet carry a risk of progression to more severe states. The prevalence of thyroid disorders is particularly high among women, and their incidence increases with age, making them a major public health concern.

A critical challenge shared by both PCOS and thyroid disorders is the delay between the onset of underlying physiological changes and formal clinical diagnosis. This delay is not merely a matter of timing; it has profound implications for patient outcomes. In the case of PCOS, prolonged exposure to hyperandrogenism and insulin resistance can increase the risk of type 2 diabetes, cardiovascular disease,

infertility, and endometrial cancer. Similarly, untreated thyroid dysfunction can lead to dyslipidaemia, cognitive impairment, reproductive complications, and, in severe cases, life-threatening conditions such as myxoedema coma or thyroid storm. Early identification and intervention are therefore essential to mitigate these risks, yet current diagnostic pathways often fail to detect disease in its initial stages. One of the primary reasons for delayed diagnosis lies in the reliance on symptom-driven clinical evaluation. Patients typically seek medical attention only after experiencing noticeable symptoms, which may already reflect advanced disease. Even then, the non-specific nature of many endocrine symptoms can lead to misinterpretation or under-recognition, particularly in primary care settings where time constraints and competing clinical priorities limit comprehensive evaluation. For example, fatigue and weight gain common complaints in hypothyroidism are frequently attributed to lifestyle factors, while menstrual irregularities in PCOS may be normalised or overlooked. This reactive approach to diagnosis inherently limits the ability to identify disease early and highlights the need for more proactive strategies. The increasing digitisation of healthcare systems offers a potential solution to this challenge. Electronic Health Records (EHRs) have become a cornerstone of modern healthcare, providing a comprehensive and longitudinal record of patient interactions, laboratory results, diagnoses, and treatments. Unlike traditional paper-based systems, EHRs enable the aggregation and analysis of large volumes of data across time and clinical contexts. This wealth of information holds significant untapped potential for improving disease detection and management. However, in practice, much of this data remains underutilised, serving primarily as a repository for documentation rather than a resource for predictive insight.

Artificial intelligence (AI), and particularly machine learning, has emerged as a powerful tool for extracting meaningful patterns from complex datasets such as EHRs. By analysing relationships between multiple variables simultaneously, AI systems can identify subtle signals that may not be apparent to human observers. In recent years, AI has demonstrated considerable success in a range of healthcare applications, including medical imaging, disease risk prediction, and clinical decision support. Its ability to process high-dimensional data and model non-linear

relationships makes it particularly well-suited for addressing the diagnostic challenges associated with endocrine disorders. In the context of PCOS and thyroid disorders, AI offers the possibility of shifting the diagnostic paradigm from reactive to predictive. Rather than waiting for symptoms to manifest, predictive models can analyse historical patient data to estimate the likelihood of disease development. This approach is especially relevant for conditions that evolve gradually over time, as is the case with many endocrine disorders. Hormonal fluctuations, metabolic changes, and subtle clinical signs may precede diagnosis by months or even years, and these patterns are often captured within EHR data. By leveraging this information, AI models can generate early warnings that prompt further investigation and timely intervention. A key advantage of using EHR data for predictive modelling is its longitudinal nature. Unlike cross-sectional datasets, which provide a snapshot of patient health at a single point in time, EHRs capture the dynamic progression of disease. This temporal dimension is critical for understanding endocrine disorders, which are characterised by gradual changes in physiological regulation. For example, a slow but consistent increase in thyroid-stimulating hormone levels may indicate emerging hypothyroidism, even if individual measurements fall within normal ranges. Similarly, progressive alterations in hormonal ratios or metabolic markers may signal the early stages of PCOS. AI models that incorporate temporal data are therefore better equipped to detect these trends and provide more accurate predictions. Despite its potential, the application of AI in healthcare is not without challenges. One of the most significant barriers to adoption is the issue of interpretability. Many advanced machine learning models, particularly deep learning systems, are often described as “black boxes” due to the difficulty of understanding how they arrive at their predictions. In a clinical setting, where decisions have direct implications for patient care, this lack of transparency can undermine trust and limit usability. Clinicians require not only accurate predictions but also clear explanations that align with established medical knowledge. As a result, there is growing interest in explainable AI techniques that provide insight into model decision-making processes. Another important consideration is the quality and representativeness of the data used to train AI models. EHR data are inherently heterogeneous, with

variations in data entry practices, measurement standards, and patient populations across different healthcare settings. Missing data, inconsistencies, and biases can all affect model performance and generalisability. For instance, a model trained on data from urban tertiary care centres may not perform equally well in rural or resource-limited settings. Addressing these issues requires careful data preprocessing, validation across diverse populations, and ongoing monitoring to ensure equitable performance.

Ethical and regulatory considerations also play a crucial role in the development and deployment of AI in healthcare. The use of patient data raises concerns about privacy, consent, and data security, particularly in the context of large-scale data analysis. Regulatory frameworks are evolving to address these challenges, but there remains a need for clear guidelines that balance innovation with patient protection. Additionally, the potential for algorithmic bias must be carefully managed to avoid exacerbating existing health disparities. Within this complex landscape, the integration of AI-powered predictive modelling into clinical practice represents both a significant opportunity and a formidable challenge. For PCOS and thyroid disorders, in particular, the potential benefits are substantial. Early detection can enable timely intervention, reduce the risk of complications, and improve quality of life for affected individuals. Moreover, predictive models can support clinicians by providing additional insights that complement clinical judgment, rather than replacing it.

This paper seeks to explore the application of AI-driven predictive modelling for the early diagnosis of PCOS and thyroid disorders using EHR data. It aims to synthesise current knowledge from the fields of endocrinology, data science, and health informatics to develop a comprehensive understanding of how these technologies can be effectively utilised. By examining the conceptual foundations, methodological approaches, and practical considerations involved, the study provides a framework for future research and implementation in this rapidly evolving field. In doing so, it also highlights the broader implications of AI in healthcare, particularly its role in advancing precision medicine and shifting the focus from treatment to prevention. As healthcare systems continue to evolve, the ability to anticipate disease and intervene early will become increasingly important. AI-powered

predictive modelling, when applied thoughtfully and responsibly, has the potential to play a central role in this transformation, offering new possibilities for improving patient outcomes and addressing some of the most persistent challenges in modern medicine.

## II. AI IN PCOS DIAGNOSIS

The application of artificial intelligence in the diagnosis of Polycystic Ovary Syndrome has gained considerable attention over the past decade, largely due to the inherent complexity and heterogeneity of the condition. Traditional diagnostic approaches rely on the Rotterdam criteria, which combine clinical, biochemical, and imaging findings. However, these criteria often fail to capture the full spectrum of PCOS manifestations, leading to underdiagnosis or misclassification. In response to these challenges, researchers have increasingly explored machine learning techniques as a means to improve diagnostic accuracy and consistency. Early studies in this domain primarily employed conventional machine learning models such as support vector machines (SVM), k-nearest neighbours, and logistic regression to classify patients based on structured clinical and biochemical data. These models demonstrated promising results, with several studies reporting diagnostic accuracies exceeding 80%, indicating that algorithmic approaches can effectively identify patterns associated with PCOS that may not be readily apparent through traditional clinical assessment (1). As the field progressed, more advanced ensemble methods, including random forests and gradient boosting algorithms, were introduced, offering improved performance through their ability to model complex, non-linear relationships within the data.

More recently, deep learning approaches have further expanded the capabilities of AI in PCOS diagnosis. Convolutional neural networks (CNNs), particularly in imaging-based studies, have achieved exceptionally high accuracy rates, in some cases approaching 98–99%, by analysing ovarian ultrasound images and identifying subtle morphological features associated with the disorder (2). These advancements highlight the potential of AI to enhance diagnostic precision while reducing reliance on operator-dependent techniques. However, it is important to note that many of these high-performing models are developed using relatively small and homogeneous datasets, which

limits their generalisability to broader populations. Despite the encouraging results, several limitations persist within the existing body of research. A significant proportion of studies rely on retrospective data and lack external validation, raising concerns about overfitting and reproducibility. Furthermore, inconsistencies in the application of diagnostic criteria across studies introduce variability in model training and evaluation, which may affect the reliability of reported outcomes (1). Another critical limitation is the limited integration of multi-modal data. While some studies incorporate imaging and biochemical markers, few have successfully combined these with longitudinal clinical data, thereby restricting the ability of models to capture the dynamic nature of PCOS progression.

In addition, the issue of interpretability remains a major barrier to clinical adoption. Many AI models function as black-box systems, providing predictions without clear explanations of the underlying decision-making process. Although recent research has begun to incorporate explainable AI techniques such as SHAP and LIME, their use is still not widespread, with only a minority of studies addressing the need for transparency in model outputs (2). This gap underscores the need for future research to prioritise not only accuracy but also interpretability and clinical relevance

## III. AI IN THYROID DISORDER DETECTION

The application of artificial intelligence in thyroid disorder detection has followed a trajectory similar to that observed in PCOS research, with initial studies focusing on traditional machine learning techniques and more recent work exploring deep learning and hybrid models. Thyroid disorders, including hypothyroidism and hyperthyroidism, present a unique opportunity for AI-based detection due to the availability of well-defined biochemical markers, such as thyroid-stimulating hormone (TSH) and thyroid hormones (T3 and T4), which can be quantitatively analysed. A substantial body of literature has demonstrated the effectiveness of machine learning algorithms in classifying thyroid disorders based on clinical and laboratory data. Supervised learning models, including decision trees, random forests, and neural networks, have been widely used, often achieving high levels of diagnostic accuracy. A

systematic review analysing multiple studies reported an average accuracy of approximately 96.8% in detecting thyroid diseases, highlighting the strong potential of AI in this domain (3). These findings suggest that AI models can serve as valuable tools for screening and early detection, particularly in settings where access to specialised endocrinological expertise may be limited.

