

Generative Health Data Science for Data-to-Dose Precision Pharmacology: A Critical Methodological Review and Translational Framework

Alisha Singh¹, Md Faiyaz^{2*}, Ujjwal khatua³, Anjali Mishra⁴, Ravi Chandra Kumar⁵, Rupesh kumar⁶,
Md Gulzar Ahmad⁷, Lalit Kumar⁸

¹Assistant Professor, School of Pharmacy, Sandip University, Sijoul, Madhubani, Bihar, India, ORCID: 0009-0001-2081-9676,

E-mail: alishabaghel241@gmail.com

²PG Scholar, Department of Pharmacology, Faculty of Pharmacy, Uttar Pradesh University of Medical Sciences, Saifai, Etawah, Uttar Pradesh, India,

ORCID: 0009-0001-7948-5068, E-mail: drx.mdfaiyaz@gmail.com

³PG Scholar, Department of Pharmacy, Central University of South Bihar, India, E-mail: khatuaujjwal435@gmail.com

⁴Assistant Professor, School of Pharmacy, Sandip University, Sijoul, Madhubani, Bihar, India, E-mail: 14.misanjali@gmail.com

⁵Assistant Professor, SNS College of Pharmacy, Motihari, Bihar, India, E-mail: rkumar1734@gmail.com

⁶PG Scholar, Department of Pharmacy, Central University of South Bihar, India, E-mail: rupesh48793@gmail.com

⁷Pharmacist at Apollo Pharmacies Limited, New Delhi, E-mail: mdgulzarahmad89@gmail.com

⁸Assistant Professor, Shri R.L.T. Institute of Pharmaceutical Science and Technology, Ekdil, Etawah, Uttar Pradesh, India, E-mail: lalitlalts7@gmail.com

Corresponding author: Md Faiyaz, drx.mdfaiyaz@gmail.com

Abstract:

Precision pharmacology aims to select the most appropriate medicine, dose and monitoring strategy for an individual patient. Its translation is limited by fragmented records, incomplete biological characterization, time-varying physiology and the difficulty of converting high-dimensional evidence into a safe action. Generative health data science extends conventional prediction by learning structured representations of molecules, cells, patients and clinical trajectories and by producing reviewable candidates such as molecular structures, treatment options, synthetic records, counterfactual courses and dose scenarios. This critical review evaluates large language models, multimodal and biological foundation models, graph and diffusion architectures, synthetic-patient generators, digital twins and bounded agentic workflows across target discovery, repurposing, pharmacogenomics, treatment selection, individualised dosing and pharmacovigilance. Evidence is strongest for representation learning, structure-enabled discovery and selected pharmacogenomic pathways; autonomous patient-level treatment and dosing remain investigational. Recent advances include explainable drug-repurposing foundation models, multimodal clinical foundation models and a randomised phase 2a trial of a generative-AI-discovered TNiK inhibitor. We propose a Data-to-Dose framework linking intended-use definition, data provenance, longitudinal patient-state modelling, constrained generation, mechanistic and causal checks, uncertainty-aware ranking, professional review and lifecycle monitoring. Generation must remain distinct from clinical authorisation, and synthetic-data fidelity must not be equated with privacy, fairness or external validity. Generative systems may become an enabling layer for precision pharmacology only when their creative capacity is bounded by pharmacology, evidence and accountable clinical judgment.


Keywords: digital twins; generative artificial intelligence; health data science; multimodal foundation models; pharmacogenomics; precision pharmacology; synthetic patients.

How to cite this article: Saxena S, Jain SK. Development and Validation of a Stability-Indicating RP-HPLC Method for Simultaneous Estimation of Dapagliflozin, Metformin Hydrochloride, and Vildagliptin in a Fixed-Dose Combination Tablet. *Int J Drug Deliv Technol*. 2026;16(86s): 1591-1601; DOI: 10.25258/ijddt.16.86s.140

Abbreviations: ADMET, absorption, distribution, metabolism, excretion and toxicity; AI, artificial intelligence; CDS, clinical decision support; EHR, electronic health record; EU, European Union; FDA, US Food and Drug Administration; FM, foundation model; GenAI, generative artificial intelligence; LLM, large language model; MFM, multimodal foundation model; PGx, pharmacogenomics; PK/PD, pharmacokinetics/pharmacodynamics; RAG, retrieval-augmented generation; RWE, real-world evidence; WHO, World Health organisation

Graphical abstract:

Generative health data science for precision pharmacology



1. Introduction

Precision pharmacology integrates patient, disease and drug information to select an intervention and individualise exposure while preserving a favourable benefit-risk balance [1-3]. Its evidence base spans medication history, laboratory trajectories, organ function and pharmacogenomics, while multimodal AI extends representation to imaging and physiological signals [2,4]. Most established clinical AI systems classify observations or estimate predefined outcomes, whereas generative models can propose molecular structures or learn reusable representations across heterogeneous biomedical data [5-7]. Patient-linked simulation is an emerging objective of digital-twin research, but any output that can influence treatment requires intended-use definition and workflow evaluation [8,9]. Previous reviews described broad healthcare, pharmaceutical and AI-enabled drug-discovery applications; the present article focuses on the narrower pathway from multimodal patient data to pharmacological action [10,11].

Drug response reflects inherited determinants as well as dynamic factors such as inflammation, renal or hepatic function, co-medication, disease stage and adherence [2,3]. Actionable PGx pathways exist, but genotype alone does not

capture the complete treatment phenotype [2,3,12]. Multimodal modelling can integrate clinical and molecular layers, although retrospective performance does not establish safe medication use [4,5,9]. Medication-focused reviews identify hallucination, omission and limited prospective evidence as unresolved concerns for GenAI and LLM systems [13-15]. Computationally generated molecules likewise require chemical-validity checks, synthesis and experimental confirmation before pharmacological claims are justified [7,11,16]. Synthetic records require task-specific evaluation and must not substitute for validation on real patients [17,18]. Generative systems should therefore be treated as proposal engines within an accountable workflow rather than autonomous prescribers [9,14,15].

This review defines generative health data science for precision pharmacology, evaluates applications from target discovery to post-prescription monitoring, summarizes recent advances, introduces a Data-to-Dose framework and specifies validation and governance requirements. As illustrated in **Figure 1**, the framework links input modalities and generative representations with pharmacological outputs and cross-cutting safeguards.




Figure 1. Generative health data science landscape for precision pharmacology.

Note: Patient records, multi-omics, imaging and physiological signals can be represented by language and multimodal models [4,6,19]. Chemical and structural data can be represented by latent-space, graph, diffusion and sequence models [7,20,21]. The lower band presents provenance, uncertainty, privacy, equity, human authorisation and controlled updating as cross-cutting design requirements [9,14,22].

2. Review approach and evidence selection

This article was developed as a targeted critical methodological review rather than a systematic review or meta-analysis. PubMed/MEDLINE, primary publisher platforms and official regulatory repositories were searched for English-language literature available through 21 July 2026. Search terms combined generative artificial intelligence, foundation models, multimodal health data, single-cell models, drug discovery, pharmacogenomics, precision dosing, synthetic patients, digital twins, clinical validation, bias, federated learning and lifecycle regulation. Eligible sources included original methodological or clinical studies, prospective or randomised evaluations, authoritative reporting standards and official governance documents. Opinion pieces without methodological or pharmacological relevance, duplicate reports, non-English records and sources whose claims could not be verified against an original or official record were excluded. Reference chaining was used for foundational studies. Bibliographic metadata and DOI identifiers were checked against publisher pages, PubMed, official records or the supplied published PDFs.

Evidence was interpreted across four maturity levels: enabling performance, translational confirmation, clinical utility and lifecycle evidence. Because the included technologies, intended uses and outcomes were heterogeneous, quantitative pooling was inappropriate. Claims were therefore calibrated to the maturity and directness of the supporting evidence. The review does not claim exhaustive systematic coverage; instead, it provides a transparent, purpose-driven

synthesis centred on the proposed Data-to-Dose pathway.

3. Foundations of generative precision pharmacology

3.1 From prediction to decision-focused generation

Predictive models estimate outcomes conditional on observed inputs, whereas generative models learn distributions or structured representations from which new instances can be sampled or proposed [6,7]. Pharmacological outputs may include molecular graphs, biomolecular complexes or designed proteins [20,21,25]. Other outputs include generated single-cell representations, clinical text or patient-linked simulations [8,26,27]. Generation is not equivalent to a clinical decision because treatment selection requires a defined comparator, benefit-risk objective, contraindication checks and uncertainty thresholds [9,24].

3.2 Multimodal patient-state representation

Structured EHR fields capture diagnoses, prescriptions and laboratory results; clinical notes preserve context; imaging and pathology capture spatial phenotype; omics describe inherited and dynamic biology; and wearables provide longitudinal physiology [4,6,28]. Generalist and specialty foundation models can reuse representations across tasks, as illustrated by medical foundation-model frameworks, PanDerm and SleepFM [6,19,29]. Retrospective benchmark performance does not establish medication safety because temporal ambiguity, incomplete reconciliation, local terminology and missing data can distort the inferred patient state [9,13,24]. Retrieval grounding can improve traceability, but retrieved evidence can still be incomplete, outdated or misapplied [13-15].

