

Review Article

The role of artificial intelligence in clinical trials: opportunities

Vijaykumar Gawali*

Bhaktivedanta Hospital and Research Institute, Mira Road, Thane, Maharashtra, India

Received: 18 June 2026

Accepted: 06 October 2026

*Correspondence:

Dr. Vijaykumar Gawali,

E-mail: drvijaykumar.g@bhaktivedantahospital.com

Copyright: © the author(s), publisher and licensee Medip Academy. This is an open-access article distributed under the terms of the Creative Commons Attribution Non-Commercial License, which permits unrestricted non-commercial use, distribution, and reproduction in any medium, provided the original work is properly cited.

ABSTRACT

Artificial intelligence (AI) is now embedded across the clinical trial lifecycle, from target discovery and protocol design to recruitment, conduct, and post-marketing surveillance. Adoption has outpaced validation, and the randomised controlled trial (RCT) evidence base remains geographically uneven. For India, the question is no longer whether to adopt AI in trials, but how to do so without importing algorithms never validated on Indian patients. Narrative review of peer-reviewed literature (2018–2026) indexed in PubMed and major journal families, supplemented by official Indian regulatory documents (CDSCO, ICMR, MeitY). The review maps AI applications across the trial lifecycle, examines RCT evidence, and situates findings in India's regulatory context. A 2024 scoping review identified 86 AI RCTs (2018–2023); 81 per cent reported a positive primary endpoint, although 63 per cent were single-centre and only 19 per cent of post-2021 trials cited CONSORT-AI.¹ Machine learning predicts early trial termination with area under the curve 0.80.² The first generative-AI-discovered small molecule, rentosertib for idiopathic pulmonary fibrosis, completed phase 2a in 2025.³ Counter-examples persist: the Epic Sepsis Model showed sensitivity of 33 per cent on external validation.⁴ India's sepsis mortality (36.3 per cent overall; 50.8 per cent in septic shock) and projected antimicrobial-resistance deaths (11.8 million in South Asia by 2050) make trial-grade AI validation a national imperative. AI will reshape clinical trials, but only prospective, externally validated, reporting-standard-compliant studies, increasingly in Indian populations, will turn potential into evidence that regulators and clinicians can act on.

Keywords: Artificial intelligence, Machine learning, Clinical trials, CONSORT-AI, India, Pharmacovigilance, Antimicrobial resistance

INTRODUCTION

In two decades of supervising international drug trials at Indian tertiary-care hospitals, I have watched clinical research move through three distinct technological eras: the paper case-report-form era of the late 1990s, the electronic data capture era of the 2000s, and the cloud-and-API era of the 2010s. We are now in a fourth, the era of AI, and unlike the previous three it is not merely changing how we collect, store, or transmit trial data. What it changes is the questions we are able to ask of that data.

Over the last 5 years The use of AI in healthcare has expanded , with multiple studies reporting that medical

AI models can match or exceed clinician performance on specific diagnostic tasks.⁷ However only a small minority have been evaluated in prospective randomised controlled trials which is observed in 300 AI-enabled medical devices approved by The US Food and Drug Administration (RCT).¹ The gap between performance on a held-out dataset and performance in real clinical workflow is wider than is commonly acknowledged. That gap is where clinical trials, the discipline I have spent my career inside, must do their work.

This review is written from the standpoint of a practising clinician-researcher in a 750-bed tertiary-care hospital in suburban Mumbai. Two facts shape the perspective. First, India remains under-represented in AI clinical trial datasets despite contributing a disproportionate share of

the global disease burden in oncology, infectious disease, and antimicrobial resistance.^{5,6} Second, India's regulatory infrastructure for AI in health research—anchored by the New Drugs and Clinical Trials Rules, 2019 (NDCT Rules) of the Central Drugs Standard Control Organisation, the Indian Council of Medical Research's 2023 Ethical Guidelines for AI in Biomedical Research and Healthcare, and the Cabinet-approved IndiaAI Mission of March 2024—has matured to a point where indigenous AI clinical trials are not just possible but expected.⁸⁻¹⁰

The review has three objectives. First, to map where AI now operates across the clinical trial lifecycle. Second, to summarise the evidence—including the RCTs, the negative examples, and the methodological gaps. Third, to argue, with as much specificity as a single review can carry, what the Indian clinical research community should be building, validating, and publishing in the next five years. The review is narrative rather than systematic; it is intended to orient researchers, ethics committee members, sponsors, and regulators to a fast-moving field, not to substitute for primary-source reading on any specific application.