In addition to structured data, AI has also been applied to imaging-based diagnosis of thyroid conditions, particularly in the detection of thyroid nodules and cancer using ultrasound imaging. Deep learning models have demonstrated the ability to automatically extract relevant features from imaging data, thereby reducing the dependence on manual interpretation and improving diagnostic consistency. These approaches have been particularly useful in differentiating benign from malignant nodules, a task that traditionally requires significant clinical expertise.

However, despite these advancements, important gaps remain in the application of AI to thyroid disorder detection. One of the primary limitations is the reliance on static datasets that do not capture temporal variations in hormone levels. Thyroid function is inherently dynamic, and single-point measurements may not accurately reflect underlying physiological changes. As a result, models trained on cross-sectional data may fail to identify early or subclinical stages of disease. Another limitation is the widespread use of private or institution-specific datasets, which limits the generalisability of findings. Many studies lack external validation and are conducted on relatively small sample sizes, raising concerns about the robustness of the models. Furthermore, while accuracy metrics are often reported, there is less emphasis on clinically relevant outcomes such as early detection or progression prediction.

The issue of explainability is also pertinent in this context. Although AI models can achieve high accuracy, their clinical utility depends on the ability to provide interpretable results that can inform decision-making. Without transparency, clinicians may be hesitant to rely on AI-generated predictions, particularly in cases where treatment decisions carry significant risks

#### IV. USE OF EHR DATA IN PREDICTIVE MODELING

The integration of Electronic Health Record data into predictive modelling represents a significant advancement in the application of artificial intelligence in healthcare. Unlike traditional datasets, which are often limited in scope and duration, EHRs provide comprehensive, longitudinal records of patient health, capturing a wide range of variables over time. This includes demographic information, laboratory results, medication histories, diagnostic codes, and clinical notes, all of which contribute to a rich and multidimensional dataset. One of the key advantages of EHR data is its ability to reflect the temporal progression of disease. Endocrine disorders such as PCOS and thyroid dysfunction do not develop abruptly; rather, they evolve gradually through a series of physiological changes. EHR data enable the tracking of these changes over time, providing valuable insights into disease trajectories. Predictive models that leverage this temporal information can identify early warning signs and generate risk predictions before the onset of overt symptoms.

Machine learning models applied to EHR data have demonstrated considerable success in predicting a wide range of clinical outcomes. For example, studies using recurrent neural networks have shown the ability to predict future diagnoses based on historical patient records, highlighting the potential of temporal modelling in healthcare (4). In the context of PCOS, recent research has utilised EHR data to develop predictive models capable of identifying at-risk individuals within large outpatient populations, thereby facilitating earlier diagnosis and intervention (5). Despite these advantages, the use of EHR data presents several challenges. Data quality is a major concern, as EHRs often contain missing, inconsistent, or erroneous entries. Variations in data recording practices across institutions further complicate the integration of datasets. Additionally, the presence of unstructured data, such as clinical notes, requires advanced natural language processing techniques to extract meaningful information. Another important consideration is data privacy and security. The use of patient data for predictive modelling raises ethical concerns, particularly with respect to consent and confidentiality. Regulatory frameworks governing the use of health data vary across regions, and compliance

with these regulations is essential for the responsible implementation of AI systems. Nevertheless, the potential benefits of leveraging EHR data for predictive modelling are substantial. By enabling early detection and personalised risk assessment, these models have the potential to transform clinical practice and improve patient outcomes.

## V. GAPS IN EXISTING RESEARCH

Despite the rapid growth of research in AI-based diagnosis of endocrine disorders, several critical gaps remain that limit the translation of these technologies into clinical practice. One of the most notable gaps is the lack of integrated models that simultaneously address both PCOS and thyroid disorders. Given the known association between these conditions, particularly in the context of hormonal and metabolic dysregulation, the development of combined predictive models represents an important yet underexplored area of research. Current studies tend to focus on individual conditions in isolation, thereby missing the opportunity to identify shared patterns and comorbid risk factors. Another significant gap is the limited incorporation of temporal modelling in existing studies. While some research has utilised longitudinal data, the majority of models are based on cross-sectional datasets that do not fully capture the dynamic nature of endocrine disorders. Temporal models, such as recurrent neural networks and LSTM architectures, offer the ability to analyse sequential data and identify trends over time. However, their application in this domain remains relatively limited, highlighting an important direction for future research. The issue of explainability also represents a major challenge. Although AI models have demonstrated high levels of accuracy, their lack of transparency limits their clinical applicability. Only a small proportion of studies incorporate explainable AI techniques, and even fewer evaluate the clinical relevance of the explanations provided. This gap is particularly important in healthcare, where trust and accountability are essential for adoption.

In addition to methodological limitations, there are also gaps related to data diversity and representativeness. Many studies are conducted using datasets from specific geographic regions or healthcare settings, which may not be representative of broader populations. This raises concerns about the

generalisability of models and their ability to perform effectively across diverse patient groups.

Finally, there is a lack of real-world validation and implementation studies. While many models demonstrate strong performance in controlled research settings, few have been tested in clinical environments. Bridging this gap requires collaboration between researchers, clinicians, and healthcare institutions to evaluate the practical feasibility and impact of AI-based predictive systems.

## VI. STUDY DESIGN

The present study is structured around a retrospective analytical design combined with a conceptual modelling framework aimed at developing an AI-driven predictive system for the early detection of Polycystic Ovary Syndrome (PCOS) and thyroid disorders. A retrospective approach is particularly appropriate in this context, as it allows for the utilisation of existing Electronic Health Record (EHR) data to identify patterns that precede clinical diagnosis. Unlike prospective studies, which require long follow-up periods and extensive resource allocation, retrospective analyses enable the efficient examination of large patient cohorts over extended timeframes, thereby providing valuable insights into disease progression and early indicators.

At the core of the study lies a conceptual framework that integrates clinical endocrinology with advanced computational modelling. This framework is designed to capture the multidimensional and temporal nature of endocrine disorders by mapping patient-level clinical data onto predictive outputs. The conceptual model assumes that both PCOS and thyroid disorders evolve through a sequence of measurable physiological changes, many of which are recorded in routine healthcare interactions. By systematically analysing these changes, the study seeks to establish a predictive relationship between early clinical markers and eventual diagnosis. Furthermore, the study adopts a comparative modelling approach, wherein multiple machine learning and deep learning algorithms are evaluated to determine their relative effectiveness in early prediction. This design not only enhances the robustness of the findings but also provides a comprehensive understanding of how different modelling techniques perform in the context of complex, real-world healthcare data (6,7).

## VII. DATA SOURCE (ELECTRONIC HEALTH RECORDS)

The primary data source for this study is Electronic Health Records, which provide a comprehensive and longitudinal repository of patient information. EHR systems have become an integral component of modern healthcare infrastructure, capturing a wide range of clinical data across different stages of patient care. These records include both structured and semi-structured data, encompassing demographic details, laboratory results, diagnostic codes, medication histories, and clinical observations. The use of EHR data offers several advantages for predictive modelling. Most notably, it enables the analysis of longitudinal patient trajectories, allowing for the identification of trends and patterns that develop over time. For endocrine disorders such as PCOS and thyroid dysfunction, which often progress gradually, this temporal dimension is particularly valuable. Subtle changes in hormonal levels or metabolic parameters may not be clinically significant when considered in isolation but can provide important predictive signals when analysed over extended periods. In addition to structured data, EHRs also contain unstructured information in the form of clinical notes and narrative descriptions. Although these data require advanced processing techniques such as natural language processing, they can offer additional context that enhances the predictive capacity of the model. However, the heterogeneity and variability of EHR data across institutions necessitate careful standardisation and preprocessing to ensure consistency and reliability (8).

## VIII. STUDY POPULATION

The study population is derived from individuals whose health records are available within the selected EHR databases. The inclusion criteria are defined to ensure that sufficient and relevant data are available for predictive modelling. Specifically, individuals included in the study must have documented endocrine-related evaluations, including at least one set of hormonal or metabolic laboratory results, as well as follow-up records that allow for the assessment of disease progression. The population primarily focuses on individuals within the reproductive and adult age groups, as these are the populations most affected by

PCOS and thyroid disorders. Both male and female patients may be included for thyroid disorder analysis, while PCOS-specific modelling is restricted to female patients. The inclusion of diverse demographic groups is essential to ensure that the predictive models are generalisable and capable of capturing variations across different populations.

Exclusion criteria are applied to minimise confounding factors that may distort the predictive relationships. These include patients with incomplete or inconsistent records, individuals with pre-existing endocrine disorders diagnosed prior to the study period, and cases involving conditions or treatments that significantly alter hormonal profiles, such as pregnancy or endocrine surgeries. By carefully defining the study population, the research aims to create a dataset that is both representative and suitable for robust predictive modelling (6).

## IX. VARIABLES AND FEATURE SPACE

The predictive framework is built upon a comprehensive set of variables that reflect the multifactorial nature of endocrine disorders. These variables are broadly categorised into hormonal, metabolic, and clinical features, each contributing unique insights into patient health. Hormonal variables form the core of the feature space, as both PCOS and thyroid disorders are fundamentally characterised by hormonal imbalances. Key indicators include thyroid-stimulating hormone (TSH), free triiodothyronine (T3), free thyroxine (T4), luteinising hormone (LH), follicle-stimulating hormone (FSH), and androgen levels such as testosterone. These biomarkers provide direct measures of endocrine function and are essential for both diagnosis and prediction.

Metabolic variables complement hormonal data by capturing the broader physiological impact of endocrine dysfunction. Parameters such as fasting glucose, insulin levels, lipid profiles, and body mass index (BMI) are included to reflect metabolic status and its interaction with hormonal regulation. In the case of PCOS, insulin resistance is a key feature, while thyroid disorders are closely linked to metabolic rate and lipid metabolism. Clinical variables, including patient-reported symptoms, menstrual history, and comorbid conditions, provide additional context that enhances model performance. These variables help bridge the gap between biochemical measurements

and real-world clinical presentation, allowing the model to capture a more holistic view of patient health (5).