3.3 Biological and chemical foundation models

Geneformer and scGPT demonstrate that large-scale single-cell pretraining can support network prediction, perturbation analysis and multi-omics integration [26,30]. Latent-space and diffusion approaches can generate molecules or proteins under specified objectives [7,21,25]. AlphaFold 3

predicts structures of biomolecular complexes rather than establishing drug efficacy [20]. Generated candidates remain hypotheses until synthesis or expression, assay confirmation, selectivity, ADMET and reproducibility are

demonstrated [7,16,31]. As summarised in **Table 1**, each health-data modality requires a distinct generative task, precision-pharmacology use and minimum validity assessment.

Table 1. Health-data modalities, generative tasks and pharmacological validity requirements.

Data modality	Generative contribution	Precision-pharmacology use	Minimum validity checks
Clinical records and medication history	Longitudinal representation, summarization and candidate explanation [13,27,28]	Medication reconciliation, evidence synthesis and option review [10,13]	Temporal accuracy, source attribution and reconciliation with actual administration [9,23,24]
Genomics and multi-omics	Variant, cell-state and perturbation representations [2,26,30]	Drug avoidance, stratification, biomarker discovery and dose adjustment [2,3,12]	Validated assays, ancestry coverage, batch control and guideline concordance [2,12]
Imaging and physiology	Cross-modal representation and longitudinal phenotype modelling [4,19,29]	Phenotyping, response assessment and adaptive monitoring [4,19,29]	External-site, device-shift, subgroup and prospective validation [9,19,32]
Chemical and structural data	Molecule, protein or complex generation and structure prediction [7,20,21]	Target discovery, lead design and repurposing hypotheses [11,31,33]	Chemical validity, synthesizability, orthogonal assays and off-target testing [7,16,21]

[13-15].

Note: ADMET, absorption, distribution, metabolism, excretion and toxicity; EHR, electronic health record. A generated value should remain explicitly labelled and should not replace an observed measurement without task-specific validation [9,17,24].

4. Model families and bounded workflows

4.1 Language and retrieval-grounded models

LLMs can encode clinical knowledge and support health-system prediction tasks [27,28]. Medication-focused studies have separately evaluated information extraction, decision support and pharmacovigilance applications [13]. A risk-controlled role is to assemble an inspectable patient state, retrieve current evidence and format a proposal for professional review rather than independently authorize treatment [9,14,15]. Prompting alone is not a safety architecture; pharmacology-focused systems require source traceability and deterministic checks for active ingredient, formulation, unit, dose limits, organ-function thresholds, contraindications and interactions

4.2 Graph, sequence, diffusion and causal

models

Molecules, proteins, pathways and longitudinal patients have graph, sequence or temporal structure that can be represented by specialized models [7,20,30]. Deep learning identified antibacterial candidates, TxGNN supported clinician-centred drug-repurposing prioritisation, and generative discovery produced rentosertib for a randomised phase 2a trial [16,31,33]. These studies demonstrate feasibility in specific settings, not general superiority of generative pipelines over conventional discovery or clinical development [7,31,33]. Counterfactual treatment modelling additionally requires explicit assumptions about treatment assignment, confounding, adherence and time-varying covariates [8,34].

4.3 Bounded agentic systems

A bounded agent could retrieve laboratory

results, query approved knowledge sources, invoke a validated dose model and draft a recommendation for review; this is a proposed workflow rather than an established standard of care [13-15]. Because errors can propagate across tool calls, we propose permission controls, audit logs, independent checking, human confirmation and incident pathways before an agent influences treatment [14,15]. High-risk systems should not place orders or alter doses unless the exact workflow has prospective evidence and explicit institutional authorisation [9,23,24].

5. Applications across the pharmacology continuum

5.1 Target discovery, molecular generation and repurposing

Single-cell foundation models can prioritize disease-associated programs and perturbation responses [26,30]. Protein and molecular generators can propose interaction structures and optimize candidate properties [7,20,21]. Credible discovery workflows require experimental measurement and independent confirmation [7,16,31]. The rentosertib phase 2a study is a translational milestone, but its findings support further testing rather than general claims of clinical superiority [31]. TxGNN can prioritize repurposing candidates and provide biomedical paths, although prospective treatment benefit remains to be established [33,34].

