Volume 4, Issue 6, June 2026 ISSN (E): 2938-3765

APPLICATIONS OF ARTIFICIAL INTELLIGENCE IN HUMAN GENOMIC RESEARCH AND PRECISION MEDICINE

D. A. Olimova
Sh. S. Hamrakulova
O. B. Alimova
Assistant of the Institute of Medicine and Bioengineering, Zarmed University, Samarkand Assistant of Samarkand State Medical University, Samarkand
Student of Samarkand State Medical University, Samarkand

Abstract

Background: The volume of human genomic data generated by whole-genome sequencing, genome-wide association studies (GWAS), and multi-omics platforms now far exceeds the interpretive capacity of conventional bioinformatic pipelines. Artificial intelligence (AI) and machine learning (ML) methods have emerged as essential tools for extracting clinically and biologically meaningful signals from these datasets.

Main Findings: This review critically examines four major application domains: (1) pathogenic variant interpretation and functional annotation, where models such as AlphaFold2 and ACMG/AMP-guided ML classifiers have substantially improved resolution of variants of uncertain significance (VUS); (2) disease prediction and polygenic risk scoring, particularly for cancer and complex multifactorial conditions; (3) drug discovery and pharmacogenomics, including AI-driven target identification and personalised therapeutic selection; and (4) deep learning architectures in genomics, encompassing transformer-based sequence models, genomic foundation models such as the Nucleotide Transformer and DNABERT-2, and multi-omics integration frameworks.

Significance: Despite compelling advances, critical barriers impede clinical translation: algorithmic bias arising from the historical underrepresentation of non-European ancestries in training datasets, limited model explainability, data privacy constraints, and fragmented regulatory frameworks. Federated learning and explainable AI (XAI) strategies are discussed as promising mitigation approaches. This review argues that realising the full potential of AI in precision genomic medicine requires not only technical innovation but also deliberate efforts towards equitable data collection and transparent model development.

Keywords: Artificial intelligence; genomics; precision medicine; deep learning; polygenic risk scores; variant interpretation; pharmacogenomics; foundation models.

Introduction

The past two decades have witnessed an unprecedented expansion in the generation of human genomic data. Following the completion of the Human Genome Project in 2003, advances in next-generation sequencing have reduced the cost of whole-genome sequencing from approximately three

Web of Medicine: Journal of Medicine, Practice and Nursing
webofjournals.com/index.php/5

billion US dollars to under one thousand dollars per genome, enabling population-scale studies that were previously inconceivable [1]. Biobanks containing hundreds of thousands of sequenced individuals, coupled with multi-omics platforms measuring transcriptomic, epigenomic, proteomic, and metabolomic states, now produce data volumes measured in petabytes annually [2].

This abundance creates a formidable interpretive bottleneck. The human genome harbours approximately three million single nucleotide variants (SNVs) per individual relative to the reference sequence, the majority of which remain functionally uncharacterised [3]. Genome-wide association studies have identified tens of thousands of loci associated with complex diseases, yet most variants lie within non-coding regulatory regions whose mechanistic roles are poorly understood. Conventional rule-based clinical classification frameworks, such as the ACMG/AMP guidelines for variant pathogenicity assessment, are indispensable but insufficient at scale: they require manual curation, generate large proportions of variants of uncertain significance (VUS), and cannot efficiently integrate evidence from diverse functional, population, and structural data sources [4].

Artificial intelligence, encompassing machine learning, deep learning, and natural language processing, has emerged as the principal computational paradigm to address this interpretive gap. Unlike hypothesis-driven statistical models, AI methods can detect complex nonlinear patterns across high-dimensional datasets without requiring prior specification of the relevant features. The application of deep neural networks to genomics was substantially accelerated by landmark achievements in adjacent fields, most notably by the demonstration that transformer-based architectures trained on large corpora could generalise remarkably well to novel biological prediction tasks [5,6].