#### X. DATA PREPROCESSING

The preprocessing of EHR data represents a critical step in the modelling pipeline, as the quality and consistency of input data directly influence model performance. EHR datasets are inherently complex and often contain missing values, inconsistencies, and noise, which must be addressed through systematic data cleaning procedures. Missing data are handled using advanced imputation techniques that preserve the statistical properties of the dataset while minimising bias. Methods such as multiple imputation and model-based imputation are commonly employed to estimate missing values based on observed data patterns. These approaches are preferred over simple techniques such as mean substitution, as they provide more accurate and reliable estimates (9).

Normalization is applied to continuous variables to ensure comparability across different measurement scales. This is particularly important when combining data from multiple sources, where variations in units and reference ranges may exist. Standardisation techniques, such as z-score normalization, are used to transform variables into a common scale, thereby facilitating efficient model training. In addition, categorical variables are encoded into numerical representations suitable for machine learning algorithms. The preprocessing stage also involves the identification and removal of outliers, which may arise from data entry errors or rare clinical conditions. By ensuring data quality and consistency, preprocessing lays the foundation for accurate and reliable predictive modelling.

#### XI. FEATURE ENGINEERING

Feature engineering is a crucial component of the modelling process, as it enhances the predictive capacity of the model by transforming raw data into meaningful representations. In the context of endocrine disorders, feature engineering involves the creation of derived variables that capture clinically relevant relationships between biomarkers. One of the most commonly used derived features in PCOS research is the ratio of luteinising hormone to follicle-

stimulating hormone (LH/FSH), which serves as an indicator of hormonal imbalance. Similarly, indices of insulin resistance, such as the Homeostatic Model Assessment (HOMA-IR), are calculated to reflect metabolic dysfunction. These derived variables provide additional layers of information that are not directly available in raw data.

Temporal feature engineering is particularly important for capturing trends in longitudinal data. By analysing changes in hormone levels over time, the model can identify patterns indicative of disease progression. For example, a gradual increase in TSH levels may signal the onset of hypothyroidism, while fluctuations in androgen levels may indicate emerging PCOS. The incorporation of interaction terms and composite features further enhances the model's ability to capture complex relationships within the data. By integrating domain knowledge with computational techniques, feature engineering bridges the gap between clinical understanding and machine learning (7).

#### XII. MODEL DEVELOPMENT

The development of predictive models in this study involves the implementation and comparison of multiple machine learning and deep learning algorithms. Each model is selected based on its ability to capture different aspects of the data and provide complementary insights into disease prediction.

Logistic regression serves as a baseline model due to its simplicity and interpretability. It provides a transparent framework for understanding the relationship between input variables and disease outcomes, making it a valuable reference point for evaluating more complex models. However, its ability to capture non-linear relationships is limited, which may restrict its performance in complex datasets.

Ensemble methods such as Random Forest and XGBoost are employed to address these limitations. Random Forest models leverage multiple decision trees to improve predictive accuracy and reduce overfitting, while XGBoost utilises gradient boosting techniques to optimise model performance. These methods are particularly effective in handling high-dimensional data and capturing non-linear interactions between variables (10). To account for the temporal nature of EHR data, deep learning models are incorporated, particularly Long Short-Term Memory (LSTM) networks. LSTM models are designed to

process sequential data and retain information over time, making them well-suited for analysing longitudinal patient records. Hybrid architectures that combine convolutional neural networks (CNNs) with LSTM layers are also explored, enabling the simultaneous extraction of spatial and temporal features from the data (11,12).

### XIII. VALIDATION STRATEGY

The validation of predictive models is essential to ensure their reliability and generalisability. In this study, a multi-level validation strategy is employed to evaluate model performance under different conditions. The dataset is divided into training, validation, and testing subsets, with careful attention to maintaining representative distributions of patient characteristics. Cross-validation techniques are used to optimise model parameters and assess performance across multiple data partitions. This approach reduces the risk of overfitting and ensures that the model performs consistently across different subsets of data. In addition to cross-validation, temporal validation is incorporated to simulate real-world deployment scenarios. Temporal validation involves training the model on historical data and testing it on more recent data, thereby assessing its ability to predict future outcomes. This approach is particularly relevant for healthcare applications, where models must operate on continuously evolving data streams. By incorporating both cross-sectional and temporal validation methods, the study ensures a comprehensive evaluation of model performance (13).

### XIV. PERFORMANCE METRICS

The evaluation of predictive models is conducted using a range of performance metrics that reflect both statistical accuracy and clinical relevance. The area under the receiver operating characteristic curve (AUC-ROC) is used as a primary metric, as it provides a measure of the model's ability to distinguish between positive and negative cases across different thresholds. In addition to AUC, metrics such as accuracy, sensitivity, specificity, precision, and F1-score are calculated to provide a comprehensive assessment of model performance. Sensitivity, or true positive rate, is particularly important in the context of early disease detection, as it reflects the model's ability to

correctly identify individuals at risk. Specificity, on the other hand, measures the ability to correctly identify non-diseased individuals, thereby reducing false positives. Precision and F1-score provide additional insights into the balance between sensitivity and specificity, while calibration metrics assess the alignment between predicted probabilities and observed outcomes. By using a combination of these metrics, the study ensures that model evaluation is both rigorous and clinically meaningful (10).

### XV. SYSTEM ARCHITECTURE

The proposed AI framework for the early diagnosis of Polycystic Ovary Syndrome (PCOS) and thyroid disorders is designed as a multi-layered architecture that systematically transforms raw Electronic Health Record (EHR) data into clinically meaningful predictive outputs. At a conceptual level, the system is structured into three primary components: input, processing, and output. However, within each of these components lies a series of interconnected processes that collectively enable accurate and interpretable disease prediction. The input layer represents the point of entry for all patient-related data derived from EHR systems. This includes structured variables such as demographic information, laboratory results, and diagnostic codes, as well as semi-structured data like clinical observations. The richness of this input layer is fundamental to the effectiveness of the model, as endocrine disorders are inherently multifactorial and require the integration of diverse data types for accurate representation. By incorporating a wide range of variables, the system is able to capture both direct biomarkers and indirect indicators of disease progression.

The processing layer forms the core of the framework, where data undergo transformation, analysis, and modelling. This layer is not a single step but rather a sequence of operations that include preprocessing, feature engineering, and model execution. Within this stage, machine learning algorithms analyse complex relationships between variables, identifying patterns that may indicate early disease development. Advanced models such as ensemble learning techniques and deep neural networks are particularly effective in this context, as they can capture non-linear interactions and high-dimensional dependencies within the data (10,12).

The output layer translates the results of computational analysis into clinically actionable insights. Rather than providing binary classifications alone, the system generates probabilistic risk scores that indicate the likelihood of a patient developing PCOS or thyroid disorders. These outputs can be integrated into clinical decision support systems, enabling healthcare professionals to make informed decisions based on both quantitative predictions and underlying data trends. Importantly, the inclusion of explainability mechanisms ensures that these predictions are interpretable and aligned with clinical reasoning, thereby enhancing trust and usability (7).

The architecture is designed to be modular and scalable, allowing for the integration of additional data sources and modelling techniques as the system evolves. This flexibility is essential in the context of rapidly advancing AI technologies and continuously expanding healthcare datasets. By adopting a layered and adaptable design, the proposed framework provides a robust foundation for predictive modelling in endocrine disorders.

#### XVI. DATA FLOW PIPELINE

The functionality of the proposed AI system is further clarified through the data flow pipeline, which illustrates the sequential transformation of raw EHR data into predictive outputs. This pipeline begins with the extraction of patient data from EHR systems, followed by a series of preprocessing and analytical steps that culminate in disease prediction. At the initial stage, raw EHR data are collected and aggregated from various clinical sources. These data are often heterogeneous, comprising multiple formats and levels of completeness. As such, the preprocessing stage plays a critical role in preparing the data for analysis. This involves cleaning the dataset, handling missing values, normalising variables, and ensuring consistency across different records. Without this step, the reliability of the predictive model would be significantly compromised. Following preprocessing, the data undergo feature engineering, where raw variables are transformed into meaningful representations that enhance model performance. This includes the creation of derived features such as hormonal ratios and temporal trends, which capture clinically relevant relationships between variables. Feature engineering serves as a bridge between raw

data and model input, ensuring that the information fed into the algorithm reflects both statistical patterns and clinical significance.

The processed data are then passed into the modelling stage, where machine learning algorithms are applied to generate predictions. Depending on the complexity of the data and the specific objectives of the study, different models may be used either independently or in combination. Ensemble methods such as Random Forest and XGBoost provide strong baseline performance, while deep learning models such as CNN-LSTM architectures offer enhanced capabilities for handling sequential data (10,11). Once the model generates predictions, the results are translated into clinically interpretable outputs. This includes risk scores, probability estimates, and classification labels that indicate the likelihood of disease presence or development. These outputs can be visualised through dashboards or integrated into clinical workflows, enabling real-time decision support.

#### XVII. ROLE OF TEMPORAL DATA IN PREDICTIVE MODELING

One of the most critical aspects of the proposed AI framework is the incorporation of temporal data, which fundamentally distinguishes it from traditional diagnostic approaches. Endocrine disorders such as PCOS and thyroid dysfunction are not static conditions; they develop gradually through a series of physiological changes that unfold over time. As such, the ability to analyse longitudinal data is essential for accurate early prediction. Traditional models that rely on cross-sectional data provide only a snapshot of patient health at a single point in time. While such models can be useful for classification, they are inherently limited in their ability to detect early-stage disease or predict future outcomes. In contrast, temporal models analyse sequences of data points, enabling the identification of trends and patterns that may indicate emerging pathology.

For example, a gradual increase in thyroid-stimulating hormone levels over several months may signal the onset of hypothyroidism, even if individual values fall within the normal range. Similarly, fluctuations in hormonal ratios and metabolic markers over time may provide early indicators of PCOS. By capturing these temporal dynamics, AI models can generate predictions that are both more accurate and more

clinically meaningful. Recurrent neural networks, particularly Long Short-Term Memory (LSTM) models, are specifically designed to handle sequential data and retain information over time. These models can learn dependencies between past and present data points, allowing them to identify patterns that span multiple time steps. When combined with convolutional layers in a CNN-LSTM architecture, the model can simultaneously capture spatial and temporal features, further enhancing predictive performance (11,12). The importance of temporal modelling is also reflected in the validation strategy, where models are evaluated based on their ability to predict future outcomes rather than simply classify existing cases. This approach aligns with real-world clinical scenarios, where the goal is not merely to diagnose disease but to anticipate its development and intervene at an early stage.

