

A Multi-Level Framework for AI Adoption in Health Care: Integrating Behavioural, Organisational and Ethical Perspectives

Syed Kavish Ali

*Research Scholar, Department of Business Administration
Aligarh Muslim University, Aligarh, Uttar Pradesh, India*

ABSTRACT

Purpose: *Research on the adoption of artificial intelligence in health care has proceeded in four streams that seldom communicate: acceptance theorists study clinicians, organisational theorists study institutions, fairness theorists study algorithmic discrimination, and sustainability researchers study distributional consequences. This review integrates these four streams to build a conceptual framework for empirical work.*

Design / Methodology: *A narrative literature review was conducted using Scopus, Web of Science, PubMed, and Google Scholar (January 2015–March 2026). Around 430 records were filtered to 130 papers through title-abstract-fulltext screening, then thematically reviewed along four theoretical axes: UTAUT2, the Technology-Organisation-Environment (TOE) framework, the AI Bias Taxonomy, and the Triple Bottom Line (TBL).*

Findings: *The review reveals five gaps: fragmentation across individual and organisational levels; weak integration of bias into adoption theory; narrow outcome framings dominated by financial measures; scarce India-specific evidence; and under-use of mediation and moderation methods. The paper develops a multi-level conceptual model with 17 testable propositions linking acceptance, organisational context, bias, trust, and sustainability outcomes.*

Originality / Value: *The framework proposed here is the first to recognise algorithmic bias as a first-class adoption construct—a direct barrier to intention, a constraint on trust, and a moderator of the equity benefits of adoption. It offers a systematic empirical research agenda for scholars and a common language for leaders and policymakers assessing institutional AI decisions.*

Keywords: *Healthcare AI; UTAUT2; TOE Framework; Algorithmic Bias; Triple Bottom Line; Conceptual Framework.*

1. Introduction

A physician deciding whether to trust an AI-generated diagnostic suggestion, a hospital administrator deciding whether to invest in an AI platform, and a regulator establishing a bias audit requirement for medical software are all confronting the same phenomenon: the adoption of artificial intelligence in healthcare. They are, however, contending with it using dramatically different conceptual vocabularies. The clinician's concerns include usability, trust, and professional autonomy. The administrators are concerned with return on investment, integration with existing systems, and regulatory risk. The regulators care about fairness, accountability, and population-level outcomes. Each viewpoint is legitimate. Each has created its own body of scholarship. And each, on its own, provides only a partial picture of what AI adoption in healthcare implies.

This review begins with the subject of partiality. Healthcare AI research has grown at an incredible rate over the last decade—a recent bibliometric analysis shows a staggering, roughly two orders of magnitude increase in annual publication volume between 2010 and 2024 (Xie et al., 2025), but the theoretical infrastructure through which adoption is studied has remained fragmented in many ways. Acceptance researchers who follow in the footsteps of Venkatesh and colleagues continue to map individual level intention with the Unified Theory of Acceptance and Use of Technology 2 (UTAUT2). Organisational academics use the Technology-Organisation-Environment (TOE) framework developed

by Tornatzky, Fleischer, and Chakrabarti in 1990. Fairness academics, many of those are trained in machine learning or applied ethics rather than management studies, have developed a comprehensive taxonomy of algorithmic bias in medical settings. A smaller but rapidly growing minority is beginning to question if the outcomes of AI deployment, which adoption is purportedly aiming for, are best represented by financial measures alone, or if a broader sustainability perspective, such as the Triple Bottom Line (TBL), is more appropriate.

These four streams coexist peacefully since they rarely cite each other. It is possible to read several years of UTAUT-based healthcare-AI studies without encountering the fairness literature, and several years of algorithmic-fairness work without encountering organisational adoption theory. The intellectual loss is substantial. Questions such as "Does clinicians' awareness of bias affect their intention to use a tool?" or "Do institutions that adopt AI most aggressively also deliver the most equitable outcomes?" fit exactly in the gaps between both literatures and are so under-researched. This is the combined task covered in this review. It integrates the four literatures, pinpoints what is known and what is not known in each and forms a coherent conceptual structure - a set of seventeen propositions, which explain how individual acceptance, organisational context, bias perceptions, trust and sustainability outcomes relate to each other. The contribution it aims to make is twofold. The paradigm offers researchers a broader framework of empirical research than a one-theory study. To practitioners and policymakers, the synthesis offers a common conceptual map between those constituencies whose vocabularies have previously tended largely to pass over each other.

