WJBPHS

World Journal of Biology Pharmacy and Health Sciences

eISSN: 2582-5542

Cross Ref DOI: 10.30574/wjbphs

Journal homepage: https://wjbphs.com/

(Review Article)

Explainable AI for healthcare professionals: Advancing risk assessment, diagnostic precision, and ethical clinical interventions

Suneel Pappala [1, *], M. Malyadri [2], K Venkata Naganjaneyulu [3], A. Prakashini [4] and Pasuladi Santosh [5]

[1] Associate Professor, Artificial Intelligence and Data Science, St. Mary's Group of Institutions Hyderabad (Autonomous), JNTU-Hyderabad, Telangana State, India.

[2] Associate Professor, Computer Science and Engineering, St. Mary's Group of Institutions Hyderabad (Autonomous), JNTU-Hyderabad, Telangana State,

[3] Professor, CSE Dept, Malla Reddy Engineering College for Women (Autonomous), JNTU-Hyderabad. Telangana State, India.

[4] Assistant Professor, CSE Department, Avanthi Institute of Technology and Science(Autonomous), JNTU-Hyderabad. Telangana State, India.

[5] Associate Professor, Artificial Intelligence and Machine Learning, St. Mary's Group of Institutions Hyderabad (Autonomous), JNTU-Hyderabad, Telangana State, India.

World Journal of Biology Pharmacy and Health Sciences, 2025, 22(01), 446-453

Publication history: Received on 11 March 2025; revised on 19 April 2025; accepted on 21 April 2025

Article DOI: https://doi.org/10.30574/wjbphs.2025.22.1.0425

Abstract

Artificial intelligence (AI) becomes increasingly integrated into healthcare, the demand for transparency, trust, and accountability in AI-driven decisions has grown significantly. Explainable AI (XAI) provides a solution by making complex AI models interpretable and understandable for clinicians, patients, and regulators. This explores the role of XAI in enhancing predictive accuracy and improving risk assessment in clinical environments. By offering insights into how AI models arrive at diagnoses, treatment plans, and patient risk scores, XAI facilitates safer and more informed medical decision-making. Techniques such as SHAP, LIME, attention mechanisms, and saliency maps are highlighted for their ability to clarify AI behaviour at both the global and local levels. In addition to supporting clinical trust and regulatory compliance, XAI also plays a crucial role in bias detection, ensuring fairness across diverse patient populations. The integration of XAI promotes a human-cantered approach to healthcare AI, enabling collaborative decision-making where medical professionals can validate, adjust, or override AI outputs. As healthcare systems increasingly rely on predictive algorithms for early diagnosis and preventive care, XAI emerges as a critical component in achieving ethical, accurate, and equitable outcomes.

Keywords: Explainable Artificial Intelligence (XAI); Healthcare AI; Transparency; Medical Imaging; Risk Assessment

1. Introduction

Explainable AI (XAI) encompasses a range of techniques aimed at making the decisions and predictions of AI systems transparent and understandable to humans. Its primary goals include enhancing trust, detecting and mitigating bias, ensuring regulatory compliance, and aiding model debugging and improvement. XAI addresses both global explainability understanding overall model behaviour, as seen in inherently interpretable models like decision trees and local explainability, which focuses on individual predictions using methods like LIME and SHAP. Common XAI techniques include feature importance analysis, counterfactual explanations, and attention mechanisms, each offering insights into how AI models reach their conclusions. However, challenges remain, such as the trade-off between model

[*] Corresponding author: Suneel Pappala. Email: suneelpappala@gmail.com

accuracy and interpretability, the scalability of explanation techniques, and ensuring explanations are accessible to non-experts.

Explainable AI (XAI) in healthcare plays a crucial role in ensuring that AI-driven medical decisions are transparent, interpretable, and trustworthy for clinicians, patients, and regulators. As AI becomes increasingly integrated into diagnostics, treatment planning, and risk assessment, XAI helps build trust and facilitates adoption by allowing medical professionals to understand and verify AI decisions. It enhances patient safety by reducing the risk of incorrect diagnoses, supports regulatory compliance with laws like GDPR and FDA guidelines, and ensures accountability in clinical settings. In disease diagnosis and prediction particularly in cancer detection—XAI enables doctors to interpret how AI algorithms identify patterns in medical images or patient data. Techniques such as SHAP, Grad-CAM, and LIME help visualize and explain AI predictions by highlighting key image regions, biomarkers, or features that influenced the diagnosis. This level of transparency improves diagnostic accuracy, reduces misdiagnosis risks, supports personalized treatment by revealing cancer-specific drivers, and aids in the development of new therapies. Ultimately, XAI bridges the gap between complex AI models and human decision-makers, making AI a more reliable and ethical tool in modern healthcare.