However, the use of temporal data also introduces additional challenges, including increased computational complexity and the need for well-structured longitudinal datasets. Missing data and irregular time intervals can complicate analysis, requiring sophisticated techniques for data imputation and sequence alignment. Despite these challenges, the benefits of temporal modelling far outweigh its limitations, making it a central component of the proposed AI framework.

In summary, the proposed AI framework integrates a structured system architecture, a dynamic data flow pipeline, and advanced temporal modelling techniques to enable early and accurate prediction of PCOS and thyroid disorders. By leveraging the full potential of EHR data and combining it with state-of-the-art machine learning approaches, the framework represents a significant step toward proactive and data-driven healthcare

### XVIII. EXPLAINABLE AI IN CLINICAL DECISION MAKING

The integration of artificial intelligence into healthcare has brought unprecedented advancements in diagnostic accuracy and predictive modelling. However, alongside these benefits arises a fundamental challenge that directly impacts clinical adoption: the need for interpretability. In medical practice, decisions are rarely made on the basis of outcomes alone; rather, they are grounded in a clear

understanding of underlying reasoning, supported by evidence, clinical judgment, and patient context. The introduction of complex machine learning models, particularly deep learning systems, has disrupted this paradigm by producing highly accurate predictions that are often difficult to interpret.

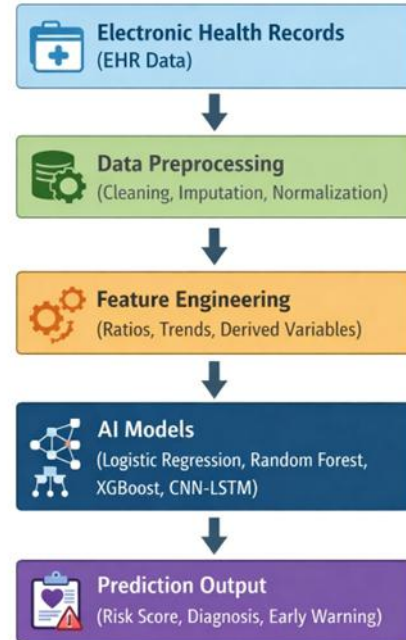


Fig 1 AI work flow

In the context of endocrine disorders such as Polycystic Ovary Syndrome (PCOS) and thyroid dysfunction, interpretability becomes even more critical. These conditions are characterised by multifactorial interactions among hormonal, metabolic, and clinical variables. Clinicians must not only determine whether a patient is at risk but also understand which factors contribute to that risk. A predictive model that identifies a patient as high-risk without providing an explanation offers limited clinical utility, as it does not support informed decision-making or guide subsequent interventions (17). The concept of interpretability in AI encompasses both transparency and explainability. Transparency refers to the inherent simplicity of a model, where its internal workings are directly understandable, as seen in linear regression or decision trees. Explainability, on the other hand, involves the use of additional techniques to interpret complex models, often referred to as “black-box” systems. While simpler models are inherently interpretable, they may lack the predictive power required for

complex clinical datasets. Conversely, advanced models such as neural networks offer superior performance but require external methods to make their predictions understandable (18). The lack of interpretability has significant implications for trust and accountability in healthcare. Clinicians are trained to critically evaluate evidence and justify their decisions, particularly in cases where treatment carries potential risks. Without clear explanations, AI-generated predictions may be perceived as unreliable or even unsafe. This concern is further amplified by the ethical and legal responsibilities associated with patient care, where decisions must be transparent and defensible (19).

Moreover, interpretability plays a crucial role in identifying biases and errors within AI models. Healthcare data are often subject to systemic biases arising from demographic, socioeconomic, and institutional factors. Without interpretability, these biases may remain hidden, leading to unequal or suboptimal outcomes for certain patient groups. By providing insights into model behaviour, explainable AI enables the detection and mitigation of such biases, thereby promoting fairness and equity in healthcare delivery (20). Another important aspect of interpretability is its role in knowledge discovery. Beyond prediction, AI models have the potential to uncover novel associations and patterns within clinical data. For instance, an explainable model may reveal previously unrecognised relationships between hormonal markers and disease risk, contributing to a deeper understanding of endocrine disorders. This dual function of prediction and discovery underscores the broader value of explainable AI in advancing medical science (21)

#### XIX. METHODS FOR EXPLAINABILITY: SHAP AND LIME

To address the challenges associated with black-box models, a range of explainable AI techniques has been developed, among which SHapley Additive exPlanations (SHAP) and Local Interpretable Model-agnostic Explanations (LIME) are the most widely used in healthcare applications. These methods provide a framework for interpreting complex model predictions by quantifying the contribution of individual features to the overall output. SHAP is based on cooperative game theory and assigns each

feature an importance value representing its contribution to the prediction. The method considers all possible combinations of features and calculates the marginal contribution of each variable, ensuring a fair and consistent allocation of importance. One of the key advantages of SHAP is its ability to provide both global and local explanations. Global explanations offer insights into overall model behaviour, identifying the most influential features across the dataset, while local explanations focus on individual predictions, explaining why a specific patient is classified as high or low risk (22).

In the context of PCOS and thyroid disorder prediction, SHAP can be used to identify which hormonal or metabolic variables are driving model predictions. For example, elevated TSH levels or abnormal LH/FSH ratios may be highlighted as key contributors to risk scores. By presenting these contributions in a visual and interpretable format, SHAP enables clinicians to understand and validate the model's reasoning, thereby facilitating its integration into clinical workflows. LIME, on the other hand, adopts a different approach by approximating complex models with simpler, interpretable models in the vicinity of a specific prediction. Rather than attempting to explain the entire model, LIME focuses on local interpretability, providing explanations for individual predictions. This is achieved by generating perturbed samples around the input data and fitting a simple model, such as a linear regression, to approximate the behaviour of the complex model in that region (23).

The strength of LIME lies in its flexibility and model-agnostic nature, meaning it can be applied to any machine learning model regardless of its internal structure. This makes it particularly useful in healthcare settings, where a variety of models may be used depending on the application. However, LIME also has limitations, including potential instability in explanations and sensitivity to parameter selection, which may affect its reliability in certain contexts. In addition to SHAP and LIME, other explainability techniques have been explored in healthcare, including attention mechanisms in neural networks, feature importance scores, and partial dependence plots. Each of these methods offers unique advantages and limitations, and their selection depends on the specific requirements of the application. For instance, attention-based models inherently provide a form of

interpretability by highlighting the parts of the input data that are most relevant to the prediction (24). The integration of explainability techniques into AI frameworks requires careful consideration of both technical and clinical factors. While these methods enhance interpretability, they also introduce additional computational complexity and may require specialised expertise to implement and interpret. Despite these challenges, their inclusion is essential for ensuring that AI models are not only accurate but also transparent and clinically meaningful.

## XX. CLINICAL RELEVANCE AND IMPACT ON DECISION MAKING

The clinical relevance of explainable AI extends beyond technical considerations, directly influencing how predictive models are utilised in real-world healthcare settings. For AI systems to be effectively integrated into clinical practice, they must align with existing workflows, support decision-making processes, and ultimately improve patient outcomes. Explainability plays a central role in achieving these objectives by bridging the gap between computational predictions and clinical reasoning. In the diagnosis and management of PCOS and thyroid disorders, explainable AI can enhance clinical decision-making in several ways. First, it provides clinicians with a clearer understanding of patient-specific risk factors, enabling more personalised and targeted interventions. For example, if a predictive model identifies insulin resistance as a **प्रमुख** contributor to PCOS risk, clinicians can prioritise metabolic interventions alongside hormonal treatment. This level of insight supports a more holistic approach to patient care, addressing both symptoms and underlying causes.

Second, explainable AI facilitates shared decision-making between clinicians and patients. In modern healthcare, patient engagement and informed consent are essential components of care. By providing transparent explanations of risk predictions, AI systems enable clinicians to communicate findings more effectively, helping patients understand their condition and participate in treatment decisions. This is particularly important in chronic conditions such as PCOS, where long-term management requires sustained patient involvement (25). The integration of explainable AI also has implications for clinical training and education. As AI becomes increasingly

prevalent in healthcare, clinicians must develop the skills **اللازمة** to interpret and utilise these tools effectively. Explainable models can serve as educational resources, illustrating how different variables interact to influence disease risk and reinforcing clinical knowledge. This synergy between AI and medical education represents an important avenue for future development (26). From a systems perspective, explainable AI can contribute to improved healthcare efficiency and quality. By providing early and accurate predictions, these models can reduce diagnostic delays, minimise unnecessary testing, and optimise resource allocation. At the same time, the transparency of explainable models ensures that these benefits are achieved without compromising patient safety or ethical standards.

However, the implementation of explainable AI in clinical practice is not without challenges. One of the **प्रमुख** barriers is the potential mismatch between technical explanations and clinical understanding. While methods such as SHAP provide quantitative measures of feature importance, these may not always align with clinical reasoning or be easily interpretable by healthcare professionals. Bridging this gap requires the development of user-friendly interfaces and visualisation tools that translate complex outputs into clinically meaningful insights (27).

Another challenge is the need for standardisation and validation of explainability methods. Unlike traditional performance metrics, there is currently no universally accepted **معیار** for evaluating the quality of explanations. This lack of standardisation can hinder the comparison of different models and limit the adoption of explainable AI in clinical settings. Ongoing research is therefore focused on developing frameworks for assessing explainability and ensuring that it meets the needs of both clinicians and patients (28).