LITERATURE SEARCH

As a narrative review, this article does not follow PRISMA. The literature search was conducted in PubMed, Embase, the Lancet, JAMA, BMJ, Nature and Springer journal families, and the American Society of Clinical Oncology Educational Book, covering the period from 1 January 2018 to 30 April 2026. Search terms combined 'artificial intelligence', 'machine learning', 'deep learning', 'generative AI' and 'large language model' with 'clinical trial', 'randomised controlled trial', 'drug discovery', 'trial design', 'patient recruitment', 'eligibility', 'pharmacovigilance' and 'reporting guideline'. Indian regulatory documents were retrieved directly from the official sites of the Central Drugs Standard Control Organisation (CDSCO), the Indian Council of Medical Research (ICMR), and the Ministry of Electronics and Information Technology (MeitY).

Studies were prioritised if they reported prospective clinical evaluation (preferably RCT), proposed or evaluated a methodological framework with regulatory uptake (CONSORT-AI, SPIRIT-AI, DECIDE-AI, TRIPOD+AI), or characterised the Indian healthcare-AI landscape. Editorials and pre-prints were excluded. Where multiple secondary sources discussed the same primary study, the primary publication was cited. Every reference cited in this manuscript was verified by the author against PubMed identifiers, DOIs, or an official URL with an access date.

AI ACROSS THE CLINICAL TRIAL LIFECYCLE

In trials, AI is less a single intervention than a layer of computational support running across every stage of the

process. Table 1 maps these touchpoints; the rest of this section takes five lifecycle phases in turn.

Pre-trial: Target identification and drug discovery

Drug discovery has been the most visible early proving ground for AI in trials. The cost of bringing a new molecule to approval, once failures are capitalised in, now exceeds one billion US dollars per approved compound, and success probabilities are particularly low in oncology.¹¹ A 2026 scoping review of 19 applied studies on AI in drug discovery found that machine learning is most consistently used for hit discovery (DNA-encoded library screening, virtual screening), generative de novo design, lead optimisation, and pharmacokinetic and toxicity prediction.¹¹

The case most often raised here is the 2025 Nature Medicine publication of the GENESIS-IPF phase 2a trial of rentosertib, a first-in-class TNIK inhibitor whose target and molecule were both identified through generative AI.³ A randomised, placebo-controlled, multicentre trial across 22 sites in China showed a mean improvement in forced vital capacity of 98.4 mL in the 60 mg once-daily arm, compared with a 20.3 mL decline in placebo, with a treatment-emergent adverse event profile broadly comparable across arms.³ The trial enrolled 71 patients; it is too small to be a regulatory pivot, but it is the first time a generative-AI-discovered small molecule has produced randomised efficacy data in humans. The discovery-to-phase-2a timeline was compressed from the conventional five-to-seven years to under three.

Trial design and protocol optimisation

Approximately one in seven interventional trials registered on ClinicalTrials.gov terminates prematurely.² The economic and ethical waste this represents (patients who consented to research that never produced an answer) has motivated several machine-learning approaches to protocol optimisation. Kavalci and Hartshorn developed an ensemble model on 420,268 ClinicalTrials.gov records with 24 study-design features plus 12,864 eligibility-criteria features derived from the CHIA dataset; their gradient-boosting classifier achieved an area under the receiver operating characteristic curve of 0.80 and balanced accuracy of 0.70 on the held-out set, with Shapley value interpretability flagging the protocol elements most strongly associated with predicted termination.² The clinical value of such a tool is not as a verdict on a protocol, but as a peer-review aid that surfaces design choices likely to cause site failure or under-enrolment before the first patient is screened.