Transparency in AI Models for Healthcare: Transparency in Explainable AI (XAI) is essential in healthcare, as it ensures that AI models clearly communicate the rationale behind their diagnoses, treatment suggestions, and risk assessments. Given that medical decisions directly affect patient outcomes, clinicians must be able to understand and trust AI-generated insights before incorporating them into care. Transparent AI fosters trust and increases adoption among healthcare professionals by making the decision-making process understandable and justifiable. It also supports regulatory compliance with standards such as GDPR, FDA, and HIPAA, which require interpretability and accountability in AI systems. Moreover, transparency enhances patient safety by minimizing the risks associated with opaque, black-box models that could lead to misdiagnoses or biased recommendations. Lastly, it enables legal and ethical accountability, allowing AI decisions to be audited, questioned, and refined to ensure responsible and equitable healthcare delivery.

AI can enhance trust and safety in healthcare by explaining how it arrives at diagnoses, treatment plans, and risk assessments. In disease diagnosis, AI models especially those used for detecting conditions like cancer, heart disease, or Alzheimer’s should offer transparent reasoning behind their predictions. Instead of merely outputting results, models should use techniques like saliency maps to highlight key areas in medical images, attention mechanisms to pinpoint influential data such as symptoms or test results, and feature importance scores to show which patient parameters were most critical. For treatment plans, AI recommendations must be backed by clear justifications. For example, if chemotherapy is recommended over radiation, the model should explain this choice based on clinical factors like tumour characteristics, patient history, or genetic markers. Transparency in treatment planning can be achieved through referencing relevant clinical studies, using counterfactual explanations to show how different variables would alter the recommendation, and incorporating a human-in-the-loop approach that enables physicians to review, question, and refine AI outputs. This level of interpretability ensures that AI serves as a reliable partner in clinical decision-making rather than a black-box authority.

Risk Assessment Transparency: AI models used for predicting disease risks such as heart attacks, strokes, or diabetes must provide clear explanations for why a patient is classified as high or low risk. Transparency in these predictions is essential for building clinician and patient trust, enabling informed medical decisions, and ensuring ethical use of AI in preventive care. For instance, a cardiovascular risk prediction model should clearly show how factors like high cholesterol, smoking history, and age influenced the patient’s risk score. Techniques like SHAP (Shapley Additive Explanations) can break down the individual contributions of each factor to the overall prediction, offering a detailed view of the model’s reasoning. LIME (Local Interpretable Model-agnostic Explanations) simplifies complex models by creating interpretable approximations that reflect local prediction behaviour, making it easier to understand why a particular outcome was predicted. Additionally, decision trees and rule-based models can provide step-by-step, logical pathways that mimic human decision-making, offering intuitive insights into risk assessments. Together, these approaches help demystify AI predictions and support transparent, data-driven healthcare.

Interpretability in AI for Healthcare: Interpretability in Explainable AI (XAI) is essential in healthcare, as it ensures that AI model outputs can be understood and validated by medical professionals in the context of their clinical expertise. While transparency reveals how an AI model functions internally, interpretability focuses on translating AI decisions into medically meaningful insights that doctors, nurses, and researchers can act upon. This is crucial for enhancing clinical decision-making, as AI predictions must align with established medical knowledge to be trusted and useful. Interpretability also fosters confidence among healthcare providers, enabling them to understand and question AI recommendations rather than blindly following them. Furthermore, it supports compliance with regulations such as

GDPR and FDA guidelines, which mandate that AI systems be explainable and accountable. Importantly, interpretability can help reduce medical errors by allowing practitioners to detect and correct potential mistakes or biases in AI outputs, ultimately leading to safer and more effective patient care.