5.2 Patient stratification and drug selection

PGx provides a mature example of data-to-treatment translation because selected gene-drug pairs have validated assays and explicit actions [2,12,35]. A 12-gene panel reduced clinically relevant adverse reactions in a multicountry implementation study, and genotype-guided P2Y12 selection was evaluated in primary percutaneous coronary intervention [12,35]. Generative systems may

explain recommendations or identify missing variables, but they should not replace validated laboratory interpretation or current guidance [2,3,9]. Multimodal treatment selection should estimate comparative benefit from trials, causal analyses or validated models rather than linguistic similarity [4,24,34].

5.3 individualised dose and schedule

Dose individualization is high risk because numerical outputs can be directly actionable and must account for formulation, indication, body size, organ function, genotype, interacting medicines and therapeutic monitoring [2,3,36]. The GIFT trial illustrates that genotype-guided warfarin dosing can influence anticoagulation control when embedded in a prespecified protocol [36]. A generative dose interface should operate within a bounded action space, reject incompatible units, expose uncertainty and defer to validated PK/PD or Bayesian calculations where available [3,9,24]. Clinical evaluation should measure target attainment and toxicity rather than agreement with historical prescribing [9,23,36].

5.4 Trial emulation and pharmacovigilance

Representation learning can prioritize repurposing hypotheses through trial emulation in longitudinal records, but causal validity remains vulnerable to confounding, treatment channeling and missing outcomes [34]. Medication-related GenAI studies have evaluated adverse-event extraction, interaction checking and pharmacovigilance, but the evidence base remains dominated by retrospective or benchmark evaluations [13]. Clinical deployment should therefore be assessed for serious-event recall, false-alert burden, reviewer workload and downstream action [9,13,15]. As summarised in **Table 2**, evidence maturity and minimum validation requirements differ substantially across the major pharmacological use cases.

Table 2. Evidence maturity and minimum requirements for major applications.

Use case	Current evidence	Minimum evidence before clinical use	Principal failure mode
Target and compound generation	Experimental confirmation and selected early clinical translation [16,21,31]	Independent synthesis or expression, assay confirmation, selectivity, ADMET and reproducibility [7,16,31]	Optimization to model artefacts or invalid chemistry [7]
PGx decision support	Clinical utility exists for selected gene-drug pathways; the	Validated result, current guidance, interaction checks and local	Incorrect phenotype translation or obsolete recommendation

	generative interface remains early [2,12,35]	workflow evaluation [2,9,12]	
Dose individualization	Clinical evidence exists for selected protocolized algorithms; generative dose interfaces remain early [13,36]	External calibration, prospective target attainment, safety limits and unit testing [9,24,36]	Unit error or unsupported extrapolation
Synthetic patients and digital twins	Evidence is predominantly enabling or retrospective [8,17,18]	Fidelity, utility, privacy, dynamic updating and prospective evaluation [8,9,17]	Realistic but clinically impossible or privacy-leaking trajectory
Pharmacovigilance	Medication-related GenAI evidence is mainly retrospective or benchmark-based [13]	High-sensitivity evaluation, subgroup testing and reviewer-workload analysis [9,13,23]	Omission of serious events or excessive false alerts

Note: Evidence maturity is intentionally conservative; publication of a model or proof of concept is not equivalent to demonstrated clinical utility [9,23,24].

6. Synthetic patients and digital twins

Synthetic patient data reproduce selected statistical or structural properties of source populations and can support software testing, education and method development [17,18]. Fidelity, task utility, privacy and clinical plausibility require separate evaluation; realistic marginal distributions do not establish safe use for a clinical task [17]. Synthetic records should remain labelled and versioned and should not replace external validation on real patients [9,17,18].

A pharmacological digital twin is a patient-linked model initialized with observed data, updated over time and used to simulate responses under alternative interventions [8]. Mechanistic components can preserve physiological and PK/PD constraints, whereas learned components can represent residual nonlinear effects [8]. A twin intended to influence treatment should be evaluated for longitudinal calibration, sensitivity to missing observations, update stability and the

consequences of erroneous simulations [8,9,24].

6.1 Positioning the Data-to-Dose framework

The proposed framework differs from existing clinical-AI reporting and regulatory frameworks by placing pharmacological action boundaries, mechanistic PK/PD constraints, candidate

generation and post-treatment adaptation within one sequential pathway. DECIDE-AI primarily guides early clinical evaluation, CONSORT-AI and TRIPOD+AI guide reporting, and FDA change-control guidance addresses controlled modification [9,22-24]. The Data-to-Dose framework connects these governance principles to the specific transition from patient-state representation to a reviewable drug, dose or schedule proposal. Supplementary Table S3 compares the purpose, stages and residual gaps of these frameworks.