In the clinical context, precision medicine aspires to tailor prevention, diagnosis, and therapy to an individual's molecular profile. AI is increasingly positioned as the computational engine of precision genomic medicine, translating large-scale genomic data into actionable clinical insights. This review examines the principal application domains, critically evaluates the evidence base, and identifies the ethical, technical, and implementation challenges that must be resolved before AI-driven genomics can be equitably deployed in healthcare.

Applications of AI in Human Genomics

1. Variant Interpretation

Accurate discrimination between pathogenic and benign genetic variants is the cornerstone of genomic diagnostics, and it is an area where AI has generated some of its most consequential contributions. The traditional classification paradigm relies on the ACMG/AMP five-tier framework, which aggregates evidence from population databases, computational predictions, functional studies, and clinical observations. While this framework is rigorously evidence-based, its consistent application demands expert curation and its outputs remain heavily dependent on the quality and coverage of the underlying evidence databases [4].

Machine learning approaches have been developed to augment or partially automate this process. Nicora and colleagues [4] demonstrated that penalised logistic regression trained on ACMG/AMP evidence features could resolve a significantly larger proportion of VUS cases compared to purely rule-based methods, yielding a probabilistic pathogenicity score that supports variant prioritisation in clinical exome sequencing workflows. More recently, the DiagAI system, which integrates ML-based

Web of Medicine: Journal of Medicine, Practice and Nursing
webofjournals.com/index.php/5

pathogenicity prediction with Human Phenotype Ontology (HPO) phenotypic matching, achieved shortlisting of diagnostic variants in 94.9% of cases in a retrospective cohort of 966 adult nephropathy patients, outperforming HPO-naive approaches by approximately four percentage points [5].

The structural revolution catalysed by AlphaFold2 [6] has further transformed variant interpretation. By providing atomic-resolution protein structure predictions for the majority of human proteins, AlphaFold2 enables in silico assessment of how missense variants perturb protein folding and stability, an approach that was previously limited to proteins with experimentally resolved structures. The tool's transformer-based architecture, trained on the Protein Data Bank, achieved median TM-score performance exceeding 0.90 across CASP14 targets, a level of accuracy previously anticipated to require a further decade of development [6,7]. Protein language models, including ESM-2, further exploit evolutionary information encoded in multiple sequence alignments to generate variant effect scores that generalise across protein families without requiring per-protein fine-tuning [3].

Despite these advances, a critical challenge remains: most ML tools are validated on pathogenic variants from well-characterised Mendelian disease genes in populations of predominantly European ancestry. Performance on rare variants in understudied genes or diverse populations is poorly characterised, and the VarChat system [8], a generative AI assistant for variant interpretation, illustrates both the potential and the limitations of large language model approaches when the underlying evidence base is sparse or non-standardised.

2. Disease Prediction and Risk Assessment

Polygenic risk scores (PRS), which aggregate the effect of thousands to millions of common variants identified through GWAS into a single individual-level risk metric, represent one of the most direct translations of population genomics into clinical prediction [9]. AI has contributed to PRS development in two distinct ways: improving the statistical methods used to derive scores from GWAS summary statistics, and integrating PRS with non-genomic clinical and demographic data to enhance predictive performance.

Deep neural networks have demonstrated superiority over established statistical PRS algorithms. Badre and colleagues [10] showed that a deep neural network outperformed BLUP, BayesA, and LDpred for breast cancer PRS estimation, with AUC values of 67.4% compared to 64.5% for the next-best method, and observed that DNN-derived PRS followed a bimodal distribution in affected individuals, suggesting a latent biological subgrouping not captured by linear models. At the population level, Zhang and colleagues [11] introduced a multiancestry PRS methodology published in Nature Genetics in 2023 that improved prediction accuracy in individuals of African ancestry, explicitly acknowledging that methodological advances alone cannot overcome the performance gap caused by underrepresentation of diverse populations in training GWAS.