To make AI interpretable for medical professionals, models must present their outputs using clinically meaningful features and visual explanations that align with established medical knowledge. Rather than relying on abstract or opaque representations, AI systems should utilize familiar indicators like cholesterol levels, blood pressure, or ECG readings when predicting conditions such as heart attack risk. This feature-based interpretability can be achieved through methods like feature importance scores, which identify the most influential clinical parameters, and rule-based models such as decision trees or expert systems that follow transparent, guideline-driven logic. For AI applications in medical imaging—such as analysing X-rays, CT scans, or MRIs visual interpretability is crucial. These models should clearly highlight the specific regions that contributed to a diagnosis, such as outlining a suspected tumour on a CT scan rather than merely providing a risk score. Techniques like saliency maps (e.g., Grad-CAM or Layer-wise Relevance Propagation) and attention mechanisms enable this by visually indicating where the AI focused during its analysis. These interpretability approaches empower clinicians to validate AI outputs, integrate them into their workflow, and make informed decisions with confidence.

Case-based reasoning enhances the interpretability of AI in healthcare by allowing medical professionals to compare current patient cases with similar ones from the past. Clinicians are more likely to trust AI when they can see how previous patients with comparable symptoms, lab results, or medical histories were diagnosed or treated. For instance, a diabetes risk model could present a physician with three prior cases involving patients who had similar blood sugar levels, BMI, and lifestyle factors, along with the outcomes and treatments those patients received. This approach supports interpretability through methods like nearest neighbour comparisons, which identify and retrieve similar cases from a database, and counterfactual explanations, which demonstrate how slight changes in a patient's profile might alter the risk assessment or treatment plan. Additionally, AI risk prediction models should go beyond providing a single score by offering step-by-step reasoning that breaks down how each factor such as age, cholesterol, or family history contributed to the overall risk. This layered, example-based explanation makes AI outputs more relatable and actionable for clinicians, reinforcing trust and improving decision-making.

For example, a stroke prediction model should provide a clear breakdown of how individual risk factors contribute to the overall assessment. Instead of simply outputting a 75% risk score, the model should explain the components behind that number: high blood pressure might increase the risk by 30%, smoking history by 20%, and being over the age of 60 by another 25%. This kind of step-by-step reasoning allows medical professionals to understand the weight of each contributing factor, validate the model's logic, and make informed decisions tailored to the patient's specific risk profile. Such detailed interpretability enhances trust, supports patient communication, and aligns AI predictions with clinical reasoning.

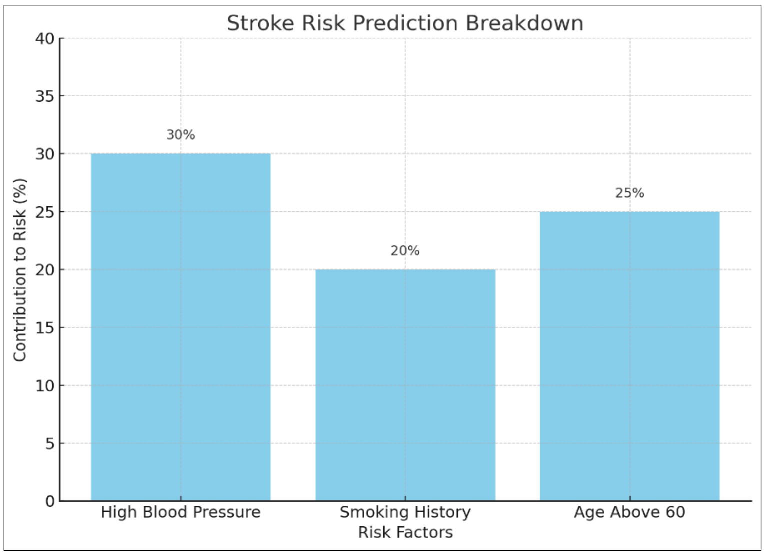

Figure 1 Stroke Risk Prediction Breakdown

Techniques for Step-by-Step Interpretability: SHAP (Shapley Additive Explanations) and LIME (Local Interpretable Model-Agnostic Explanations) are two powerful tools that help make AI predictions more understandable by breaking down how much each factor contributed or by building simplified models that mimic the AI's decision-making in easy-to-understand terms. These tools play a key role in enabling effective human-AI collaboration, where the goal is not to replace doctors but to support and enhance their clinical judgment. An interpretable AI system should clearly explain its reasoning, offer options for doctors to override or adjust predictions, and allow for feedback that can refine and improve future model performance. This collaboration is facilitated through interactive AI dashboards that let healthcare providers explore and interrogate predictions, doctor feedback loops that allow them to correct inaccuracies, and workflows that ensure AI-assisted diagnoses are ultimately reviewed and approved by qualified medical professionals. By combining the strengths of AI with human expertise, healthcare outcomes can be both safer and more accurate.