7. Proposed Data-to-Dose framework

We propose an eight-step Data-to-Dose pathway that makes every clinically consequential transition inspectable and reversible. The pathway specifies intended use; harmonizes data with provenance; constructs a longitudinal patient state; generates candidate drugs, doses or schedules; applies mechanistic, causal and safety

constraints; ranks benefit-risk with calibrated uncertainty; obtains professional and patient authorisation; and monitors outcomes. These controls align with clinical-evaluation and lifecycle-governance principles [9,15,22]. The system should abstain when required variables are missing, evidence conflicts or uncertainty exceeds

a prespecified threshold [9,14,15]. Controlled updates require versioning, locked safety tests, revalidation and the option to restrict or withdraw the system [15,22,37]. As illustrated in **Figure 2**, the framework uses sequential decision gates and a controlled feedback loop.




Figure 2. Proposed Data-to-Dose translation framework.

Note: A prespecified clinical question is converted into candidate actions only after provenance review, patient-state construction, pharmacological constraints, uncertainty-aware ranking and professional authorisation. Treatment outcomes enter a controlled monitoring and update cycle rather than unrestricted self-modification [9,15,22].

8. Validation, safety and governance

Technical evaluation must match the output. Molecular generators require chemical validity and experimental confirmation [7,16,21]. Medication-facing text systems require factuality, completeness and harmful-omission testing [13-15]. Synthetic records require task-specific fidelity, utility, privacy and plausibility assessment [17,18]. Dose systems require unit integrity, calibration and safety-limit testing [9,24,36]. External and temporal validation

should preserve the intended workflow and evaluate changes in population, measurements, formularies and practice [9,23,24]. Subgroup analysis is necessary because deployed health algorithms have reproduced structural inequity and differential underdiagnosis [32,38].

Silent or shadow-mode evaluation can reveal data latency, workflow mismatch and unexpected failure before outputs influence care [9]. Prospective studies should measure clinically relevant outcomes, user interaction and net benefit rather than benchmark accuracy alone [9,23,24]. The evidence threshold should increase with autonomy, patient exposure, irreversibility and potential severity of harm [9,14,22]. As illustrated in **Figure 3**, the validation pathway progresses from pre-deployment testing to post-deployment surveillance.




Figure 3. Validation and governance lifecycle.

Evidence progresses from technical and mechanistic validation through silent workflow testing, comparative clinical-utility studies, controlled deployment and post-deployment surveillance [9,15,22]. The evidence threshold rises with autonomy, patient exposure, irreversibility and potential severity of harm [9,14,22].

Generative medication systems can produce unsupported content or omit clinically important information [13-15]. Bias can enter through training data, outcome definitions and historical

treatment patterns, while sensitive health data create privacy and cybersecurity obligations [15,38,39]. Federated learning can reduce data centralization but does not remove heterogeneity, inference or governance risks [39]. WHO guidance, FDA change-control guidance and the EU AI Act support risk management, transparency, accountability and controlled modification [14,22,37]. Responsibilities should be explicit across developers, institutions, laboratories, clinicians and governance bodies [14,15,22]. As summarised in **Table 3**, the principal failure modes require layered detection and safeguard strategies.

Table 3. Major failure modes and layered safeguards.

Failure mode	Potential consequence	Detection	Required safeguards
Unsupported evidence or harmful omission	Incorrect drug, dose or monitoring plan [13-15]	Source-text comparison, omission testing and expert review [9,13,14]	Curated retrieval, source traceability, deterministic checks and abstention [9,14,15]
Invalid chemical or biological generation	False target confidence, wasted synthesis or unexpected toxicity [7,11]	Chemical and structural checks plus orthogonal assays [7,16,21]	Hard constraints, synthesis planning, ADMET and experimental confirmation [7,16,21]
Temporal, formulation or unit error	Obsolete evidence or order-of-magnitude dose error [9,13,24]	Timeline reconstruction, normalization and boundary tests [9,13,24]	Deterministic unit engine, freshness limits and mandatory confirmation [9,14,15]

Population bias or privacy leakage	Unequal benefit, excess harm or disclosure [32,38,39]	Subgroup calibration, inference testing and audit [32,38,39]	Diverse data, local validation, access controls and privacy testing [15,32,39]
Model or workflow drift	Progressive performance decline after data, guideline or workflow change [15,22,37]	Outcome, calibration, override and incident monitoring [9,15,22]	Version control, controlled updates, revalidation, restriction or withdrawal [15,22,37]

Note: No single safeguard is sufficient; controls should be layered and proportionate to the intended use and potential harm [9,14,15].