Adaptive trial designs, long advocated in oncology, gain additional traction when paired with AI-driven simulation. Real-time modification of dose, sample size, or stratification rules, based on accumulating safety and efficacy signals, moves from a statistically demanding manual exercise to one supported by continuous

machine-learning inference.¹² In oncology specifically, ML-based patient outcome simulations are being used to enhance stratification strategies and statistical power before the trial opens.¹²

Patient recruitment and eligibility screening

Patient-trial matching has long been one of the most labour-intensive steps in clinical research. In early-phase oncology trials, eligible patients were estimated 30 per cent of referred, and because of disease progression or toxicity 20 per cent of enrolled patients discontinued early.¹² TrialGPT, developed in collaboration with United States National Institutes of Health, uses large language models to find appropriate trials from free-text patient descriptions and outputs a ranked list of matches.¹² A second large language model, OncoLLM, was 1st to be federally funded specifically for clinical-trial matching, and is integrated into a patient records interpretation for semantic clinical trial matching system.¹²

From an Indian perspective, the more important application is not matching patients to existing trials but identifying, ahead of time, those who would qualify for a planned one. This matters most in conditions like sepsis, where eligibility windows are short and where Indian patients are systematically under-represented in global datasets. Natural-language processing applied to electronic health records can compress what is currently a manual chart-review exercise into minutes. The catch is that the underlying record must be in structured English text, a constraint that is partly cultural and partly infrastructural in most Indian hospitals.

Trial conduct, monitoring and safety

Once a trial is open, AI's contribution is largely operational: automated data entry validation, anomaly detection across sites, predictive monitoring of likely protocol deviations, and risk-based monitoring algorithms that decide which sites need an in-person visit. These tools improve data fidelity and reduce investigator burden, and they enable the practical decentralisation of trials.¹² Computational pathology, radiomics and spatial omics, all of which depend on deep learning, are now used both as eligibility criteria (for stratification) and as response assessments (as biomarkers).¹²

Pharmacovigilance is the area where AI's promise must be most carefully circumscribed. A widely deployed proprietary sepsis prediction model, the Epic Sepsis Model v1.0, was externally validated across 38,455 hospitalisations at the University of Michigan in 2018–2019. The model achieved a sensitivity of 33 per cent, a positive predictive value of 12 per cent, and an area under the curve of 0.63-and produced a substantial burden of alert fatigue.⁴ A subsequent 2024 external validation in two county emergency departments in Houston, Texas, reported a sensitivity of just 14.7 per cent at the recommended threshold.⁴

It is the cautionary tale clinical AI keeps returning to: a model that performs adequately in development can fail badly in deployment, and the failures are often invisible until they have already affected care. The lesson for trials is blunt. External validation in the trial population, using the trial's actual workflow, is not optional.

Table 1: Artificial intelligence applications across the clinical trial lifecycle.

Trial phases	AI application	Representative evidence (citation)
Pre-trial: discovery	Target identification; generative de novo molecular design; virtual screening; ADME/PK and toxicity prediction	Xu and Abaneme ^{3,11}
Protocol design	Termination-risk prediction; adaptive design simulation; outcome simulation for power calculation	Kavalci and Azenkot ^{2,12}
Recruitment and eligibility	NLP-based patient-trial matching (TrialGPT, OncoLLM); EHR mining for cohort identification	Azenkot ¹²
Trial conduct	Automated data validation; risk-based monitoring; computational pathology and radiomics as imaging endpoints	Han and Azenkot ^{1,12}
Analysis	Risk stratification for sub-cohort analysis; prediction-augmented endpoints	Kang ¹³
Safety and post-marketing	Real-time deterioration alerting (with caveats); NLP-based pharmacovigilance signal detection	Wong ⁴

*ADME, absorption, distribution, metabolism, excretion; EHR, electronic health record; NLP, natural language processing; PK, pharmacokinetics.