1.1 Why now, and why India in particular

The timing of such a synthesis matters. Industry surveys from the opening months of 2026 describe an inflection point in healthcare AI: the NVIDIA State of AI in Healthcare and Life Sciences 2026 report places active AI use among healthcare professionals at 63 percent, with a further 31 percent in evaluation (NVIDIA, 2026); the Doximity 2026 State of AI in Medicine Report records similar numbers from the physician side (Doximity, 2026). Analysts are increasingly describing 2026 as the year in which AI in healthcare moves from pilot projects into enterprise-scale deployment (BCG, 2026; Chief Healthcare Executive, 2026; Wolters Kluwer, 2026). A transition of this scale invites theoretical stocktaking. The frameworks used to study adoption now will shape the questions asked of it for years to come.

That argument takes on additional force when one considers India specifically. Over the last two years the country has assembled an unusually dense policy stack around healthcare AI. A national AI mission, approved in 2024, sits alongside a digital health backbone (the Ayushman Bharat Digital Mission) that has been expanding in parallel, and the most recent additions are the two initiatives rolled out at the India AI Impact Summit in Delhi in February 2026—one, known as SAHI, is principally a governance document setting out how AI in Indian healthcare should be developed and deployed; the other, BODH, is a benchmarking platform built to validate such tools against Indian-population data before they reach clinical users (PIB, 2026; Analytics Insight, 2026). What makes India an interesting empirical case is not any one of these pieces individually, but the fact that adoption, validation, and governance are all being worked out at the same time, in full public view, across a health system whose internal variation is striking. A tertiary hospital in a metropolitan area and a semi-urban district facility may both, on paper, be adopting "AI," but they are doing so under conditions so dissimilar that findings from Western hospital settings cannot simply be imported (Chettri et al., 2025). A review that centres India, then, is not making a provincial gesture; it is registering the fact that a literature built largely on Western empirical material now needs to be stress-tested against a substantially different kind of environment.

1.2 Objectives of this review

Four objectives guide the review. The first is descriptive: it outlines how each of the four theoretical streams currently contributes to the research of healthcare AI adoption. The second step is diagnostic: identify the points where the streams deviate and the substantive concerns that arise as a result. The third step is constructive: offer an integrated framework that connects the pathways and specifies testable propositions across person, organisational, inhibitor-based, and

outcome-based dimensions. The fourth step is to establish an empirical research programme that the framework enables and that, if pursued, would provide healthcare organisations and policymakers with a far stronger evidentiary basis for AI adoption decisions.

The remaining paper is organised in the following way. The review approach is described in Section 2. The theoretical bases of the four literatures are discussed in section 3. Their main empirical results are summarised in section 4. There are five gaps in research described in Section 5. Section 6 provides the suggested integrated framework and seventeen propositions. Part 7 examines policy implications, theoretical and practical implications. Section 8 recognises the limitations of the review. Section 9 suggests guidelines for future empirical research. Section ten concludes.

2. Review Methodology

A narrative, framework-oriented review method has been used in this paper instead of a proper PRISMA-compliant systematic review. The choice is not random. Systematic reviews are best suited to the research question that is sufficiently narrow so that all the possible eligible studies can be identified, coded, and compared on a set of variables that is shared and thus comparatively can be seen as a synthesis, but less appropriate to a synthesis across four different theoretical traditions with their own methodological norms. Contrastingly, a narrative synthesis enables the reviewer to tap into a range of sources—quantitative survey research, qualitative field research, meta-analyses, taxonomy articles, and policy reports—and balance their input against a conceptual scheme instead of a predetermined outcome measure. Since the intention in this case is to develop an integrated framework, but not to provide an answer to a single empirical question, the narrative method is the right one (Grant & Booth, 2009).

2.1 Search strategy and study selection

Literature was retrieved from four databases—Scopus, Web of Science, PubMed, and Google Scholar—covering publications between January 2015 and March 2026. The 2015 start date reflects the point at which deep-learning applications in medical imaging began to be documented in the clinical literature; earlier work, where relevant, was captured through reference chaining. Search strings combined framework-specific terms ("UTAUT2", "technology acceptance", "TOE framework", "AI bias", "algorithmic fairness", "triple bottom line") with domain-specific terms ("healthcare", "clinical", "hospital", "medical AI", "digital health"). The initial yield of approximately 430 records was reduced through duplicate removal (about 80 records), title and abstract screening (a further 140 excluded as either tangential or non-empirical commentary), and full-text review (another 80 excluded for lacking adequate methodological reporting or falling outside the scope). The 130 studies retained for thematic analysis span empirical surveys, systematic reviews, meta-analyses, theoretical essays, and policy briefings. A summary of the search protocol is provided in Table 1.