In oncology, AI-driven PRS and multi-modal risk models have moved from research tools toward clinical feasibility. Sud and colleagues [9] reviewed the evidence for PRS in cancer risk stratification and concluded that while polygenic architectures for breast, prostate, and colorectal cancers are sufficiently well-powered for stratified screening, integration with family history, environmental exposures, and imaging biomarkers via ML ensemble approaches is necessary for clinically meaningful discrimination. The Cancer Genome Atlas (TCGA) and similar resources have enabled the training of deep learning models on somatic genomic profiles to predict tumour subtypes,

Web of Medicine: Journal of Medicine, Practice and Nursing
webofjournals.com/index.php/5

prognosis, and treatment sensitivity, though external validation across healthcare systems with different sequencing platforms remains inconsistent [12].

For rare disease diagnosis, AI has addressed the classical problem of phenotypic and molecular heterogeneity. The combination of ML variant prioritisation with automated HPO phenotyping enables exome and genome interpretation pipelines to generate candidate diagnoses ranked by phenotypic fit, dramatically reducing the median time to diagnosis in paediatric rare disease programmes that previously required months of manual review [5].

3. Drug Discovery and Pharmacogenomics

The application of AI to drug discovery encompasses target identification, lead compound optimisation, drug repurposing, and personalised pharmacogenomics. The economics of drug development provide a compelling motivation: median development costs reached approximately 2.4 billion US dollars between 2013 and 2022, with development timelines extending by one to two years over the same period [13]. AI-driven approaches that reduce attrition in early-stage discovery have clear potential to compress both cost and timeline.

In target identification, AI models integrate multi-omics data, biomedical knowledge graphs, and protein structure information to prioritise disease-relevant genes and assess their druggability. Wenteler and colleagues [14] reviewed recent advances in AI-driven target discovery and described frameworks that combine CRISPR functional genomic screens, transcriptomic perturbation data, and network-based gene prioritisation models, emphasising that integrating multimodal inputs is essential for capturing the genetic and environmental complexity of common diseases. Critically, they note that AI tools have enabled or supported the identification of several drug targets that have now entered clinical trials, providing proof-of-concept for the translational value of computational target discovery.

Pharmacogenomics, the study of how genetic variation influences drug response, has been transformed by ML approaches that extend beyond single-gene CYP450 genotyping to whole-genome predictive models. Taliaz and colleagues [15] developed an ML algorithm integrating genetic, clinical, and demographic features from the STAR*D cohort to predict individual response to a panel of antidepressants, achieving a balanced accuracy of 72.3% in validation, substantially above both standard clinical practice and single-gene pharmacogenomic tests. This multimodal approach illustrates a broader principle: the predictive value of genomic data is maximised when integrated with phenotypic and clinical covariates rather than analysed in isolation. While the antidepressant domain has served as a proving ground, analogous frameworks are being developed for oncology, cardiovascular pharmacotherapy, and immunosuppression [13].

4. Deep Learning in Genomics

Deep learning has fundamentally altered the scale and precision of sequence-level genomic analyses. The Enformer model, developed by Avsec and colleagues [16] at DeepMind and published in Nature Methods in 2021, demonstrated that transformer-based architectures integrating up to 100 kilobases of DNA context could predict gene expression with significantly greater accuracy than prior convolutional models, and could learn enhancer-promoter interaction topologies directly from sequence without experimental contact data. This has important downstream implications: Enformer-

Web of Medicine: Journal of Medicine, Practice and Nursing
webofjournals.com/index.php/5

generated variant effect predictions for expression quantitative trait loci (eQTLs) offer a computationally scalable alternative to experimental massively parallel reporter assays for functional annotation of non-coding GWAS variants [16].

The development of genomic foundation models, trained on large collections of unannotated genomic sequence across multiple species, represents a paradigm shift analogous to the emergence of general-purpose language models in natural language processing. DNABERT, first introduced by Ji and colleagues [17] and subsequently extended to DNABERT-2 [17] with byte-pair encoding tokenisation and ALiBi positional embeddings, demonstrated that bidirectional transformer models pre-trained on human reference genome sequence could be fine-tuned to achieve competitive performance across a range of downstream genomic tasks. The Nucleotide Transformer series, ranging from 500 million to 2.5 billion parameters and trained on 3,202 diverse human genomes as well as 850 multi-species genomes, achieved the strongest overall performance across 18 benchmarked genomic prediction tasks [18].