Trial analysis and post-marketing

Within trial analysis, AI methods are still adjuncts to traditional inferential statistics rather than substitutes. A useful framing has been offered in radiation oncology: machine-learning models can classify trial participants into risk strata, and the between-arm comparison is then

performed using conventional methods.¹³ In a single-institution quality-improvement trial at Duke, a gradient-boosted-tree model classified radiation-therapy patients into high- and low-risk for acute care visits; the high-risk arm was randomised to twice-weekly versus standard once-weekly clinic visits during radiation, and the intensive arm reduced acute care visits from 22 per cent

to 12 per cent.¹³ The AI did not replace the statistician. It shaped the cohort the statistician then analysed.

In post-marketing, the most concrete application is signal detection in spontaneous adverse-event databases and electronic health records, an area where natural-language-processing models can pre-process narrative free text at a scale no human pharmacovigilance team can match. The regulatory question is how to qualify such tools so that the signals they raise, and equally the signals they miss, can be defended in front of a drug regulator. India's CDSCO, through the NDCT Rules, currently requires sponsor-and investigator-driven serious-adverse-event reporting within 14 days of knowledge.⁸ An AI tool that learns from this stream is plausible, but only with auditable methodology and explicit calibration on Indian post-marketing data.

REPORTING STANDARDS FOR AI TRIALS

The history of clinical trials is, in large part, a history of reporting standards: CONSORT for randomised trials, SPIRIT for protocols, STROBE for observational studies, PRISMA for systematic reviews. AI extensions to these standards are now established and should be treated by Indian ethics committees and sponsors as a baseline expectation, not an aspirational extra. Table 2 summarises them.

The CONSORT-AI extension, published in *Nature Medicine* and *BMJ* in 2020, adds 14 AI-specific items to CONSORT 2010, covering algorithm description, training and validation datasets, version control, handling of poor-quality or unavailable input data, human-AI interaction, and analysis of error cases.¹⁴ Its companion SPIRIT-AI extension applies the same logic to trial protocols.¹⁵ DECIDE-AI, published in 2022, fills a gap that CONSORT and SPIRIT do not cover—the early-stage clinical evaluation that sits between in silico performance and large-scale RCT.¹⁶ It is the standard most relevant for first-in-human deployment of a decision-support algorithm in any Indian hospital. TRIPOD+AI, published in *BMJ* in 2024, supersedes the original TRIPOD 2015 statement and provides harmonised guidance for reporting clinical prediction models irrespective of whether they use regression or machine learning.¹⁷

Uptake is the weak link. In the 2024 scoping review of 86 AI RCTs, only 12 of the 63 trials published after 2021 (19 per cent) cited the CONSORT-AI guideline.¹ For the Indian clinical research community, embedding these standards into institutional ethics committee templates and CDSCO submission packs would do more for trial quality than any new policy.

Table 2: AI-specific reporting standards relevant to clinical trials.

Standards	Scope	Citation/ year	Indian uptake (author observation)
CONSORT-AI	Reporting of completed RCTs evaluating AI interventions	Liu, <i>Nat Med</i> ¹⁴	Low—~19% citation among AI RCTs published 2021-2023 ¹
SPIRIT-AI	Protocols for RCTs of AI interventions	Cruz Rivera, <i>Nat Med</i> ¹⁵	Low—rarely required by EC templates
DECIDE-AI	Early-stage (pre-RCT) clinical evaluation of AI decision-support systems	Vasey, <i>Nat Med</i> ¹⁶	Very low—most first-in-Indian-hospital pilots are unpublished
TRIPOD+AI	Development and validation of clinical prediction models, ML or regression	Collins, <i>BMJ</i> ¹⁷	Increasing in journal peer review

*EC, ethics committee; ML, machine learning; RCT, randomised controlled trial.

THE INDIAN IMPERATIVE

The case for indigenous AI clinical trial activity in India rests on four pillars: disease burden, dataset diversity, regulatory readiness, and validation gaps in imported algorithms. Each deserves brief attention.

Disease burden

India's sepsis burden alone makes the point. Incidence is estimated at 540-640 per 100,000 population, with a total burden of approximately 89.6 lakh cases annually and mortality of 25 to 30 per cent.¹⁸ The SEPSIS-INDIA prospective multicentre registry, conducted across 19 Indian intensive care units between August 2022 and July

2023, reported a hospital mortality of 36.3 per cent for sepsis overall and 50.8 per cent for septic shock.⁵ The Global Research on Antimicrobial Resistance project forecasts that South Asia, led by India, will see 11.8 million deaths attributable to antimicrobial resistance between 2025 and 2050, the highest of any world region.⁶ An AI early-warning system that detects sepsis four to six hours sooner, the timescale reported in international AI sepsis trials, would have a public-health impact in India that no imported, untranslated algorithm can deliver.