Despite these challenges, the potential of explainable AI to transform clinical decision-making is substantial. By combining predictive accuracy with interpretability, these systems offer a powerful tool for improving the diagnosis and management of endocrine disorders. As research continues to advance, the integration of explainable AI into healthcare is likely to play a central role in the transition toward precision and preventive medicine.

## XXI. EARLY DIAGNOSIS BENEFITS

The integration of AI-powered predictive modelling into the clinical landscape has the potential to fundamentally reshape the way endocrine disorders such as Polycystic Ovary Syndrome (PCOS) and thyroid dysfunction are identified and managed. One of the most significant clinical implications of this approach lies in its ability to enable earlier diagnosis, shifting the paradigm from reactive treatment to proactive care. Traditionally, the diagnosis of these conditions is often delayed due to the gradual onset of symptoms and their non-specific nature. Patients may experience subtle physiological changes for months or even years before receiving a definitive diagnosis, during which time the disease continues to progress and complications may develop. Early diagnosis, facilitated by predictive modelling using Electronic Health Record (EHR) data, offers a critical window of opportunity for intervention. By identifying individuals at risk before the manifestation of overt clinical symptoms, clinicians can initiate targeted monitoring and preventive strategies. For instance, in the case of PCOS, early identification of hormonal imbalance and insulin resistance can lead to timely lifestyle interventions, pharmacological management, and counselling aimed at preventing long-term complications such as infertility, metabolic syndrome, and cardiovascular disease (29). Similarly, early detection of thyroid dysfunction allows for prompt correction of hormonal imbalances, thereby preventing the progression to more severe clinical states and reducing the risk of systemic complications. Beyond individual patient outcomes, early diagnosis also contributes to improved diagnostic efficiency within healthcare systems. Predictive models can serve as screening tools that prioritise high-risk individuals for further evaluation, thereby optimising the use of clinical resources. This is particularly relevant in resource-constrained settings, where access to specialised diagnostic services may be limited. By guiding clinicians toward patients who are most likely to benefit from further testing, AI-driven systems can reduce unnecessary investigations while ensuring timely care for those in need (30).

Another important aspect of early diagnosis is its impact on patient experience. Delayed diagnosis is often associated with frustration, uncertainty, and reduced quality of life, particularly in conditions such

as PCOS where symptoms may affect physical appearance, reproductive health, and psychological well-being. Early identification and intervention can alleviate these burdens by providing clarity and enabling patients to take control of their health at an earlier stage. This not only improves clinical outcomes but also enhances patient satisfaction and engagement in the care process (31).

## XXII. REDUCTION IN DISEASE BURDEN

The broader impact of AI-driven early diagnosis extends to the reduction of disease burden at both individual and population levels. Endocrine disorders such as PCOS and thyroid dysfunction are associated with a wide range of complications that contribute to long-term morbidity and healthcare costs. These include metabolic disorders, cardiovascular disease, reproductive complications, and mental health conditions, all of which place a significant strain on healthcare systems. By enabling earlier detection and intervention, predictive modelling has the potential to reduce the incidence and severity of these complications. For example, early management of insulin resistance in PCOS can prevent the progression to type 2 diabetes, a condition that carries substantial health and economic consequences. Similarly, timely treatment of thyroid disorders can prevent complications such as dyslipidaemia, hypertension, and cognitive impairment, thereby improving overall health outcomes (32).

At a population level, the implementation of AI-based predictive systems can contribute to more effective disease surveillance and public health planning. By analysing large-scale EHR data, healthcare systems can identify trends and patterns in disease prevalence, enabling targeted interventions and resource allocation. This is particularly important in regions where endocrine disorders are highly prevalent and healthcare resources are limited. Predictive modelling can support the development of screening programs and preventive strategies that address the specific needs of different populations (33). The economic implications of reducing disease burden are also significant. Chronic conditions such as PCOS and thyroid disorders require ongoing management, including regular monitoring, medication, and treatment of complications. By preventing or delaying the onset of these conditions, early diagnosis can

reduce long-term healthcare costs and improve the sustainability of healthcare systems. Studies have shown that preventive interventions are often more cost-effective than treating advanced disease, highlighting the value of predictive modelling as a tool for healthcare optimisation (34). In addition to physical health outcomes, the reduction of disease burden also encompasses improvements in mental health and quality of life. Endocrine disorders are often associated with psychological challenges, including anxiety, depression, and reduced self-esteem. By addressing these conditions at an earlier stage, healthcare providers can mitigate their psychological impact and support holistic patient well-being. This integrated approach to care reflects a shift toward more patient-centred healthcare, where the focus extends beyond disease management to overall quality of life (35).

### XXIII. PERSONALIZED MEDICINE

One of the most transformative implications of AI in healthcare is its contribution to the advancement of personalised medicine. Traditional medical approaches often rely on generalised treatment protocols that may not account for individual variability in disease presentation and response to therapy. In contrast, AI-driven predictive modelling enables the development of tailored interventions based on individual patient data, thereby enhancing the precision and effectiveness of care. In the context of PCOS and thyroid disorders, personalised medicine is particularly relevant due to the heterogeneity of these conditions. Patients with PCOS, for example, may present with varying combinations of hormonal imbalance, metabolic dysfunction, and reproductive issues, each requiring a different therapeutic approach. Similarly, thyroid disorders can manifest in diverse forms, ranging from subclinical abnormalities to overt dysfunction, with varying responses to treatment. By analysing comprehensive EHR data, AI models can identify patient-specific patterns and risk factors, enabling the development of customised treatment plans. For instance, a patient identified as having a high risk of PCOS driven primarily by insulin resistance may benefit from targeted metabolic interventions, while another patient with predominant hormonal imbalance may require a different therapeutic strategy. This level of precision not only

improves treatment outcomes but also reduces the risk of adverse effects associated with inappropriate or unnecessary interventions (36). Personalised medicine also extends to the monitoring and management of disease over time. AI systems can continuously analyse patient data to assess treatment response and adjust interventions accordingly. This dynamic approach to care ensures that treatment remains aligned with the patient's evolving health status, thereby improving long-term outcomes. In the case of thyroid disorders, for example, AI models can assist in optimising hormone replacement therapy by predicting changes in thyroid function and adjusting dosage accordingly (37).

Furthermore, personalised medicine enhances patient engagement by involving individuals in their own care. By providing tailored insights and recommendations, AI systems empower patients to make informed decisions about their health. This collaborative approach fosters a sense of ownership and responsibility, which is essential for the successful management of chronic conditions.

However, the implementation of personalised medicine also presents challenges, including the need for robust data integration, advanced analytical capabilities, and careful consideration of ethical and privacy concerns. Ensuring that personalised interventions are equitable and accessible to all patients is a critical consideration, particularly in diverse and resource-limited settings (38).

### XXIV. CHALLENGES AND LIMITATIONS

The successful implementation of AI-powered predictive modelling in healthcare is fundamentally dependent on the quality of the underlying data. Electronic Health Records (EHRs), while rich in information, are not originally designed for research or machine learning applications. Instead, they are developed primarily for clinical documentation and administrative purposes, which often results in datasets that are heterogeneous, incomplete, and prone to inconsistencies. These limitations pose significant challenges for the development of reliable and generalisable predictive models. One of the most pervasive issues in EHR data is missing information. Clinical data are often recorded selectively, based on immediate clinical needs rather than systematic data collection protocols. As a result, important variables

such as hormonal levels, metabolic indicators, or patient history may be absent or inconsistently recorded across patients. Missing data can introduce bias into predictive models, particularly if the absence of data is not random but related to specific patient characteristics or clinical practices. Although advanced imputation techniques can mitigate this issue to some extent, they cannot fully replace the value of complete and accurate data (39).

In addition to missing data, inconsistencies in data entry present another major challenge. Variations in measurement units, laboratory reference ranges, and coding practices across different healthcare institutions can lead to discrepancies that complicate data integration and analysis. For example, thyroid hormone levels may be reported using different units or reference standards, making it difficult to compare values across datasets. Without proper standardisation, these inconsistencies can distort model training and reduce predictive accuracy (40).

The presence of noise and errors in EHR data further complicates the modelling process. Data entry errors, duplicate records, and misclassification of diagnoses are common in large-scale healthcare databases. Such inaccuracies can mislead machine learning algorithms, leading to incorrect pattern recognition and unreliable predictions. While data cleaning and preprocessing techniques can address some of these issues, they require significant effort and expertise, and may not always be able to fully eliminate underlying inaccuracies (41). Another important consideration is the lack of structured data in certain components of EHRs. Clinical notes, which often contain valuable information about patient symptoms and physician observations, are typically recorded in unstructured text format. Extracting meaningful insights from these data requires the use of natural language processing techniques, which introduce additional complexity and potential sources of error. Moreover, the variability in language and terminology used by different clinicians can further complicate the extraction process (42). Temporal irregularity is another characteristic of EHR data that poses challenges for predictive modelling. Unlike controlled experimental datasets, EHR data are collected at irregular intervals, depending on patient visits and clinical needs. This results in uneven time series data, which can complicate the application of temporal models such as recurrent neural networks. Handling

such irregularities requires sophisticated modelling techniques that can accommodate varying time gaps and incomplete sequences (43). Beyond technical challenges, data quality issues also have implications for the generalisability of predictive models. Models trained on data from a specific institution or population may not perform well when applied to different settings, particularly if there are differences in patient demographics, clinical practices, or data recording standards. This lack of external validity limits the broader applicability of AI systems and underscores the need for diverse and representative datasets (44).

## XXV. BIAS IN AI MODELS

Bias in AI models represents one of the most critical challenges in the application of machine learning to healthcare. Bias can arise at multiple stages of the modelling process, including data collection, feature selection, model training, and evaluation. If not properly addressed, these biases can lead to unequal performance across different patient groups, potentially exacerbating existing health disparities. One of the primary sources of bias is the dataset itself. EHR data often reflect the demographics and characteristics of the populations served by specific healthcare institutions. As a result, certain groups may be overrepresented or underrepresented in the data. For example, studies conducted in urban tertiary care centres may predominantly include patients from higher socioeconomic backgrounds, while rural or marginalised populations may be underrepresented. Models trained on such datasets may not accurately capture the health patterns of diverse populations, leading to reduced performance in underrepresented groups (45). Algorithmic bias can also arise from the way features are selected and weighted within the model. Certain variables may have different clinical significance across populations, and failing to account for these differences can lead to biased predictions. For instance, body mass index (BMI) thresholds associated with metabolic risk may vary across ethnic groups, and using a uniform threshold may result in misclassification. Similarly, hormonal profiles and clinical presentations of PCOS may differ across populations, requiring context-specific modelling approaches (46).