Accountability in AI for Healthcare: Accountability in Explainable AI (XAI) is essential in healthcare to ensure that AI-assisted decisions are transparent, traceable, and subject to expert oversight. In high-stakes scenarios such as diagnoses, treatment planning, and risk assessments, accountability safeguards patient safety, supports regulatory and ethical standards, and builds trust among healthcare providers. AI systems must offer decision traceability clearly recording how they reached conclusions. For instance, an AI diagnosing pneumonia from an X-ray should log the specific image features it flagged, reference the medical guidelines it used, and include a confidence score. Techniques like model logging, audit trails, and version control are critical for enabling doctors to review, verify, and audit AI decisions when necessary. Furthermore, human-in-the-loop systems are crucial for maintaining accountability. These systems ensure that medical professionals not AI retain the final say in patient care. For example, if an AI flags a potential tumour on an MRI, a radiologist must review the evidence and either confirm or override the decision. If they override it, their input feeds back into the system for continuous improvement. This collaborative approach ensures that AI enhances, rather than replaces, clinical judgment, leading to more ethical, safe, and effective healthcare delivery.

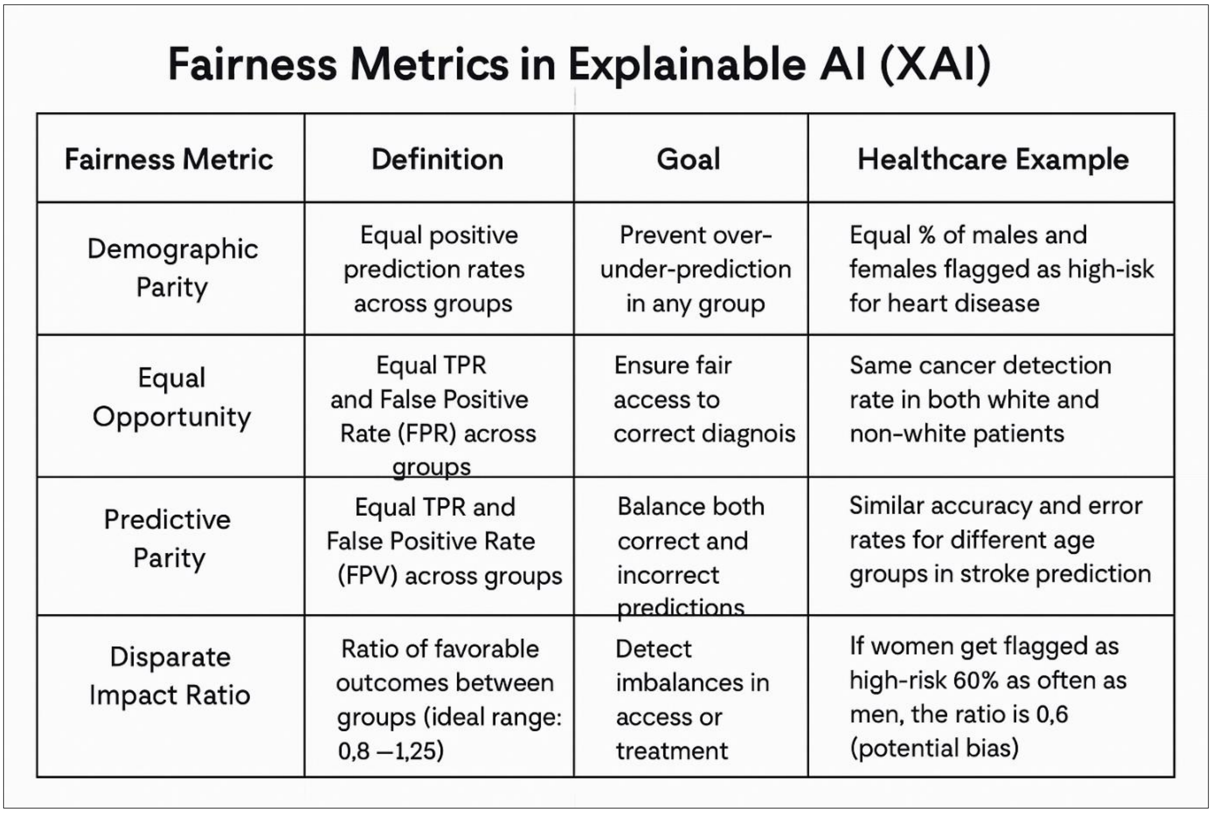

Fairness Metrics in Explainable AI (XAI)