Dataset diversity

Indian patient populations differ from the predominantly Caucasian and East Asian cohorts that dominate AI

training data. Of the 86 RCTs in the 2024 scoping review, just 22 reported race or ethnicity, and of those, 18 were US trials with a median White (non-Hispanic) proportion of 70.5 per cent.¹ Only three trials in China and one in South Korea reported a single Asian ethnicity. Indian participants are essentially absent from the published AI RCT evidence base. I do not mean this as a complaint about exclusion. It is better read as an opportunity. The Indian clinical research community can build the cohorts that the next generation of algorithms will need.

Regulatory readiness

Three Indian frameworks now make AI clinical research operationally feasible. The NDCT Rules, 2019, set timelines (30 working days for drugs developed in India, 90 for global trials), strengthened ethics committee composition and registration, and introduced statutory compensation for trial-related injury and death.⁸ The ICMR's 2023 Ethical Guidelines for AI in Biomedical Research and Healthcare provide the country's first ethics framework specific to AI applications, with principles covering autonomy, safety, trustworthiness, accountability, transparency, data privacy and validity.⁹ The IndiaAI Mission, with a budget outlay of Rs 10,371.92 crore approved by the Union Cabinet in March 2024, funds compute infrastructure, indigenous

foundational models on Indian-language datasets, and the IndiaAI Application Development Initiative-within which healthcare AI is an explicit priority.¹⁰ Together, these instruments give Indian clinical researchers something uncommon: patient access, computational infrastructure, ethics infrastructure, and a regulatory pathway, all in one place. Few global-south jurisdictions are this well-positioned.

Validation gaps in imported algorithms

Every AI tool currently marketed for clinical use in India was trained somewhere else. The Wong et al Epic Sepsis Model is the clearest illustration of why these matters, but the principle is more general.⁴ Indian hospitals run on different laboratory reference ranges, different microbiology spectra, different antibiotic-use patterns, different patient comorbidity profiles, and different documentation practices. An algorithm trained on a Boston or Beijing dataset and deployed in Mumbai or Mira Road is, formally, operating outside its validated domain. Indian Institutional Ethics Committees should require local prospective validation as a condition of approving any AI-augmented clinical study, and increasingly they do. Bhaktivedanta Hospital's ethics committee has, in the trials I have overseen over the past decade, taken this position consistently.

Table 3: Selected randomised controlled trial evidence on AI in clinical practice.

Studies/ year	Clinical setting	Key finding	Limitations
Han et al 2024 (scoping review of 86 AI RCTs)¹	All specialties, 2018-2023	81% (70/86) reported a positive primary endpoint; 69% used deep learning on imaging	63% single-centre; only 19% of post-2021 trials cited CONSORT-AI
Xu et al 2025 (GENESIS-IPF phase 2a)³	Idiopathic pulmonary fibrosis, China (22 sites)	First generative-AI-discovered small molecule (rentosertib) to randomised efficacy data; FVC +98.4 mL on 60 mg QD vs -20.3 mL placebo	N=71; 12 weeks; single country
Wong et al 2021 (Epic Sepsis Model external validation)⁴	Sepsis, University of Michigan	On 38,455 hospitalisations: sensitivity 33%, PPV 12%, AUC 0.63; substantial alert-fatigue burden	Retrospective external validation; underscores need for prospective local validation
Kavalci and Hartshorn 2023²	Trial-design optimisation (ClinicalTrials.gov)	Gradient-boosting ensemble predicted early termination with AUROC 0.80, balanced accuracy 0.70; Shapley-value interpretation	Registry data only; not an RCT outcome; methodology applies pre-trial

*AUC, area under the receiver operating characteristic curve; AUROC, area under the receiver operating characteristic curve; FVC, forced vital capacity; PPV, positive predictive value; QD, once daily; RCT, randomised controlled trial.