Another important source of bias is the historical bias embedded within healthcare systems. EHR data

capture past clinical decisions, which may themselves be influenced by systemic biases or disparities in care. If these biases are not addressed, AI models may inadvertently learn and perpetuate them. For example, if certain groups are less likely to receive diagnostic testing, the absence of data may be interpreted by the model as a lower risk of disease, leading to underdiagnosis in those populations (47). The evaluation of AI models also plays a crucial role in identifying and mitigating bias. Traditional performance metrics such as accuracy and AUC may not fully capture disparities in model performance across subgroups. A model with high overall accuracy may still perform poorly for specific populations, highlighting the need for subgroup analysis and fairness metrics. Techniques such as stratified evaluation and bias auditing are essential for ensuring equitable model performance (48).

Addressing bias in AI models requires a multifaceted approach that includes the use of diverse and representative datasets, careful feature selection, and ongoing monitoring of model performance. Techniques such as re-sampling, re-weighting, and fairness-aware algorithms can help mitigate bias, but their implementation requires careful consideration to avoid unintended consequences. Ultimately, achieving fairness in AI requires not only technical solutions but also a broader commitment to equity in healthcare (49).

## XXVI. INTEGRATION INTO HEALTHCARE SYSTEMS

While the technical development of AI models is a critical component of predictive healthcare, their successful implementation depends on effective integration into existing healthcare systems. This process presents a range of challenges that extend beyond algorithm design, encompassing organisational, technological, and human factors. One of the primary challenges is the integration of AI systems with existing clinical workflows. Healthcare environments are complex and often characterised by time constraints and high workloads. Introducing new technologies into these settings requires careful consideration of how they will be used in practice. AI systems must be designed to complement, rather than disrupt, clinical workflows, providing actionable insights at the point of care without adding to the

cognitive burden of clinicians (50). Interoperability is another significant barrier to integration. Healthcare systems often use different EHR platforms and data standards, which can limit the ability to share and analyse data across institutions. The lack of standardisation in data formats and communication protocols complicates the implementation of AI systems that rely on integrated datasets. Efforts to establish common standards, such as the Fast Healthcare Interoperability Resources (FHIR) framework, represent important خطوات toward addressing this challenge, but widespread adoption remains a work in progress (51).

The issue of clinician acceptance and trust also plays a critical role in the integration of AI systems. Even the most accurate predictive models will have limited impact if they are not trusted and utilised by healthcare professionals. Building trust requires not only high performance but also transparency, interpretability, and evidence of clinical benefit. Explainable AI techniques, as discussed in the previous section, are essential for fostering trust and facilitating adoption (27). Training and education are equally important components of successful integration. Clinicians must be equipped with the knowledge and skills اللازمة to interpret AI outputs and incorporate them into decision-making processes. This may require changes in medical education and the development of new training programs focused on digital health and data science. Without adequate training, there is a risk that AI systems may be underutilised or misinterpreted, potentially compromising patient care (52).

Regulatory and legal considerations also present significant challenges. The use of AI in healthcare raises questions about accountability, liability, and compliance with data protection regulations. Determining who is responsible for decisions influenced by AI systems is a complex issue that requires clear legal frameworks and guidelines. Additionally, ensuring compliance with data privacy regulations is essential for protecting patient information and maintaining public trust (53).

Economic factors must also be considered in the integration of AI systems. The development, implementation, and maintenance of AI technologies require substantial investment, and healthcare institutions must evaluate the cost-effectiveness of these systems. While AI has the potential to reduce long-term costs through improved efficiency and early

intervention, the initial investment and ongoing operational costs may pose barriers to adoption, particularly in resource-limited settings (54).

Finally, the dynamic nature of healthcare data and clinical practice necessitates continuous monitoring and updating of AI models. Models must be regularly retrained to reflect new data and evolving clinical guidelines, ensuring that their predictions remain accurate and relevant. This requires the establishment of robust monitoring systems and governance structures to oversee the lifecycle of AI technologies (55).

#### XXVII. ETHICAL AND REGULATORY CONSIDERATIONS

The integration of artificial intelligence into healthcare, particularly through the use of Electronic Health Records (EHRs), brings patient privacy to the forefront of ethical considerations. Healthcare data are among the most sensitive forms of personal information, encompassing not only medical histories and diagnostic results but also demographic, behavioural, and sometimes genetic data. The use of such data for predictive modelling introduces complex challenges related to consent, ownership, and ethical utilisation. At the core of patient privacy is the principle of autonomy, which emphasises an individual's right to control how their personal information is used. In traditional clinical settings, patients provide consent for the use of their data in diagnosis and treatment. However, the secondary use of EHR data for research and AI development often occurs without explicit, case-specific consent, particularly when data are de-identified. While de-identification reduces the risk of direct identification, it does not entirely eliminate the possibility of re-identification, especially when datasets are large and combined with other sources of information (56).

The concept of informed consent becomes increasingly complex in the context of AI. Patients may not fully understand how their data are being used, particularly when advanced algorithms and large-scale data processing are involved. This raises questions about the adequacy of current consent frameworks, which may not be designed to address the nuances of AI-driven research. There is a growing recognition of the need for more dynamic and transparent consent models, such as broad consent or

dynamic consent, which allow patients to have ongoing control over the use of their data (57). Privacy concerns are particularly relevant in the context of endocrine disorders such as PCOS, where patient data may include sensitive information related to reproductive health, hormonal status, and mental well-being. The potential misuse or unintended disclosure of such information can have significant social and psychological implications, including stigma and discrimination. Ensuring robust privacy protections is therefore not only a technical requirement but also a moral obligation (58). Another dimension of patient privacy is data minimisation, which involves collecting and using only the data necessary for a specific purpose. In AI modelling, there is often a tendency to include as many variables as possible to improve predictive performance. However, this approach must be balanced against the ethical imperative to limit data collection and reduce the risk of privacy breaches. Achieving this balance requires careful consideration of the trade-offs between model accuracy and privacy protection (59).

Furthermore, the increasing use of cloud-based platforms and data-sharing initiatives introduces additional privacy risks. While these technologies enable large-scale data analysis and collaboration, they also increase the potential for unauthorised access and data breaches. Ensuring that patient data remain secure throughout their lifecycle from collection and storage to analysis and sharing is essential for maintaining trust in AI systems (60).

#### XXVIII. DATA SECURITY

Closely related to patient privacy is the issue of data security, which focuses on protecting health data from unauthorised access, breaches, and misuse. As healthcare systems become increasingly digitised and interconnected, the volume and complexity of data being generated and stored have grown exponentially. This expansion creates new vulnerabilities that must be addressed through robust security measures. One of the primary challenges in data security is safeguarding EHR systems against cyber threats. Healthcare organisations are frequent targets of cyberattacks due to the high value of medical data, which can be used for identity theft, insurance fraud, and other malicious activities. Data breaches can have severe consequences, not only for patients but also for

healthcare institutions, leading to financial losses, legal liabilities, and reputational damage (61). To mitigate these risks, a range of technical measures are employed, including encryption, access control, and intrusion detection systems. Encryption ensures that data are protected both at rest and in transit, making it difficult for unauthorised individuals to access sensitive information. Access control mechanisms restrict data access to authorised personnel בלבד, based on predefined roles and permissions. These measures are essential for maintaining the confidentiality and integrity of health data (62).

However, technical solutions alone are not sufficient to ensure data security. Human factors, such as inadequate training, poor password practices, and insider threats, can undermine even the most advanced security systems. Healthcare organisations must therefore adopt a comprehensive approach to security that includes staff training, awareness programs, and organisational policies aimed at promoting secure practices (63). The use of AI itself introduces additional security considerations. Machine learning models can be vulnerable to adversarial attacks, where malicious actors manipulate input data to produce incorrect predictions. In the context of clinical decision-making, such attacks could have serious consequences, potentially leading to incorrect diagnoses or inappropriate treatments. Ensuring the robustness of AI models against such threats is an emerging area of research and a critical component of secure AI deployment (64). Another important aspect of data security is the management of data sharing and interoperability. While the integration of data across systems is essential for comprehensive analysis, it also increases the risk of data exposure. Secure data-sharing frameworks must be established to ensure that information is exchanged safely and in compliance with regulatory standards. Technologies such as blockchain have been proposed as potential solutions for enhancing data security and transparency, although their practical implementation in healthcare is still evolving (65).

#### XXIX. AI REGULATIONS AND GOVERNANCE

The rapid advancement of AI technologies in healthcare has outpaced the development of regulatory frameworks, creating a complex landscape for governance and oversight. Effective regulation is

essential to ensure that AI systems are safe, ethical, and aligned with public interests, while also fostering innovation and technological progress. One of the key challenges in regulating AI is the diversity of applications and technologies involved. AI systems in healthcare range from simple decision support tools to complex predictive models, each with different levels of risk and impact. Regulatory frameworks must therefore be flexible and adaptive, capable of addressing a wide range of use cases while maintaining consistent standards for safety and effectiveness (66).