Table 1: Summary of Literature Search and Selection Protocol

Parameter	Specification
Databases searched	Scopus; Web of Science; PubMed; Google Scholar
Date range	January 2015 – March 2026
Search term combinations	Framework terms (UTAUT2, TOE, AI bias, algorithmic fairness, triple bottom line) × domain terms (healthcare, clinical, hospital, medical AI, digital health)
Inclusion criteria	Peer-reviewed empirical or theoretical studies examining AI adoption, acceptance, bias, trust, or sustainability outcomes in healthcare settings; English-language

Exclusion criteria	Editorials and commentaries lacking methodology; studies on non-medical AI applications; duplicate publications; studies where AI was tangential to the stated research question
Initial records identified	≈ 430
Duplicates removed	≈ 80
Excluded at title/abstract screening	≈ 140
Excluded at full-text review	≈ 80
Studies included in thematic synthesis	130

Source: Author's synthesis based on database searches conducted between November 2025 and March 2026.

2.2 Synthesis approach

Selected studies were coded along two axes. The first axis was the theoretical framework most heavily drawn upon—UTAUT2, TOE, AI Bias Taxonomy, TBL, or a combination. The second axis was the substantive focus—individual acceptance, organisational adoption decisions, bias and trust dynamics, or sustainability outcomes. Cross-cutting patterns were then examined: where did the frameworks converge on shared findings, where did they diverge, and where did they simply fail to address each other's concerns? This coding produced the four thematic sections that follow in Section 4, the gap analysis in Section 5, and the consolidated framework in Section 6.

There should be an important exception of the evidence base. The four literatures generate methodologically unbalanced evidence. UTAUT2 studies are mostly cross-sectional surveys evaluated with structural equation modelling; TOE studies have surveys, meta-analyses and case studies; the bias literature exists between algorithmic audit and qualitative clinician-attitude research and the sustainability literature is mostly a literature of conceptual essays with a limited number of empirical tests. Such imbalance is an observation, which informs the research-gap analysis as well as the recommendations of further research.

3. Theoretical Foundations

It would be critical to define what each of the four frameworks is supposed to achieve before delving into the findings of the research. Readers who have read any of them might find it unnecessary to follow this section; this is not an attempt at a complete illustration, but a vocabulary that is liable to be commonly used and will render the subsequent synthesis intelligible.

3.1 UTAUT2: individual-level acceptance

UTAUT2 (Venkatesh et al., 2012) evolved from a line of technology-acceptance models stretching back to Davis's (1989) Technology Acceptance Model. Its predecessor, UTAUT (Venkatesh et al., 2003), consolidated eight prior models into four predictors of behavioural intention: performance expectancy, effort expectancy, social influence, and facilitating conditions. The UTAUT2 extension added three further constructs—hedonic motivation, price value, and habit—and was validated in consumer-technology contexts where such factors were theoretically salient. Behavioural intention, the dependent variable, captures an individual's stated willingness to use a technology and has been shown across many studies to be a reasonably strong proxy for subsequent actual use (Tamilmani et al., 2019, 2021).

When applied to healthcare AI, the seven UTAUT2 constructs translate into familiar clinical concerns. Performance expectancy asks whether the tool is perceived to improve diagnostic accuracy, speed, or workflow. Effort expectancy captures the learning curve and ease of use. Social influence reflects pressure from peers, supervisors, and patients. Facilitating conditions denote the infrastructure, training, and institutional support available. Hedonic motivation is the degree of enjoyment or intellectual engagement the technology provides. Price value captures the perceived trade-off between cost (of the tool, of training time, of potential errors) and benefit. Habit reflects the degree to which use has become automatic. The framework's utility in healthcare has been repeatedly demonstrated (Ammenwerth, 2019; Lambert et al., 2023; Cobelli & Blasioli, 2023), though not without debate about the weights that different constructs should carry in clinical rather than consumer contexts.