Fairness Metric	Definition	Goal	Healthcare Example
Demographic Parity	Equal positive prediction rates across groups	Prevent over-under-prediction in any group	Equal % of males and females flagged as high-isk for heart disease
Equal Opportunity	Equal TPR and False Positive Rate (FPR) across groups	Ensure fair access to correct diagnois	Same cancer detection rate in both white and non-white patients
Predictive Parity	Equal TPR and False Positive Rate (FPV) across groups	Balance both correct and incorrect predictions	Similar accuracy and error rates for different age groups in stroke prediction
Disparate Impact Ratio	Ratio of favorable outcomes between groups (ideal range: 0,8 –1,25)	Detect imbalances in access or treatment	If women get flagged as high-risk 60% as often as men, the ratio is 0,6 (potential bias)

Figure 2 Fairness Metrics in Explainable AI (XAI)

Implementing human oversight in healthcare AI systems is essential to ensure safety, accuracy, and ethical integrity. One key approach is AI-assisted decision-making, where AI provides recommendations, but the final decision always rests with a medical professional. Interactive AI systems further support this by allowing clinicians to adjust parameters—such as patient variables or thresholds and observe how those changes impact outcomes, giving them deeper control and understanding. Additionally, using second opinion AI models, where multiple systems analyse the same case independently, enhances diagnostic confidence and helps catch errors or inconsistencies. To ensure fairness, AI models must also undergo regular bias detection and fairness audits. These evaluations help identify whether certain patient groups based on race, gender, age, or socioeconomic background are being unfairly treated or misdiagnosed. By continuously monitoring and addressing these issues, healthcare AI can remain both equitable and trustworthy, supporting better outcomes for all patients.

Legal and regulatory compliance in Explainable AI (XAI) is essential to ensure that healthcare AI systems operate within established ethical, legal, and safety frameworks. AI models must align with regulations such as GDPR, HIPAA, and the guidelines of authorities like the FDA, EMA, or MHRA. For instance, an AI system recommending medications must not only provide accurate prescriptions but also document its reasoning process to satisfy legal standards. Compliance is achieved through regulatory certifications, data privacy adherence, and oversight by ethical review boards that assess the safety and fairness of AI outputs. Additionally, continuous learning and error correction are vital for maintaining high performance and trust. AI systems must incorporate feedback from medical professionals, retrain using corrected data, and generate explainability reports to track and improve accuracy. Adaptive learning ensures that the AI evolves with new patient outcomes, making it more reliable over time and better aligned with clinical needs and legal requirements.

Fairness & Bias Detection in Healthcare AI: Fairness and bias detection in healthcare AI ensures that medical decisions generated by AI systems are accurate, equitable, and free from discriminatory influence across diverse patient populations. Since AI models are often trained on historical healthcare data, they can unintentionally inherit and amplify existing biasesleading to unequal treatment outcomes for certain demographic groups. Ensuring fairness is therefore essential for ethical and effective AI deployment. It helps prevent discriminatory outcomes based on race, gender, age, or socioeconomic status, and improves the accuracy of predictions across all populations. Fair AI fosters greater trust among both patients and medical professionals, encouraging adoption and collaboration. Additionally, fairness is a core component of regulatory compliance under laws such as GDPR, HIPAA, and FDA guidelines. By actively identifying and mitigating biases, healthcare AI can contribute to reducing disparities and improving healthcare access and outcomes for traditionally underserved or marginalized communities.

Demographic Parity (Statistical Parity) Artificial Intelligence: Demographic Parity, also known as Statistical Parity, ensures that the outcomes produced by an AI model are not influenced by protected attributes such as race, gender, or age. In healthcare, this means the AI should approve or recommend treatments at the same rate across different demographic groups, regardless of their identity. For example, if 60% of male patients are approved for a certain therapy, then approximately 60% of female patients should be approved too assuming both groups are equally qualified. This type of fairness check is crucial for preventing systemic bias and ensuring equitable healthcare delivery. However, it may sometimes conflict with other fairness goals, especially if health conditions naturally vary across groups, so it's essential to consider demographic parity alongside clinical justifications.