In many regions, existing medical device regulations are being extended to include AI-based systems. For example, regulatory bodies such as the U.S. Food and Drug Administration (FDA) and the European Medicines Agency (EMA) have begun to develop guidelines for the evaluation and approval of AI-driven medical technologies. These guidelines emphasise the need for rigorous validation, transparency, and ongoing monitoring of AI systems throughout their lifecycle (67). A particularly important aspect of AI regulation is the concept of accountability. Determining responsibility for decisions influenced by AI systems is a complex issue that involves multiple stakeholders, including developers, healthcare providers, and regulatory authorities. Clear guidelines are needed to define roles and responsibilities, ensuring that accountability is maintained even when decisions are supported by automated systems (68). The issue of algorithmic transparency is also central to regulatory efforts. Regulators are increasingly recognising the importance of explainability in AI systems, particularly in high-stakes domains such as healthcare. Requirements for transparency may include the documentation of model development processes, validation results, and mechanisms for interpreting predictions. These measures are intended to ensure that AI systems can be evaluated and trusted by both clinicians and patients (69).

Data protection regulations play a crucial role in governing the use of health data for AI applications. Frameworks such as the General Data Protection Regulation (GDPR) in Europe establish strict guidelines for data processing, including requirements for consent, data minimisation, and the right to explanation. Similar regulations are being developed and implemented in other regions, reflecting a global

effort to protect patient data and ensure ethical use of AI (53,70). Despite these advancements, significant challenges remain in the implementation of AI regulations. The pace of technological innovation often exceeds the ability of regulatory bodies to adapt, leading to gaps in oversight and uncertainty for developers and healthcare providers. Additionally, differences in regulatory frameworks across regions can complicate the global deployment of AI systems, highlighting the need for international collaboration and harmonisation of standards (71). In conclusion, ethical and regulatory considerations are central to the responsible implementation of AI in healthcare. Issues related to patient privacy, data security, and regulatory governance must be carefully addressed to ensure that AI systems are not only effective but also ethical and trustworthy. By establishing robust frameworks and fostering a culture of accountability, the healthcare community can harness the potential of AI while safeguarding the rights and well-being of patients.

### XXX. CONCLUSION

The integration of artificial intelligence into healthcare represents a transformative shift from traditional, reactive models of diagnosis toward proactive, predictive, and personalised care. This study has explored the development and implications of an AI-powered predictive framework for the early diagnosis of Polycystic Ovary Syndrome (PCOS) and thyroid disorders using Electronic Health Record (EHR) data. By bringing together clinical knowledge, data-driven methodologies, and advanced machine learning techniques, the proposed framework highlights the potential of AI to address longstanding challenges in the timely identification and management of endocrine disorders. A key contribution of this work lies in its emphasis on early detection. Both PCOS and thyroid disorders are characterised by gradual onset and complex symptomatology, often leading to delayed diagnosis and suboptimal clinical outcomes. Through the analysis of longitudinal EHR data, the proposed framework demonstrates how subtle changes in hormonal, metabolic, and clinical parameters can be leveraged to predict disease risk at an earlier stage. This shift toward early diagnosis not only improves individual patient outcomes but also aligns with broader public health goals of prevention and disease management (29,32).

The study further underscores the importance of combining multiple modelling approaches to capture the complexity of clinical data. Traditional statistical methods, ensemble machine learning techniques, and deep learning architectures each contribute unique strengths, enabling the development of robust and adaptable predictive systems. In particular, the incorporation of temporal modelling reflects a deeper understanding of disease progression, allowing the system to move beyond static classification toward dynamic prediction (11,14). Equally important is the integration of explainable AI within the predictive framework. As highlighted in earlier sections, interpretability is essential for bridging the gap between computational outputs and clinical decision-making. By employing techniques such as SHAP and LIME, the framework ensures that predictions are not only accurate but also transparent and clinically meaningful. This fosters trust among healthcare professionals and supports the responsible adoption of AI technologies in clinical practice (22,23).

At the same time, the study acknowledges the significant challenges associated with implementing AI in healthcare. Issues related to data quality, bias, system integration, and ethical considerations present substantial barriers that must be addressed to ensure the effectiveness and equity of AI systems. The discussion of these challenges emphasises the need for a holistic approach that combines technical innovation with ethical responsibility, regulatory oversight, and stakeholder engagement (44,49,53).

Overall, this research highlights the potential of AI-powered predictive modelling to transform the diagnosis and management of endocrine disorders. By leveraging the richness of EHR data and the capabilities of modern machine learning, the proposed framework offers a pathway toward more timely, accurate, and personalised healthcare. However, the realisation of this potential depends on continued research, collaboration, and the development of robust systems that can be effectively integrated into clinical practice. While the proposed AI framework provides a strong foundation for early diagnosis of PCOS and thyroid disorders, it also opens several avenues for future research and development. One of the most important directions lies in the expansion and diversification of data sources. Future studies can incorporate additional data modalities, such as genomic data, wearable device data, and patient-

reported outcomes, to enhance the predictive capabilities of AI models. The integration of multimodal data has the potential to provide a more comprehensive understanding of disease mechanisms and improve the accuracy of predictions (72).

Another promising area for future research is the development of real-time predictive systems. Current models often rely on retrospective data analysis, but advances in data infrastructure and computational capabilities make it possible to implement real-time monitoring and prediction. Such systems could continuously analyse incoming patient data, providing timely alerts and recommendations to clinicians. This would further strengthen the shift toward proactive and preventive healthcare (73). The refinement of temporal modelling techniques also represents an important area for future exploration. While models such as Long Short-Term Memory networks have demonstrated strong performance in handling sequential data, there is ongoing research into more advanced architectures, such as transformer-based models, which may offer improved capabilities for capturing long-range dependencies in time-series data. Applying these models to EHR data could further enhance the ability to detect early patterns of disease progression (74). In addition to technical advancements, future work must focus on improving the interpretability and usability of AI systems. Developing intuitive visualisation tools and user interfaces that translate complex model outputs into clinically meaningful insights will be critical for facilitating adoption. Collaboration between data scientists, clinicians, and designers will play a key role in ensuring that AI systems are both effective and user-friendly (27).

The issue of generalisability also warrants further investigation. Future studies should aim to validate predictive models across diverse populations and healthcare settings, ensuring that they perform consistently and equitably. This may involve the creation of large, multi-institutional datasets and the use of federated learning approaches, which allow models to be trained on distributed data without compromising patient privacy (75). Ethical and regulatory considerations will continue to shape the future of AI in healthcare. As technologies evolve, there will be a need for adaptive regulatory frameworks that can keep pace with innovation while ensuring patient safety and data protection. Research

into ethical AI design, fairness, and accountability will remain essential for building trustworthy systems that are aligned with societal values (71).

Finally, the translation of AI research into clinical practice represents a critical step in realising its full potential. This will require not only technological development but also organisational change, including the integration of AI into clinical workflows, the training of healthcare professionals, and the establishment of supportive policies and infrastructure. Pilot studies and real-world implementations will be essential for demonstrating the clinical and economic value of AI systems and driving their adoption at scale (50).

## REFERENCES

- [1] E. W. Boyer and M. Shannon, "The serotonin syndrome," *New England Journal of Medicine*, vol. 352, pp. 1112–1120, 2005.
- [2] FDA, "FDA Drug Safety Communication: Linezolid and serotonergic psychiatric medications."
- [3] K. R. Lawrence, M. Adra, and P. K. Gillman, "Serotonin toxicity associated with linezolid," *Annals of Pharmacotherapy*, vol. 40, pp. 119–123, 2006.
- [4] J. J. Taylor, J. W. Wilson, and L. L. Estes, "Linezolid and serotonergic drug interactions," *Clinical Infectious Diseases*, vol. 42, pp. 1578–1583, 2006.
- [5] P. K. Gillman, "Monoamine oxidase inhibitors and serotonin toxicity," *British Journal of Pharmacology*, vol. 147, suppl. 1, pp. S287–S296, 2006.
- [6] E. J. Topol, "High-performance medicine: The convergence of human and artificial intelligence," *Nature Medicine*, vol. 25, pp. 44–56, 2019.
- [7] P. Rajpurkar, E. Chen, O. Banerjee, and E. J. Topol, "AI in health and medicine," *Nature Medicine*, vol. 28, pp. 31–38, 2022.
- [8] J. Adler-Milstein, C. M. DesRoches, P. Kralovec, et al., "Adoption of electronic health record systems," *Health Affairs*, vol. 33, pp. 1672–1680, 2014.
- [9] S. van Buuren and K. Groothuis-Oudshoorn, "MICE: Multivariate imputation by chained equations," *Journal of Statistical Software*, vol. 45, pp. 1–67, 2011.
- [10] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proc. ACM SIGKDD*