3.2 The Technology–Organisation–Environment framework

Where UTAUT2 focuses on the individual, TOE (Tornatzky et al., 1990) focuses on the institution. It posits that an organisation's decision to adopt a technology is shaped by three contextual dimensions. The technological context concerns attributes of the technology itself: its relative advantage over existing practice, its compatibility with current systems, and the organisation's technical readiness to deploy it. The organisational context captures internal factors: top-management support, available slack resources, organisational culture, and existing data-governance capabilities. The environmental context includes external pressures: regulatory demands, competitive dynamics, government incentives, and the maturity of the vendor ecosystem. TOE is, in effect, a framework-of-frameworks: it specifies what kinds of factors matter, but the operationalisation of each dimension varies considerably across studies (Baker, 2012; Salah & Ayyash, 2024).

Meta-analytic evidence on TOE's predictive power in AI adoption has accumulated rapidly. Pinto et al.'s (2025) meta-analysis of twelve TOE-based AI-adoption studies concluded that organisational readiness and top-management support are the most consistent predictors, with technological fit second and environmental factors third in average effect size. This ordering matters theoretically: it suggests that even in a context of strong external pressure (such as India's active AI policy push), the internal capacity of the organisation remains decisive.

3.3 The AI Bias Taxonomy

Unlike UTAUT2 and TOE, which come from the information-systems and management literatures, the AI Bias Taxonomy has its roots in machine-learning fairness research and medical ethics. The taxonomy articulated by Cross et al. (2024), building on earlier work by Obermeyer et al. (2019), Rajkomar et al. (2018), and Chen et al. (2023), distinguishes six types of bias that can compromise medical AI systems. Historical bias reflects inequities already embedded in the training data. Representation bias arises when certain populations are under-represented in the data on which the model was trained. Measurement bias occurs when the features used to predict an outcome are systematically mismeasured for some groups. Label bias arises when the outcome labels themselves are biased—as in the now-canonical example of healthcare cost being used as a proxy for healthcare need (Obermeyer et al., 2019). Technical bias arises from algorithmic or architectural choices during model development. Minority bias refers to the under-performance of models on small, distinct subpopulations.

Empirical documentation of such biases has expanded quickly, with well-known examples including race-based inaccuracies in pulse oximetry readings (Sjoding et al., 2020), diagnostic-imaging models that can infer patient race from chest radiographs (Gichoya et al., 2023), and dermatological models trained predominantly on lighter skin tones (Chen et al., 2023). What the taxonomy literature offers, and what more descriptive surveys of bias have not, is a structured vocabulary for identifying where in the AI lifecycle bias enters, and how it compounds. This matters for the integration proposed here: treating bias as a single undifferentiated concern loses the structure that the taxonomy captures.

3.4 The Triple Bottom Line

The TBL concept, articulated by Elkington (1997), proposes that organisations should be evaluated on three dimensions of performance simultaneously—economic, social, and environmental—rather than on financial returns alone. Originally developed for corporate sustainability reporting, it has since been extended into domains as varied as supply-chain management, public administration, and higher education. Its application to healthcare AI is recent and relatively thin. A handful of studies (Schiavone et al., 2022; Alim et al., 2024; Ueda et al., 2024) have used TBL or TBL-adjacent constructs to argue that AI's value in healthcare should be assessed through equity and environmental cost as well as through cost reduction or clinical accuracy. None, to the author's knowledge, has embedded the TBL within an adoption framework that also theorises bias and trust.

The environmental dimension deserves particular attention. Recent analyses have drawn attention to the substantial energy and carbon footprints associated with training and running large-scale AI models, and to the fact that these costs are rising as deployment expands rather than stabilising (Ueda et al., 2024). In the healthcare context, where inference at scale is increasingly routine—every radiology AI read, every clinical-decision support query, every LLM-generated note consumes compute—this footprint is neither trivial nor static. A theoretical framework that treats adoption and outcomes separately, with no feedback between them, misses this feedback loop entirely.

4. Synthesis Of The Literature

Having set out what each framework is designed to do, this section turns to what the empirical literature has actually found. The synthesis is organised along the four theoretical axes, with a fifth sub-section on the Indian context, which has emerged as distinct enough in the reviewed literature to warrant separate treatment.