CHALLENGES AND LIMITATIONS

Three challenges deserve specific attention because they recur in every serious discussion of AI in trials. First, bias. Algorithms trained on non-representative data perpetuate, and sometimes amplify, the disparities visible in their training sets. Without race, ethnicity, socioeconomic-status and rural-versus-urban representation in the development cohort, an algorithm

cannot be assumed to perform equitably at deployment. The CONSORT-AI item on subgroup analysis exists precisely for this reason.¹⁴

Second, explainability. Deep-learning models are not transparent in the way that a Cox proportional-hazards model is transparent. For a clinician taking treatment decisions on the basis of a model's output, and for a patient consenting to participate in a trial that uses such a

model, the ICMR principle that AI systems must be 'explainable, transparent and independent of a physician's view' is meaningful only if the technical methods to deliver explainability-SHAP values, attention-map visualisation, counterfactual reasoning-are routinely demanded and routinely reported.⁹ Kavalci and Hartshorn's use of Shapley values to explain individual predictions of trial termination is a worked example of how this can be done.²

Third, generalisability. Of the 86 AI RCTs identified in the 2024 scoping review, 79 (92 per cent) were conducted in a single country, and 54 (63 per cent) were conducted in a single centre, with a median enrolment of 359 patients.¹ Gastroenterology accounted for 37 (43 per cent) of all 86 trials, most of them video-based deep-learning systems supporting clinicians during endoscopy, and 24

of those 37 came from just four institutional groups (eight from Wuhan University, six from Wisdon AI, six from Medtronic, and four from Fujifilm).¹ That concentration is a liability rather than a strength: it raises both publication-bias and generalisability concerns.

FUTURE DIRECTIONS FOR THE INDIAN CLINICAL RESEARCH COMMUNITY

The argument for what the Indian clinical research community should be doing over the next five years is, in my view, both unambiguous and overdue. Five priorities deserve explicit institutional commitment.

First, build prospective Indian disease registries-modelled on SEPSIS-INDIA-that can serve as the training and external-validation substrate for indigenous algorithms.⁵ Second, mandate CONSORT-AI, SPIRIT-AI, DECIDE-AI and TRIPOD+AI compliance in institutional ethics committee templates and CDSCO clinical-trial submissions, so that reporting quality matches the field's rate of advance.¹⁴⁻¹⁷ Third, require local prospective validation of every imported AI tool before clinical deployment in any Indian trial, and structure that validation itself as a publishable DECIDE-AI study. Fourth, build trial-level partnerships between Indian hospitals, the India AI Innovation Centre, and the academic Centres of Excellence to co-develop algorithms rather than consume them.¹⁰ Fifth, train the next generation of clinician-researchers in AI ethics, bias auditing, and explainable machine learning-a curriculum that is conspicuously missing from most Indian postgraduate clinical research programmes.

These five priorities are not incremental. Executed seriously, they amount to a programme of work that would place India among the three or four jurisdictions globally capable of delivering AI clinical trial evidence at scale within the decade. The infrastructure is in place and the disease burden makes the case. As for the patients, in twenty years they have, overwhelmingly, given us their trust.

CONCLUSION

Artificial intelligence is no longer an emerging technology in clinical trials; it is an established layer of trial infrastructure that touches discovery, design, recruitment, conduct, analysis, and post-marketing surveillance. The published RCT evidence base is expanding rapidly but remains concentrated in a small number of countries, institutions, and specialties, and the gap between in silico performance and prospective clinical benefit is wider than commonly assumed. India's combination of disease burden, dataset diversity, regulatory readiness and ethics infrastructure makes it well placed to contribute to, and to lead, the next phase of this evidence base, particularly through prospective validation of imported algorithms in Indian populations and through indigenous discovery in sepsis, oncology, antimicrobial resistance and critical care. The reporting standards and the regulatory framework are already in place, and the patients are here. What is missing is the institutional will to do the work.