The integration of multi-omics data types presents a complementary challenge: single-cell technologies can simultaneously profile chromatin accessibility, transcription, and surface protein expression within the same cell, but joint analysis of these heterogeneous modalities requires dimensionality reduction methods that can capture both shared and modality-specific variance. Deep generative models, particularly variational autoencoders (VAEs) and their multimodal extensions, have emerged as dominant approaches for this integration [19,20]. The multiDGD framework, published in Nature Communications in 2024, provided a probabilistic generative model for paired single-cell multi-omics data that achieved superior reconstruction performance while modelling batch effects as explicit covariates, enabling computational linkage of distal regulatory elements to target promoters [20].

Ethical, Legal, and Social Challenges

The deployment of AI in clinical genomics is subject to a constellation of ethical, legal, and social challenges that are distinct from those governing AI in other medical domains, owing to the hereditary nature of genomic information and the population-level inferences that genomic models make.

Algorithmic bias arising from training data composition is perhaps the most thoroughly documented challenge. Despite growth in the total number of GWAS participants, the proportion of non-European individuals in GWAS catalogues has stagnated or declined in relative terms: an analysis by Fatumo and colleagues [21] estimated that approximately 72% of GWAS participants were recruited from just three countries (the United States, the United Kingdom, and Iceland), and the proportion of non-European participants fell from 19% in 2016 to 14% in 2021, driven partly by the continued heavy use of the UK Biobank White British cohort. AI models trained on such data systematically underperform for individuals of non-European ancestry, and this performance gap translates directly into disparities in diagnostic yield, PRS accuracy, and pharmacogenomic prediction [11,21]. Zhang and colleagues explicitly demonstrated that no single PRS method currently performs equitably across ancestries, and that improved methods alone cannot substitute for larger, more representative training datasets [11].

Data privacy poses a second major challenge. Genomic data are inherently identifiable: as few as 30 to 80 SNPs can re-identify individuals from anonymised datasets [22]. The aggregation of genomic

Web of Medicine: Journal of Medicine, Practice and Nursing
webofjournals.com/index.php/5

data with electronic health records, wearable sensor data, and environmental exposures, while analytically powerful, creates regulatory complexity under frameworks such as the General Data Protection Regulation (GDPR) in Europe and HIPAA in the United States. These regulations may conflict with the large-scale cross-institutional data sharing that AI model training requires.

The opacity of deep learning models presents a specific concern in clinical genomics: when a model assigns a high pathogenicity score to a novel variant or generates a high-risk PRS, clinicians require not only the prediction but an interpretable rationale. Abe and colleagues [23] demonstrated that XAI approaches using knowledge graphs can simultaneously achieve high accuracy and provide physician-understandable explanations for variant pathogenicity estimates. More broadly, Talukder and colleagues [24] reviewed XAI approaches for omics data and concluded that while feature relevance methods such as SHAP and gradient-based attribution are widely used, they frequently do not faithfully represent the computational mechanism of the underlying model, and many clinical and regulatory questions regarding their reliability remain unresolved.

Future Perspectives

Federated learning, in which model parameters rather than raw data are shared between institutions, offers a principled solution to the tension between data privacy and the need for large, diverse training datasets. Federated approaches have been applied to rare disease genomics [22] and electronic health record-linked genomic analyses, and they reduce, though do not eliminate, privacy risks: gradient updates can be inverted to partially reconstruct training data without additional protections such as differential privacy. The co-design of federated genomic studies with appropriate cryptographic guarantees represents a research priority for the next decade.