4.1 Individual-level acceptance: what UTAUT2 studies converge on

Across healthcare-AI studies that have deployed UTAUT2 since 2019, two findings recur with sufficient regularity to be treated as established. The first is the dominance of performance expectancy and facilitating conditions as predictors of behavioural intention. Lambert et al.'s (2023) integrative review across hospital settings identified the same pattern. Su et al. (2025), examining AI health-assistant adoption, and Dai et al. (2025), with a Chinese healthcare professional sample of 2,705, arrived at similar conclusions. The second recurring finding is the relatively modest weight of social influence in professional medical settings, which contrasts with its larger role in consumer-technology contexts. The plausible explanation is that clinicians, operating under professional autonomy norms and strong individual accountability, rely more on their own evaluation of a tool than on peer enthusiasm.

Unsettled are the influences of hedonic motivation, price value and habit. The concept of hedonic motivation, which was created as one of the consumer technologies aimed at pleasure, does not fit in professional clinical settings perfectly. It has been proposed that it should be replaced with other constructs like "professional satisfaction" or "cognitive engagement" (Cobelli & Blasioli, 2023). Theoretically, price value is crucial in resource-constrained environments but it is usually crudely measured, usually as a single perception item. Habit, by definition, is a post-adoption concept; research that simultaneously measures it and pre-adoption intention in cross-sectional studies risks confounding it with previous experience. One of the reasons why the framework is best complementary to other perspectives, and not substitutive, is these gaps in the UTAUT2 literature.

4.2 Organisational adoption: TOE's consistent and contested findings

The TOE literature on healthcare AI presents a cleaner picture at the aggregate level, largely because meta-analytic evidence now exists. Pinto et al. (2025), synthesising twelve studies, report organisational readiness and top-management support as the most consistent predictors across the aggregate evidence base, with technological fit second

and environmental factors third in average effect size. Hassan et al. (2025), working specifically on NHS settings in the UK, reach similar conclusions, though they add the particular importance of "visible value"—evidence that the AI system is producing demonstrable benefit in observable cases. Salah and Ayyash's (2024) systematic review of TOE in AI contexts confirms the pattern but notes considerable variability in how the three dimensions are operationalised, which limits cross-study comparability.

But within this overall sample consistency are tensions that the synthesis reveals. First, there are variations in the importance of the environmental dimension across countries. In the US and Western Europe, environmental issues are generally trumped by competition and vendor-ecosystem maturity. In emerging markets, it is government policy and regulatory uncertainties. The second issue relates to the interaction between organisational and technological factors. The evidence from the case studies (Thieme et al., 2025; Dwivedi et al., 2025) suggests that organisations that have skills in managing the data governance process see technological attributes differently to other organisations - they are more likely to adopt tools with incompatibilities because these can be managed. This is not captured in TOE operationalisations.

4.3 Algorithmic bias: from taxonomy to adoption inhibitor

The literature on bias in healthcare AI is where this review finds the most consequential gap. The evidence that bias exists, persists, and causes clinical harm is extensive and, at this point, not seriously contested (Obermeyer et al., 2019; Sjoding et al., 2020; Chen et al., 2023; Gichoya et al., 2023; Hasanzadeh et al., 2025; Siddique et al., 2024). What is far less developed is the translation of that evidence into the adoption literature. A clinician who is aware that dermatological AI tools perform less well on darker skin tones, or that clinical-prediction algorithms have historically disadvantaged Black patients, is making an adoption decision with those considerations in mind. Yet the standard UTAUT2 and TOE instruments have no construct that captures this concern directly. It is sometimes absorbed into perceived risk, sometimes into trust, and sometimes simply left out.

A small number of recent studies have begun to close this gap. Chettri et al. (2025), in an Indian context, treat bias as a stand-alone barrier to trustworthy AI adoption. None of these, however, develops a full theoretical account of how bias perceptions propagate through the adoption process—whether directly, through trust, or through both pathways simultaneously. Section 6 of this paper argues that the evidence supports a both/and rather than an either/or account.

4.4 Sustainability outcomes: a small but important literature

The fourth literature is the thinnest. A search combining "healthcare AI" with "sustainability" or "Triple Bottom Line" terms returns far fewer empirical studies than any of the preceding searches. What is available is nonetheless telling. Schiavone et al. (2022) use a sustainability-value lens to examine digital business models in healthcare. Alim et al. (2024) review environmental sustainability of healthcare AI specifically and highlight substantial gaps in measurement. Ueda et al. (2024) synthesise environmental, social, and governance considerations for medical AI. The shared conclusion across these studies is that AI's value in healthcare is routinely evaluated through narrow efficiency or accuracy metrics, while its distributive effects (who benefits, who is left out) and environmental effects (the energy and carbon costs of inference at scale) go largely unmeasured.