- Int. Conf. on Knowledge Discovery and Data Mining (KDD), 2016, pp. 785–794.
- [11] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural Computation*, vol. 9, pp. 1735–1780, 1997.
- [12] Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” *Nature*, vol. 521, pp. 436–444, 2015.
- [13] J. Futoma, M. Simons, T. Panch, F. Doshi-Velez, and L. A. Celi, “The myth of generalisability in clinical ML,” *The Lancet Digital Health*, vol. 2, pp. e489–e492, 2020.
- [14] E. Choi, M. T. Bahadori, A. Schuetz, W. F. Stewart, and J. Sun, “Doctor AI,” 2016.
- [15] R. Miotto, F. Wang, S. Wang, X. Jiang, and J. T. Dudley, “Deep learning for healthcare,” *Briefings in Bioinformatics*, vol. 19, pp. 1236–1246, 2018.
- [16] A. Esteva, B. Kuprel, R. A. Novoa, et al., “Dermatologist-level classification of skin cancer with deep neural networks,” *Nature*, vol. 542, pp. 115–118, 2017.
- [17] F. Doshi-Velez and B. Kim, “Towards a rigorous science of interpretable machine learning,” 2017.
- [18] Z. C. Lipton, “The mythos of model interpretability,” *Queue*, vol. 16, pp. 31–57, 2018.
- [19] M. T. Ribeiro, S. Singh, and C. Guestrin, “Why should I trust you?: Explaining the predictions of any classifier,” in *Proc. ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining (KDD)*, 2016.
- [20] D. S. Char, N. H. Shah, and D. Magnus, “Implementing machine learning in healthcare—Addressing ethical challenges,” *New England Journal of Medicine*, vol. 378, pp. 981–983, 2018.
- [21] E. J. Topol, *Deep Medicine*. New York, NY, USA: Basic Books, 2019.
- [22] S. M. Lundberg and S.-I. Lee, “A unified approach to interpreting model predictions,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [23] M. T. Ribeiro, S. Singh, and C. Guestrin, “Model-agnostic explanations (LIME),” in *Proc. ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining (KDD)*, 2016.
- [24] D. Bahdanau, K. Cho, and Y. Bengio, “Neural machine translation by jointly learning to align and translate,” 2015.
- [25] G. Elwyn, D. Frosch, R. Thomson, et al., “Shared decision making: A model for clinical practice,” *Journal of General Internal Medicine*, vol. 27, pp. 1361–1367, 2012.
- [26] P. Rajpurkar et al., “AI in healthcare education,” *Nature Medicine*, 2022.
- [27] S. Tonekaboni, S. Joshi, M. D. McCradden, and A. Goldenberg, “Clinician perspectives on explainable AI,” *Nature Medicine*, vol. 25, pp. 1089–1098, 2019.
- [28] A. B. Arrieta, N. Díaz-Rodríguez, J. Del Ser, et al., “Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI,” *Information Fusion*, vol. 58, pp. 82–115, 2020.
- [29] H. J. Teede, M. L. Misso, M. F. Costello, et al., “Recommendations from the international evidence-based guideline for the assessment and management of polycystic ovary syndrome,” *Human Reproduction*, vol. 33, pp. 1602–1618, 2018.
- [30] A. Esteva et al., “AI in clinical screening,” *Nature*, 2017.
- [31] Dokras, S. Clifton, W. Futterweit, and R. Wild, “Quality of life in women with polycystic ovary syndrome: A systematic review and meta-analysis,” *Journal of Clinical Endocrinology & Metabolism*, vol. 96, pp. E973–E980, 2011.
- [32] P. N. Taylor, D. Albrecht, A. Scholz, et al., “Global epidemiology of thyroid disease,” *The Lancet Diabetes & Endocrinology*, vol. 6, pp. 1–14, 2018.
- [33] A. L. Beam and I. S. Kohane, “Big data and machine learning in healthcare,” *JAMA*, vol. 319, pp. 1317–1318, 2018.
- [34] P. J. Neumann, J. T. Cohen, and M. C. Weinstein, “Updating cost-effectiveness—The curious resilience of the \$50,000-per-QALY threshold,” *New England Journal of Medicine*, vol. 374, pp. 203–205, 2016.
- [35] L. G. Cooney, I. Lee, M. D. Sammel, and A. Dokras, “Mental health in PCOS,” *Clinical Endocrinology*, vol. 86, pp. 770–776, 2017.
- [36] F. S. Collins and H. Varmus, “A new initiative on precision medicine,” *New England Journal of Medicine*, vol. 372, pp. 793–795, 2015.
- [37] E. J. Topol, “Individualized medicine,” *Nature Medicine*, 2019.
- [38] Z. Obermeyer, B. Powers, C. Vogeli, and S. Mullainathan, “Dissecting racial bias in an algorithm used to manage the health of populations,” *Science*, vol. 366, pp. 447–453, 2019.
- [39] N. G. Weiskopf and C. Weng, “Methods and dimensions of electronic health record data quality assessment: Enabling reuse for clinical research,” *Journal of the American Medical Association*, 2019.

- Informatics Association, vol. 20, pp. 144–151, 2013.
- [40] M. G. Kahn, T. J. Callahan, J. Barnard, et al., “A harmonized data quality assessment terminology and framework for the secondary use of electronic health record data,” *EGEMS*, vol. 4, p. 1249, 2016.
- [41] T. Botsis, G. Hartvigsen, F. Chen, and C. Weng, “Secondary use of EHR: Data quality issues and informatics opportunities,” *Journal of Biomedical Informatics*, vol. 43, pp. 1–5, 2010.
- [42] Y. Wang, L. Wang, M. Rastegar-Mojarad, et al., “Clinical information extraction applications: A literature review,” *Journal of Biomedical Informatics*, vol. 77, pp. 34–49, 2018.
- [43] Z. C. Lipton, D. C. Kale, and C. Elkan, “Learning from irregularly sampled time series: A missing data perspective,” 2016.
- [44] J. Wiens, S. Saria, M. Sendak, et al., “Do no harm: A roadmap for responsible machine learning for health care,” *Nature Medicine*, vol. 25, pp. 1337–1340, 2019.
- [45] Y. Chen, F. D. Johansson, and D. Sontag, “Why is my classifier discriminatory?” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [46] Z. Obermeyer et al., “Dissecting racial bias in an algorithm used to manage the health of populations,” *Science*, vol. 366, pp. 447–453, 2019.
- [47] M. A. Gianfrancesco, S. Tamang, J. Yazdany, and G. Schmajuk, “Potential biases in machine learning algorithms using electronic health record data,” *JAMA*, vol. 319, pp. 204–205, 2018.
- [48] A. Rajkomar, M. Hardt, M. D. Howell, et al., “Ensuring fairness in machine learning to advance health equity,” *New England Journal of Medicine*, vol. 378, pp. 981–983, 2018.
- [49] N. Mehrabi, F. Morstatter, N. Saxena, et al., “A survey on bias and fairness in machine learning,” *ACM Computing Surveys*, vol. 54, pp. 1–35, 2021.
- [50] E. H. Shortliffe and M. J. Sepúlveda, “Clinical decision support in the era of artificial intelligence,” *JAMA*, vol. 320, pp. 2199–2200, 2018.
- [51] C. Mandel, D. A. Kreda, K. D. Mandl, et al., “SMART on FHIR: A standards-based, interoperable apps platform for electronic health records,” *Journal of the American Medical Informatics Association*, vol. 23, pp. 899–908, 2016.
- [52] Paranjape, M. Schinkel, R. N. Panday, et al., “Introducing artificial intelligence training in medical education,” *The Lancet Digital Health*, vol. 1, pp. e39–e40, 2019.
- [53] W. Goodman and S. Flaxman, “European Union regulations on algorithmic decision-making and a ‘right to explanation’,” *AI Magazine*, vol. 38, pp. 50–57, 2017.
- [54] F. Jiang, Y. Jiang, H. Zhi, et al., “Artificial intelligence in healthcare: Past, present and future,” *Stroke and Vascular Neurology*, vol. 2, pp. 230–243, 2017.
- [55] P. Sendak, M. Gao, M. Nichols, et al., “A path for translation of machine learning products into healthcare delivery,” *NPJ Digital Medicine*, vol. 3, pp. 1–9, 2020.
- [56] Gymrek, A. L. McGuire, D. Golan, et al., “Identifying personal genomes by surname inference,” *Science*, vol. 339, pp. 321–324, 2013.
- [57] J. Kaye, E. A. Whitley, D. Lund, et al., “Dynamic consent: A patient interface for twenty-first century research networks,” *European Journal of Human Genetics*, vol. 23, pp. 141–146, 2015.
- [58] E. Vayena, A. Blasimme, and I. G. Cohen, “Machine learning in medicine: Addressing ethical challenges,” *PLoS Medicine*, vol. 15, p. e1002689, 2018.
- [59] D. Mittelstadt, P. Allo, M. Taddeo, et al., “The ethics of algorithms: Mapping the debate,” *Big Data & Society*, vol. 3, pp. 1–21, 2016.
- [60] J. M. Rumbold and B. K. Pierscionek, “The effect of the general data protection regulation on medical research,” *BMC Medical Ethics*, vol. 18, pp. 1–9, 2017.
- [61] S. Kruse, B. Frederick, T. Jacobson, and D. K. Monticone, “Cybersecurity in healthcare: A systematic review of modern threats and trends,” *JMIR Medical Informatics*, vol. 5, p. e23, 2017.
- [62] A. McLeod and D. Dolezel, “Cyber-analytics: Modeling factors associated with healthcare data breaches,” *Journal of Healthcare Information Management*, 2018.
- [63] J. L. Fernandez-Aleman, I. C. Señor, P. Á. Lozoya, and A. Toval, “Security and privacy in electronic health records: A systematic literature review,” *Journal of Biomedical Informatics*, vol. 46, pp. 541–562, 2013.
- [64] S. G. Finlayson, J. D. Bowers, J. Ito, et al., “Adversarial attacks on medical machine learning,” *Science*, vol. 363, pp. 1287–1289, 2019.
- [65] C. Agbo, Q. H. Mahmoud, and J. M. Eklund, “Blockchain technology in healthcare: A

- systematic review,” *Healthcare Informatics Research*, vol. 25, pp. 163–174, 2019.
- [66] W. N. Price and I. G. Cohen, “Privacy in the age of medical big data,” *Yale Journal of Law and Technology*, vol. 21, pp. 88–123, 2019.
- [67] FDA, “Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device,” 2021.
- [68] L. Floridi, J. Cowls, M. Beltrametti, et al., “AI4People—An ethical framework for a good AI society,” *Minds and Machines*, vol. 28, pp. 689–707, 2018.
- [69] European Commission, *Ethics Guidelines for Trustworthy AI*, 2019.
- [70] Voigt and A. Von dem Bussche, *The EU General Data Protection Regulation (GDPR)*, 2017.
- [71] J. Morley, C. C. V. Machado, C. Burr, et al., “The ethics of AI in health care: A mapping review,” *Science and Engineering Ethics*, vol. 26, pp. 1433–1450, 2020.
- [72] Y. Hasin, M. Seldin, and A. Lusic, “Multi-omics approaches to disease,” *Genome Biology*, vol. 18, p. 83, 2017.
- [73] W. Bates, S. Saria, L. Ohno-Machado, et al., “Big data in health care: Using analytics to identify and manage high-risk and high-cost patients,” *Health Affairs*, vol. 33, pp. 1123–1131, 2014.
- [74] A. Vaswani, N. Shazeer, N. Parmar, et al., “Attention is all you need,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [75] N. Rieke, J. Hancox, W. Li, et al., “The future of digital health with federated learning,” *Nature Machine Intelligence*, vol. 2, pp. 305–311, 2020.