9. Recent advancements (2024–2026)

AlphaFold 3 broadened structure prediction to complexes containing proteins, nucleic acids, small molecules, ions and modified residues [20]. TxGNN introduced zero-shot drug-repurposing ranking with interpretable biomedical paths [33].

PanDerm evaluated multimodal representation learning across dermatological imaging modalities, whereas SleepFM used polysomnography pretraining for disease prediction [19,29]. The randomised phase 2a rentosertib study is an important clinical milestone for an AI-discovered target and molecule, but larger and longer studies are required [31]. These advances mark movement from isolated benchmarks towards translation; they do not validate autonomous precision prescribing [9,13,31].

10. Future research priorities

Future studies should compare multimodal foundation models with validated clinical, genomic and pharmacometric baselines rather than weak algorithmic comparators [2,9,24]. Hybrid models should integrate mechanistic PK/PD constraints with learned representations while preserving uncertainty [3,8]. Synthetic-data research should establish task-specific thresholds for fidelity, privacy, fairness and clinical plausibility [17,18]. Digital twins require prospective studies of dynamic calibration across progression, adherence change and treatment switching [8]. Human-factors studies should quantify automation bias, shared decision-making, workload and accountability [9,13,15]. Multi-site prospective benchmarks should include longitudinal medication context, rare serious harms, diverse populations and explicit abstention cases [9,32,38]. Lifecycle research should define measurable triggers for update, restriction and de-implementation [15,22,37].

11. Limitations of this review

This review has limitations. The search was restricted to English-language literature and was targeted rather than exhaustive; therefore, relevant studies may have been missed. Rapid model and regulatory changes may reduce the durability of some conclusions. The evidence base spans heterogeneous molecular, clinical and governance studies, which prevents quantitative synthesis and direct comparison across use cases. Many generative medication applications remain retrospective or benchmark-based, and publication bias may favour positive demonstrations. Four author-associated articles were included only where their published scope was relevant; major claims were supported preferentially by independent original studies or authoritative guidance. These limitations reinforce the need for prospective, multi-site and independently evaluated research.

12. Conclusion

Generative health data science can represent complex patient states and formulate reviewable pharmacological options [4,6,8]. It can also generate biological and chemical hypotheses [7,21,30]. The same flexibility creates risks of unsupported content, invalid molecular proposals, biased extrapolation and cascading workflow errors [7,13,15]. The proposed Data-to-Dose framework places intended use, provenance, mechanistic checks, calibrated uncertainty, professional authorisation and monitoring at the centre of translation. Evidence is substantial for enabling technologies and selected PGx pathways, whereas generative patient-level drug selection and dosing remain investigational [12,13,31]. Clinical authority should be earned through demonstrated benefit, safety and lifecycle accountability rather than model scale or fluency [9,14,22].

Declarations

Ethics approval and consent to participate: Not applicable

Consent for publication: Not applicable.

Funding: Not applicable.