The implication of this measurement approach is not widely recognised. If the decision to adopt is made based on expected financial gain and clinical-accuracy gain, but not on the basis of equity or environmental cost, then organisations may adopt aggressively and find that they have achieved economic gains while creating externalized social and environmental costs. A model that gives these three dimensions prominence is, then, not a luxury—it is an analytical requirement.

4.5 The Indian context as a structurally distinct case

Running as a thread through each of the preceding sub-sections is the question of whether findings from Western and East Asian healthcare systems translate to Indian settings. The reviewed literature suggests the answer is a qualified no. Three distinguishing features of the Indian context recur across the studies that sample it. First, resource heterogeneity: the coexistence of well-equipped metro tertiary centres with semi-urban and rural facilities operating under sharply different constraints means that a single mean effect size for the country is less meaningful than the variance around that mean. Second, regulatory dynamism: the rapid introduction of the IndiaAI Mission (2024), the classification of diagnostic AI as Class C medical devices, the Digital Personal Data Protection Act (2023), and now SAHI and BODH (February 2026) means that the policy environment is moving faster than much of the published literature, which is already somewhat out of date by the time it appears. Third, representational concerns: the over-reliance of globally available AI tools on Western training data is a bias issue with particular force in a country whose population is both large and phenotypically distinct from those represented in many benchmark datasets (Kumar et al., 2024; Chettri et al., 2025). Each of these features argues for India-specific empirical work that cannot be substituted by the replication of Western studies.

Table 2: Representative Studies Across the Four Theoretical Streams

Stream	Representative study	Setting	Method	Key finding relevant to this review
UTAUT2	[Citation pending verification]	Singapore; clinical decision support	Survey, SEM	Performance expectancy and facilitating conditions dominate as predictors of behavioural intention
UTAUT2	Lambert et al. (2023)	Multi-country integrative review	Integrative review	Facilitating conditions, trust, and performance expectancy consistently strongest across settings
UTAUT2	Dai et al. (2025)	China; n = 2,705 healthcare professionals	Large-sample SEM	UTAUT2 explains ~62 % of variance in intention; trust emerges as complementary construct
TOE	Pinto et al. (2025)	Meta-analysis of 12 AI-adoption studies	Meta-analysis	Organisational readiness and top-management support are the most consistent cross-study predictors
TOE	Hassan et al. (2025)	UK NHS settings	Mixed methods	Visible value and governance structures are decisive in healthcare contexts
Bias	Obermeyer et al. (2019)	US hospital system	Algorithmic audit	Healthcare-cost proxy produced systematic racial bias in risk prediction
Bias	Cross et al. (2024)	Conceptual / taxonomy paper	Theoretical synthesis	Six-category bias taxonomy spanning historical, representation, measurement, label, technical, minority

Bias	[Citation pending verification]	AI-based medical devices	Survey, mediation analysis	Trust mediates the relationship between system attributes and adoption intention
TBL / Sustainability	Schiavone et al. (2022)	Healthcare value co-creation	Conceptual / case	Sustainability logic underpins emerging digital healthcare business models
TBL / Sustainability	Ueda et al. (2024)	Medical AI environmental review	Narrative review	Environmental costs of AI are material and under-measured in adoption studies
Indian context	Chettri et al. (2025)	Indian healthcare AI adoption	Narrative review	Representational bias is amplified in Indian populations; governance is evolving rapidly

Source: Author's synthesis of reviewed literature. Studies selected to illustrate the range of methodologies and findings within each stream; not exhaustive.

5. Research Gaps Identified

The thematic synthesis surfaces five gaps that, taken together, motivate the consolidated framework developed in the next section. They are not independent; each is, in part, a consequence of the others. But they are distinguishable enough to be named separately, and doing so clarifies what an integrated framework must do.