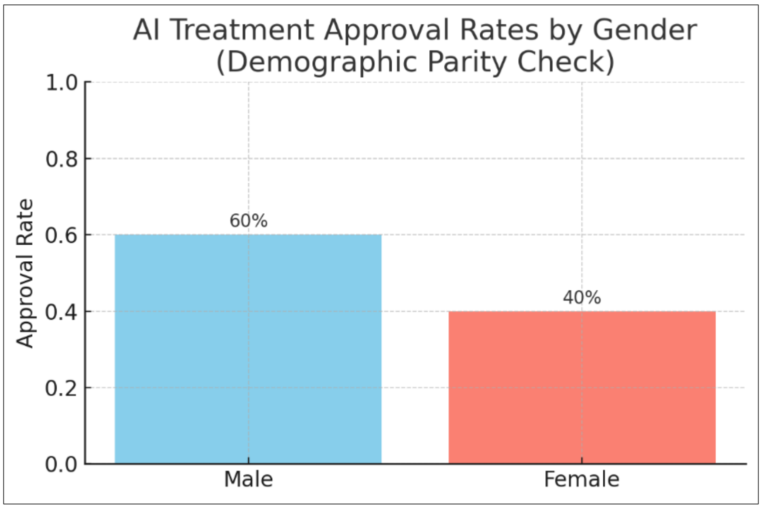

Figure 3 Artificial Intelligence Approval Rates

Drug Discovery and Development Using Explainable AI (XAI): Explainable AI is transforming drug discovery and development by enhancing transparency, speeding up research, and ensuring more informed decision-making across all stages from target identification to clinical trials. Traditional drug discovery can take over a decade and cost billions of dollars, but AI models can rapidly analyse massive datasets (e.g., chemical properties, biological interactions, clinical outcomes) to identify promising compounds. With explainability, researchers and regulators can understand why certain molecules were selected, how potential side effects were predicted, or why a candidate drug may succeed or fail. Techniques like SHAP and attention mechanisms allow scientists to trace predictions back to specific molecular features or biological markers, reducing uncertainty and boosting trust in AI-generated drug candidates. Furthermore, XAI helps meet strict regulatory standards by providing clear, interpretable evidence of how and why decisions are made at each development phase.

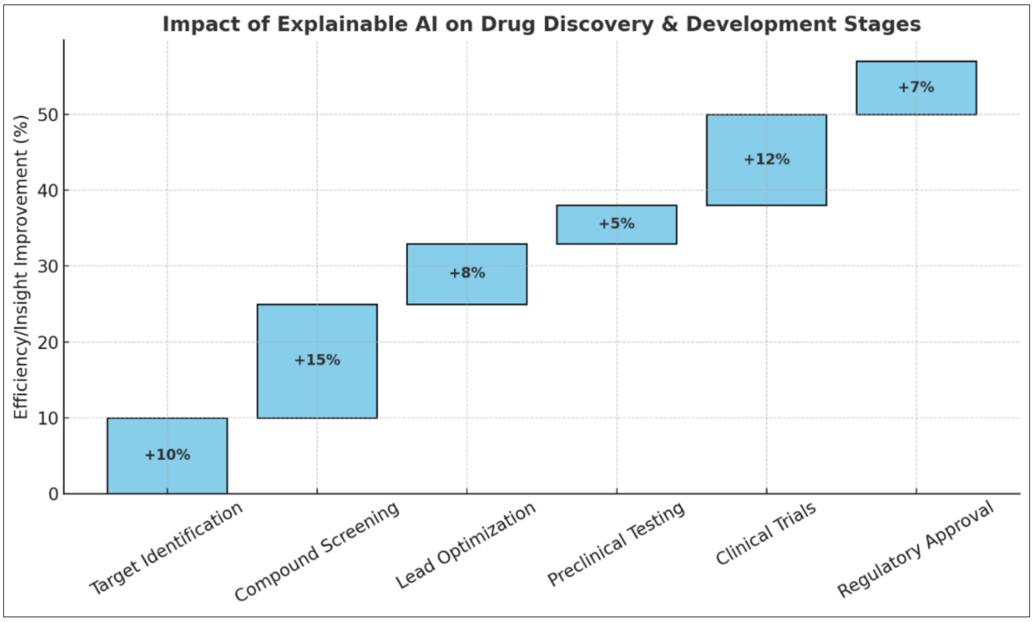

Figure 4 Drug Discovery and Development Stages

Treatment Recommendation Systems with Explainable AI: Treatment recommendation systems use artificial intelligence to assist healthcare professionals in suggesting personalized treatment plans for patients based on their medical history, diagnostics, and clinical guidelines. With the integration of Explainable AI (XAI), these systems not only provide recommendations but also explain the reasoning behind each choice. This is critical in healthcare, where trust, accountability, and transparency are essential for clinical decision-making. Explainable treatment recommendation systems analyze various patient features such as age, lab results, genetic markers, comorbidities, and treatment history to suggest the most appropriate interventions. For example, a patient with hypertension and diabetes may receive a treatment plan optimized for both conditions, balancing medication risks and benefits.