Funding: No funding sources

Conflict of interest: None declared

Ethical approval: Not required

REFERENCES

1. Han R, Acosta JN, Shakeri Z, Ioannidis JPA, Topol EJ, Rajpurkar P. Randomised controlled trials evaluating artificial intelligence in clinical practice: a scoping review. *Lancet Digit Health*. 2024;6(5):e367-73.
2. Kavalci E, Hartshorn A. Improving clinical trial design using interpretable machine learning based prediction of early trial termination. *Sci Rep*. 2023;13(1):121.
3. 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 randomized phase 2a trial. *Nat Med*. 2025;31(8):2602-10.
4. Wong A, Otles E, Donnelly JP, Krumm A, McCullough J, DeTroyer-Coolley O, et al. External validation of a widely implemented proprietary sepsis prediction model in hospitalized patients. *JAMA Intern Med*. 2021;181(8):1065-70.
5. Todi S, Mehta Y, Zirpe K, Amin P, Rao BK, Ramakrishnan N, et al. A multicentre prospective registry of one thousand sepsis patients admitted in Indian ICUs: (SEPSIS INDIA) study. *Crit Care*. 2024;28(1):375.
6. GBD 2021 Antimicrobial Resistance Collaborators. Global burden of bacterial antimicrobial resistance 1990-2021: a systematic analysis with forecasts to 2050. *Lancet*. 2024;404(10459):1199-226.
7. Topol EJ. High-performance medicine: the convergence of human and artificial intelligence. *Nat Med*. 2019;25(1):44-56.
8. Ministry of Health and Family Welfare, Government of India. New Drugs and Clinical Trials Rules, 2019.

- Gazette of India, G.S.R. 227(E). Available at: <https://cdsco.gov.in/opencms/opencms/en/Acts-and-rules/New-Drugs/>. Accessed on 15 May 2026.
9. Department of Health Research and Indian Council of Medical Research. Ethical Guidelines for Application of Artificial Intelligence in Biomedical Research and Healthcare. New Delhi: ICMR; 2023. ISBN: 978-93-5811-343-3. Available at: <https://main.icmr.nic.in/content/ethical-guidelines-application-artificial-intelligence-biomedical-research-and-healthcare>. Accessed on 15 May 2026.
 10. Press Information Bureau, Government of India. Cabinet approves ambitious IndiaAI Mission to strengthen the AI Innovation Ecosystem. PIB Release ID 2012355. Available at: <https://www.pib.gov.in/PressReleaseIframePage.aspx?PRID=2012355>. Accessed on 15 May 2026.
 11. Abaneme D, Chemjong K, Wiafe E, Moshood KO. AI-driven drug discovery and translational medicine: from computational models to clinical impact. *J Adv Med Med Res*. 2026;38(4):161-76.
 12. Azenkot T, Rivera DR, Stewart MD, Patel SP. Artificial intelligence and machine learning innovations to improve design and representativeness in oncology clinical trials. *Am Soc Clin Oncol Educ Book*. 2025;45:e473590.
 13. Kang J, Chowdhry A, Pugh S, Park J. Integrating artificial intelligence and machine learning into cancer clinical trials. *Semin Radiat Oncol*. 2023;33(4):386-94.
 14. 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):1364-74.
 15. Cruz Rivera S, Liu X, Chan AW, Denniston AK, Calvert MJ; SPIRIT-AI and CONSORT-AI Working Group. Guidelines for clinical trial protocols for interventions involving artificial intelligence: the SPIRIT-AI extension. *Nat Med*. 2020;26(9):1351-63.
 16. Vasey B, Nagendran M, Campbell B, Clifton DA, Collins GS, Denaxas S, et al; DECIDE-AI expert group. Reporting guideline for the early-stage clinical evaluation of decision support systems driven by artificial intelligence: DECIDE-AI. *Nat Med*. 2022;28(5):924-33.
 17. 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.
 18. Dedeepiya VD. Is it time to develop an Indian sepsis-related mortality prediction score? *Indian J Crit Care Med*. 2024;28(4):320-2.

Cite this article as: Gawali V. The role of artificial intelligence in clinical trials: opportunities. *Int J Clin Trials* 2026;13(4):487-93.