References

1. Topol EJ. High-performance medicine: the convergence of human and artificial intelligence. *Nat Med*. 2019;25(1):44-56. doi:10.1038/s41591-018-0300-7.
2. Relling MV, Evans WE. Pharmacogenomics in the clinic. *Nature*. 2015;526(7573):343-350. doi:10.1038/nature15817.
3. Singh A, Faiyaz M, Yadav YC, Ahmad MG, Hussain S, Shakya D, et al. Precision pharmacotherapy: pharmacogenomics, biomarkers and omics technologies: identification of drugs and doses in complex diseases in patients. *Int J Drug Discov Med Res*. 2026;15(1):8-20. doi:10.63665/ijddmr.v15.i01.02.
4. Acosta JN, Falcone GJ, Rajpurkar P, Topol EJ. Multimodal biomedical AI. *Nat Med*. 2022;28(9):1773-1784. doi:10.1038/s41591-022-01981-2.
5. Rajpurkar P, Chen E, Banerjee O, Topol EJ. AI in health and medicine. *Nat Med*. 2022;28(1):31-38. doi:10.1038/s41591-021-01614-0.
6. Moor M, Banerjee O, Abad ZSH, Krumholz HM, Leskovec J, Topol EJ, et al. Foundation models for generalist medical artificial intelligence. *Nature*. 2023;616(7956):259-265. doi:10.1038/s41586-023-05881-4.
7. Du Y, Jamasb AR, Guo J, Fu T, Harris C, Wang Y, et al. Machine learning-aided generative molecular design. *Nat Mach Intell*. 2024;6:589-604. doi:10.1038/s42256-024-00843-5.
8. De Domenico M, Allegri L, Caldarelli G, d'Andrea V, Di Camillo B, Rocha LM, et al. Challenges and opportunities for digital twins in precision medicine from a complex systems perspective. *NPJ Digit Med*. 2025;8:37. doi:10.1038/s41746-024-01402-3.
9. Vasey B, Nagendran M, Campbell B, Clifton DA, Collins GS, Denaxas S, et al. Reporting guideline for the early-stage clinical evaluation of decision support systems driven by artificial intelligence: DECIDE-AI. *Nat Med*. 2022;28(5):924-933. doi:10.1038/s41591-022-01772-9.
10. Faiyaz M, Singh A, Kumar L, Hussain S, Yadav N, Ahmad MG. AI-driven innovations in healthcare and pharmaceutical sciences. *Int J Drug Discov Med Res*. 2026;15(1):20-27.
11. Faiyaz M, Yadav YC, Singh A, Ahmad MG, Singh M, Chaudhary N. Artificial intelligence-based drug discovery, virtual surgery software and machine learning in pharmacological targets discovery. *Int J Drug Discov Med Res*. 2026;15(1):1-7. doi:10.63665/ijddmr.v15.i01.01.
12. Swen JJ, van der Wouden CH, Manson LE, Abdullah-Koolmees H, Blagec K, Blagus T, et al. A 12-gene pharmacogenetic panel to prevent adverse drug reactions: an open-label, multicentre, controlled, cluster-randomised crossover implementation study. *Lancet*. 2023;401(10374):347-356. doi:10.1016/S0140-6736(22)01841-4.
13. Ong JCL, Chen MH, Ng N, Elangovan K, Tan NYT, Jin L, et al. A scoping review on generative AI and large language models in mitigating medication related harm. *NPJ Digit Med*. 2025;8:182. doi:10.1038/s41746-025-01565-7.
14. World Health organisation. Ethics and governance of artificial intelligence for health: guidance on large multi-modal models. Geneva: World Health organisation; 2024. Available from: <https://www.who.int/publications/i/item/9789240084759>. Accessed 22 July 2026.
15. Ahmad MG, Khairi H, Dhakad R, Shankar P, Faiyaz M, Singh A. Generative artificial intelligence in digital health: clinical translation, safety governance, regulatory science, and implementation roadmap. *World J Pharm Sci Res*. 2026;5(7):303-321. doi:10.5281/zenodo.21166928.
16. Stokes JM, Yang K, Swanson K, Jin W, Cubillos-Ruiz A, Donghia NM, et al. A deep learning approach to antibiotic discovery. *Cell*. 2020;180(4):688-702.e13. doi:10.1016/j.cell.2020.01.021.
17. Tucker A, Wang Z, Rotalinti Y, Myles P. Generating high-fidelity synthetic patient data for assessing machine learning healthcare software. *NPJ Digit Med*. 2020;3:147. doi:10.1038/s41746-020-00353-9.
18. Walonoski J, Kramer M, Nichols J, Quina A, Moesel C, Hall D, et al. Synthea: an approach, method, and software mechanism for generating synthetic patients and the synthetic electronic health care record. *J Am Med Inform Assoc*. 2018;25(3):230-238. doi:10.1093/jamia/ocx079.
19. Yan S, Yu Z, Primiero C, Vico-Alonso C, Wang Z, Yang L, et al. A multimodal vision foundation model for clinical dermatology. *Nat Med*. 2025;31:2691-2702. doi:10.1038/s41591-025-03747-y.
20. Abramson J, Adler J, Dunger J, Evans R, Green T, Pritzel A, et al. Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature*. 2024;630(8016):493-500. doi:10.1038/s41586-024-07487-w.
21. Watson JL, Juergens D, Bennett NR, Trippe BL, Yim J, Eisenach HE, et al. De novo design of protein structure and function with RFdiffusion. *Nature*. 2023;620(7976):1089-1100. doi:10.1038/s41586-023-06415-8.
22. US Food and Drug Administration. Marketing submission recommendations for a predetermined change control plan for artificial intelligence-enabled device software functions: guidance for industry and Food and Drug Administration staff. Silver Spring (MD): US Food and Drug Administration; 2024. Available from: <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/marketing-submission-recommendations-predetermined-change-control-plan-artificial>. Accessed 22 July 2026.
23. Liu X, Cruz Rivera S, Moher D, Calvert MJ, Denniston AK, SPIRIT-AI and CONSORT-AI Working Group. Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: the CONSORT-AI extension. *Nat Med*. 2020;26(9):1364-1374. doi:10.1038/s41591-020-1034-x.
24. Collins GS, Moons KGM, Dhiman P, Riley RD, Beam AL, Van Calster B, et al. TRIPOD+AI statement: updated guidance for reporting clinical prediction models that use regression or machine learning methods. *BMJ*. 2024;385:e078378. doi:10.1136/bmj-2023-078378.
25. Gomez-Bombarelli R, Wei JN, Duvenaud D, Hernandez-Lobato JM, Sanchez-Lengeling B,