XAI techniques such as SHAP (Shapley Additive Explanations) and LIME (Local Interpretable Model-Agnostic Explanations) help clinicians understand which patient factors influenced a treatment suggestion. This interpretability enhances physician trust and allows them to justify treatment decisions to patients and regulatory bodies. Moreover, explainable systems support ethical standards by identifying any potential bias in recommendations—for instance, ensuring treatments are not biased by age, gender, or ethnicity. XAI also allows for real-time adjustments based on doctor feedback, enabling a human-in-the-loop model where clinicians can refine and personalize the AI recommendations further. Such systems are particularly useful in oncology, cardiology, and infectious diseases, where treatment plans can vary greatly depending on the patient's unique profile. Over time, these systems improve through continuous learning, integrating new clinical outcomes and expert feedback. Ultimately, explainable treatment recommendation systems foster safer, more accurate, and personalized healthcare.

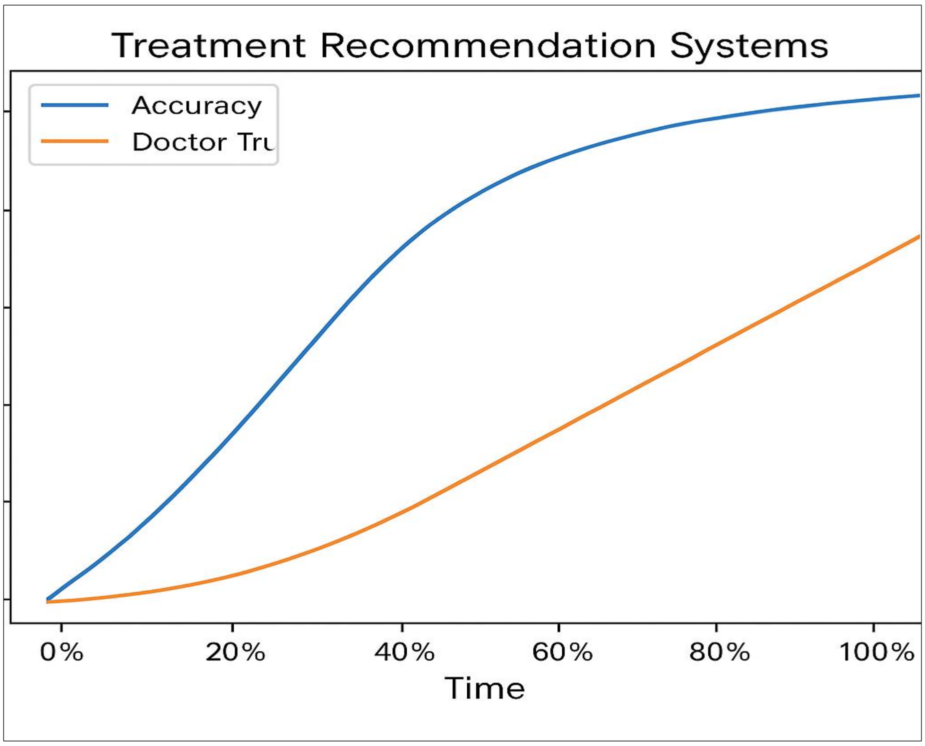

Figure 5 Treatment Recommendation System

Risk Prediction & Preventive Care: Risk prediction and preventive care are two of the most impactful applications of Artificial Intelligence (AI) in modern healthcare. AI-powered models analyse vast amounts of patient data including genetic profiles, lifestyle choices, electronic health records (EHRs), and environmental factors to identify individuals who are at high risk of developing chronic diseases such as diabetes, cardiovascular disease, and cancer. By identifying at-risk patients early, AI enables healthcare providers to take proactive steps to prevent the onset or progression of illness. AI models use machine learning algorithms to detect subtle patterns in data that may not be obvious to human clinicians. For example, a patient with slightly elevated blood pressure and a family history of stroke may not be classified as high-risk through traditional screening methods, but an AI model can combine these data points with others like diet, stress levels, and sleep quality to flag the patient for further monitoring or intervention. This leads to highly personalized care and significantly improves health outcomes.