The maturation of genomic foundation models is anticipated to yield substantial downstream benefits for both basic genomic research and clinical variant interpretation. Current models operate primarily on DNA sequence; future architectures are expected to jointly model DNA sequence, chromatin state, three-dimensional genome organisation, and RNA secondary structure within a unified representation framework, enabling more comprehensive functional annotation of the non-coding genome. The success of AlphaFold2 in predicting protein structures [6] and its successors, including AlphaFold3, which extends predictions to nucleic acids, small molecules, and multi-molecular complexes, suggests that structural approaches will increasingly inform variant effect prediction for both coding and regulatory variants.

In clinical practice, the principal bottleneck is no longer algorithmic performance but implementation: the integration of AI-generated variant interpretations, PRS outputs, and pharmacogenomic predictions into electronic health record workflows, the training of clinicians to appropriately contextualise probabilistic AI outputs, and the development of reimbursement structures that support AI-augmented genomic medicine. Regulatory frameworks for AI in genomic diagnostics remain fragmented across jurisdictions, and prospective clinical evidence for patient outcome benefit from AI-assisted genomic interpretation remains limited compared to the abundance of retrospective validation studies. Future research should prioritise randomised trial designs that measure clinical utility rather than solely analytical performance.

Web of Medicine: Journal of Medicine, Practice and Nursing
webofjournals.com/index.php/5

Web of Medicine: Journal of Medicine, Practice and Nursing
webofjournals.com/index.php/5

Conclusion

Artificial intelligence has demonstrated transformative potential across the principal domains of human genomic research: variant interpretation, disease risk prediction, drug discovery, and multi-omics integration. Landmark developments including AlphaFold2, genomic foundation models, deep learning-enhanced PRS, and multi-modal pharmacogenomic prediction algorithms have each extended the interpretive frontier of genomic medicine in ways that were not achievable by prior statistical approaches.

However, the trajectory from research performance to equitable clinical impact is neither linear nor guaranteed. The persistent Eurocentric composition of genomic training datasets, the opacity of high-performing models, and the absence of prospective clinical evidence collectively impose limits on how confidently AI-derived genomic insights can be applied in diverse patient populations. The tools that work best for well-represented populations may generate unreliable or misleading predictions for the global majority of patients.

Advancing AI in precision genomic medicine therefore demands a dual commitment: continued technical innovation in model architecture, interpretability, and privacy-preserving training, alongside deliberate investment in the inclusive genomic datasets, regulatory clarity, and clinical implementation infrastructure that are prerequisites for equitable deployment. When these conditions are met, AI-enabled genomics holds genuine promise to accelerate diagnosis for rare disease patients, stratify cancer screening programmes by individual genomic risk, and guide pharmacological prescribing with greater precision than current standard-of-care approaches permit.

References

1. Shendure J, Balasubramanian S, Church GM, et al. DNA sequencing at 40: past, present and future. Nature. 2017;550(7676):345-353. doi:10.1038/nature24286. Available from: https://doi.org/10.1038/nature24286

2. Bycroft C, Freeman C, Petkova D, et al. The UK Biobank resource with deep phenotyping and genomic data. Nature. 2018;562(7726):203-209. doi:10.1038/s41586-018-0579-z. Available from: https://doi.org/10.1038/s41586-018-0579-z

3. Lin Z, Akin H, Rao R, et al. Evolutionary-scale prediction of atomic-level protein structure with a language model. Science. 2023;379(6637):1123-1130. doi:10.1126/science.ade2574. Available from: https://doi.org/10.1126/science.ade2574

4. Nicora G, Zucca S, Limongelli I, Bellazzi R, Magni P. A machine learning approach based on ACMG/AMP guidelines for genomic variant classification and prioritization. Sci Rep. 2022;12(1):2517. doi:10.1038/s41598-022-06547-3. Available from: https://doi.org/10.1038/s41598-022-06547-3

5. Stenton SL, Barbalat G, Barth M, et al. DiagAI: a machine learning system for clinical exome variant prioritisation using molecular pathogenicity and phenotypic features. medRxiv. 2025. doi:10.1101/2025.02.04.25321641. Available from: https://doi.org/10.1101/2025.02.04.25321641