- Sheberla D, et al. Automatic chemical design using a data-driven continuous representation of molecules. *ACS Cent Sci.* 2018;4(2):268-276. doi:10.1021/acscentsci.7b00572.
26. Cui H, Wang C, Maan H, Pang K, Luo F, Duan N, et al. scGPT: toward building a foundation model for single-cell multi-omics using generative AI. *Nat Methods.* 2024;21(8):1470-1480. doi:10.1038/s41592-024-02201-0.
27. Singhal K, Azizi S, Tu T, Mahdavi SS, Wei J, Chung HW, et al. Large language models encode clinical knowledge. *Nature.* 2023;620(7972):172-180. doi:10.1038/s41586-023-06291-2.
28. Jiang LY, Liu XC, Nejatian NP, Nasir-Moin M, Wang D, Abidin A, et al. Health system-scale language models are all-purpose prediction engines. *Nature.* 2023;619(7969):357-362. doi:10.1038/s41586-023-06160-y.
29. Thapa R, Kjaer MR, He B, Covert I, Moore H IV, Hanif U, et al. A multimodal sleep foundation model for disease prediction. *Nat Med.* 2026;32(2):752-762. doi:10.1038/s41591-025-04133-4.
30. Theodoris CV, Xiao L, Chopra A, Chaffin MD, Al Sayed ZR, Hill MC, et al. Transfer learning enables predictions in network biology. *Nature.* 2023;618(7965):616-624. doi:10.1038/s41586-023-06139-9.
31. Xu Z, Ren F, Wang P, Cao J, Tan C, Ma D, et al. A generative AI-discovered TNIK inhibitor for idiopathic pulmonary fibrosis: a randomised phase 2a trial. *Nat Med.* 2025;31:2602-2610. doi:10.1038/s41591-025-03743-2.
32. Seyyed-Kalantari L, Zhang H, McDermott MBA, Chen IY, Ghassemi M. Underdiagnosis bias of artificial intelligence algorithms applied to chest radiographs in underserved patient populations. *Nat Med.* 2021;27(12):2176-2182. doi:10.1038/s41591-021-01595-0.
33. Huang K, Chandak P, Wang Q, Havaladar S, Vaid A, Leskovec J, et al. A foundation model for clinician-centred drug repurposing. *Nat Med.* 2024;30:3601-3613. doi:10.1038/s41591-024-03233-x.
34. Liu R, Wei L, Zhang P. A deep learning framework for drug repurposing via emulating clinical trials on real-world patient data. *Nat Mach Intell.* 2021;3(1):68-75. doi:10.1038/s42256-020-00276-w.
35. Claassens DMF, Vos GJA, Bergmeijer TO, Hermanides RS, van 't Hof AWJ, van der Harst P, et al. A genotype-guided strategy for oral P2Y12 inhibitors in primary PCI. *N Engl J Med.* 2019;381(17):1621-1631. doi:10.1056/NEJMoa1907096.
36. Gage BF, Bass AR, Lin H, Woller SC, Stevens SM, Al-Hammadi N, et al. Effect of genotype-guided warfarin dosing on clinical events and anticoagulation control among patients undergoing hip or knee arthroplasty: the GIFT randomised clinical trial. *JAMA.* 2017;318(12):1115-1124. doi:10.1001/jama.2017.11469.
37. European Parliament and Council of the European Union. Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Off J Eur Union.* 2024;L 2024/1689. Available from: <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>. Accessed 22 July 2026.
38. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. *Science.* 2019;366(6464):447-453. doi:10.1126/science.aax2342.
39. Rieke N, Hancox J, Li W, Milletari F, Roth HR, Albarqouni S, et al. The future of digital health with federated learning. *NPJ Digit Med.* 2020;3:119. doi:10.1038/s41746-020-00323-1.