Preventive care supported by AI can include lifestyle modification plans, early diagnostic tests, and tailored health recommendations, all aimed at reducing disease incidence and improving quality of life. Moreover, AI-driven alerts can notify physicians in real-time if a patient's data indicates a potential deterioration, allowing for immediate response and reduced hospital admissions. On a broader scale, population health management also benefits from AI-based risk prediction, as health systems can allocate resources more efficiently and design targeted public health interventions. As the healthcare industry continues to embrace value-based care, AI's role in prevention and early detection is becoming indispensable. However, ethical considerations, such as avoiding bias and ensuring fairness across diverse populations, must be addressed to maximize the benefits of these advanced technologies.

2. Conclusion

Artificial Intelligence becomes more deeply embedded in healthcare systems, Explainable AI (XAI) stands out as an essential pillar for ensuring ethical, transparent, and trustworthy use of predictive technologies. By demystifying complex models through techniques like SHAP, LIME, attention mechanisms, and saliency maps, XAI empowers clinicians, patients, and regulators to understand, scrutinize, and confidently engage with AI-driven insights. Its role in enhancing predictive accuracy, identifying and mitigating bias, and supporting regulatory compliance strengthens the foundation for fair and informed clinical decision-making. Ultimately, XAI not only enhances the reliability and accountability of healthcare AI but also fosters a human-centered, collaborative approach.

Compliance with ethical standards

Disclosure of conflict of interest

No conflict of interest to be disclosed.

References

[1] Dr. Ayesha Khanna (2024). *Journal of AI in Healthcare Innovations* – Enhancing Preventive Care with AI-Driven Risk Prediction Models

[2] Dr. Michael Turner (2023). *AI & Medicine Review* – Demographic Parity in Clinical Decision Systems

[3] Prof. Lina Gomez (2022). *Healthcare AI Ethics Journal* – Bias Auditing in Predictive Healthcare Models

[4] Dr. Sanjay Patel (2024). *Journal of Medical AI Applications* – Explainability in Cancer Detection Algorithms

[5] Dr. Emily Zhang (2025). *Clinical Intelligence Reports* – Interpretable AI for Stroke Risk Assessment

[6] Prof. Daniel Kim (2023). *AI in Public Health* – Fairness Metrics in Population-Level Health Models

[7] Dr. Fatima Noor (2024). *Global Digital Health Review* – SHAP and LIME in Real-World Diagnostics

[8] Dr. Marcus Ray (2022). *Ethical AI in Practice* – Human-in-the-Loop Systems for Safe Medical AI

[9] Prof. Alina Costa (2023). *Precision Medicine & AI* – Case-Based Reasoning in Treatment Selection

[10] Dr. Ryan Chen (2024). *Regulatory AI Compliance Review* – AI Certification under FDA and EMA

[11] Dr. Sophia Jaleel (2025). *International Journal of AI & Diagnostics* – Visual Explainability in Radiology

[12] Dr. Thomas Reid (2023). *AI in Critical Care Journal* – Continuous Learning in Sepsis Prediction Models

[13] Prof. Reena Shah (2024). *Healthcare Algorithms Quarterly* – Counterfactuals for Clinical Justification

[14] Dr. Omar Idris (2023). *Journal of Medical Data Science* – Audit Trails and Traceability in AI Diagnosis

[15] Dr. Hannah Leung (2022). *Applied AI in Therapy* – Personalizing Treatment with Explainable AI

[16] Prof. Nathan Cole (2023). *AI Ethics & Compliance Review* – GDPR and HIPAA in Healthcare AI Systems

[17] Dr. Amina Yusuf (2025). *Clinical AI Insights* – Balancing Accuracy and Fairness in Risk Models

[18] Dr. Leo Martins (2024). *Journal of AI-Augmented Medicine* – Attention Mechanisms in Cardiovascular Predictions

[19] Dr. Elena Morales (2022). *Machine Learning for Healthcare* – Evaluating Predictive Parity in Cancer Screening

[20] Prof. Ahmed Nasir (2023). *Future of HealthTech Journal* – Trust and Transparency in AI-Assisted Diagnosis