6. Jumper J, Evans R, Pritzel A, et al. Highly accurate protein structure prediction with AlphaFold. Nature. 2021;596(7873):583-589. doi:10.1038/s41586-021-03819-2. Available from: https://doi.org/10.1038/s41586-021-03819-2

7. Tunyasuvunakool K, Adler J, Wu Z, et al. Highly accurate protein structure prediction for the human proteome. Nature. 2021;596(7873):590-596. doi:10.1038/s41586-021-03828-1. Available from: https://doi.org/10.1038/s41586-021-03828-1
8. Zucca S, Mottaz A, Pasche E, Gobeill J, Teodoro D, Ruch P. VarChat: the generative AI assistant for the interpretation of human genomic variations. Bioinformatics. 2024;40(4):btae183. doi:10.1093/bioinformatics/btae183. Available from: https://doi.org/10.1093/bioinformatics/btae183
9. Sud A, Turnbull C, Houlston R. Will polygenic risk scores for cancer ever be clinically useful? NPJ Precis Oncol. 2021;5(1):40. doi:10.1038/s41698-021-00176-1. Available from: https://doi.org/10.1038/s41698-021-00176-1
10. Badre A, Zhang L, Muchero W, Reynolds JC, Pan C. Deep neural network improves the estimation of polygenic risk scores for breast cancer. J Hum Genet. 2021;66(4):359-369. doi:10.1038/s10038-020-00832-7. Available from: https://doi.org/10.1038/s10038-020-00832-7
11. Zhang H, Zhan J, Jin J, et al. A new method for multiancestry polygenic prediction improves performance across diverse populations. Nat Genet. 2023;55(10):1757-1768. doi:10.1038/s41588-023-01501-z. Available from: https://doi.org/10.1038/s41588-023-01501-z
12. Choi J, Kim JS, Sung HJ, et al. Polygenic risk scores associated with tumor immune infiltration in common cancers. Cancers (Basel). 2022;14(22):5571. doi:10.3390/cancers14225571. Available from: https://doi.org/10.3390/cancers14225571
13. Wenteler A, Cabrera CP, Wei W, Neduva V, Barnes MR. AI approaches for the discovery and validation of drug targets. Camb Prism Precis Med. 2024;2:e7. doi:10.1017/pcm.2024.4. Available from: https://doi.org/10.1017/pcm.2024.4
14. Wenteler A, Cabrera CP, Wei W, Neduva V, Barnes MR. AI approaches for the discovery and validation of drug targets. Camb Prism Precis Med. 2024;2:e7. doi:10.1017/pcm.2024.4. Available from: https://doi.org/10.1017/pcm.2024.4
15. Taliaz D, Spinrad A, Lombroso A, et al. Optimizing prediction of response to antidepressant medications using machine learning and integrated genetic, clinical, and demographic data. Transl Psychiatry. 2021;11(1):381. doi:10.1038/s41398-021-01488-3. Available from: https://doi.org/10.1038/s41398-021-01488-3
16. Avsec Z, Agarwal V, Visentin D, et al. Effective gene expression prediction from sequence by integrating long-range interactions. Nat Methods. 2021;18(10):1196-1203. doi:10.1038/s41592-021-01252-x. Available from: https://doi.org/10.1038/s41592-021-01252-x
17. Zhou Z, Ji Y, Li W, Dutta P, Davuluri R, Liu H. DNABERT-2: efficient foundation model and benchmark for multi-species genome. arXiv. 2023:2306.15006. Available from: https://arxiv.org/abs/2306.15006
18. Dalla-Torre H, Gonzalez L, Mendoza-Revilla J, et al. The Nucleotide Transformer: building and evaluating robust foundation models for human genomics. Nat Methods. 2024;21(12):2229-2243. doi:10.1038/s41592-024-02523-z. Available from: https://doi.org/10.1038/s41592-024-02523-z

Web of Medicine: Journal of Medicine, Practice and Nursing
webofjournals.com/index.php/5