5.1 Gap 1: Theoretical fragmentation between individual and organisational perspectives

The most basic gap is the near-absence of studies that test UTAUT2 and TOE constructs simultaneously within a single empirical model. Pumplun et al. (2021) is a relatively rare example of an integrated TOE–UTAUT design, but it remains an outlier in a literature dominated by single-framework studies. The separation has methodological consequences. Studies that examine only individual-level constructs treat organisational context as background noise; studies that examine only organisational constructs treat individual acceptance as a black box. Yet the empirical adoption decision is made by individuals embedded in organisations. Modelling one without the other means that the estimated effect of, say, performance expectancy may absorb variance that properly belongs to facilitating conditions or organisational readiness, and vice versa.

5.2 Gap 2: Algorithmic bias weakly integrated into adoption theory

As discussed in Section 4.3, the evidence that bias exists and causes harm is extensive, but the evidence that bias perceptions influence adoption decisions is surprisingly sparse. The dominant adoption frameworks were developed before bias became a central concern in healthcare AI, and they have not been extensively retrofitted. Where bias is included, it tends to be subsumed into perceived risk or trust, which loses the multi-dimensional structure captured by the bias taxonomy. A proper theoretical account of how bias enters adoption decisions would treat it as a distinct, multi-faceted construct with its own direct effects and its own indirect effects through trust.

5.3 Gap 3: Outcome framings that privilege financial metrics

Even when studies consider outcomes (rather than just intention), the outcomes they consider are usually narrow. Return on investment, clinical accuracy, workflow efficiency, and patient-throughput metrics dominate. Equity outcomes, workforce-wellbeing outcomes, and environmental-cost outcomes appear rarely, and when they appear they are usually measured in isolation from the adoption decision itself. A framework that specifies AI adoption as a predictor of all three TBL dimensions simultaneously—and that permits the dimensions to interact with bias and trust—is, to the author’s knowledge, absent from the published literature.

5.4 Gap 4: Scarce India-specific evidence in a strategically important context

For reasons set out in Section 4.5, India is a setting whose empirical dynamics cannot be inferred from Western studies. Yet the published empirical base on Indian healthcare AI adoption using rigorous multi-theoretical frameworks is thin. Narrative reviews and policy essays exist (Chettri et al., 2025; Adithyan et al., 2024); large-scale structural-equation studies that test integrated models in Indian samples are rare. This is a gap that the integrated framework proposed here is specifically designed to be testable against.

5.5 Gap 5: Under-use of mediation and moderation modelling

The fifth gap is methodological. The substantive claims implicit in the reviewed literature—that trust mediates the effect of bias on intention; that bias conditions the equity benefits of adoption; that organisational readiness shapes the returns to technological investment—are all mediation or moderation claims. Yet the bulk of the published work relies on multiple regression or simple path analysis, which can handle direct effects reasonably well but is ill-suited to simultaneous tests of mediation and moderation within a complex theoretical model. Covariance-based and partial-least-squares structural equation modelling are available and are increasingly used (Dai et al., 2025), but in the healthcare AI literature they remain the exception rather than the rule.

Table 3: Summary of Research Gaps and Their Implications for Framework Design

#	Gap	Consequence	Framework implication
G1	Fragmentation between individual and organisational adoption theory	Effect sizes for either level may be biased when the other is omitted	Framework must integrate UTAUT2 and TOE within a single model
G2	Bias weakly integrated into adoption theory	Bias concerns are under-theorised as adoption determinants	Framework must treat bias as a first-class, multi-dimensional construct
G3	Outcome framings dominated by financial metrics	Equity and environmental effects of adoption are under-measured	Framework must include all three TBL dimensions as outcomes
G4	Scarce India-specific evidence	Findings from Western settings may not transfer	Framework must be testable in Indian contexts and sensitive to local conditions
G5	Under-use of mediation / moderation methods	Substantive mediation and moderation claims are stated but not tested	Framework must be specified in a form amenable to SEM-based testing

Source: Author's synthesis based on the literature review.

6. Proposed Integrated Conceptual Framework

The framework advanced here integrates the four literatures into a single multi-level model of AI adoption in healthcare. It is designed to do five things: preserve the explanatory structure of each source framework, permit the simultaneous estimation of individual-level and organisational-level effects, admit bias as a distinct adoption construct with both direct and trust-mediated effects, include the full TBL as outcome space, and specify moderation relationships that capture the conditional equity benefits of adoption. Seventeen propositions follow from this design, each representing a relationship between constructs that the reviewed literature supports but that has not been empirically tested within an integrated model.

6.1 Framework architecture

At the individual level, seven UTAUT2 constructs are proposed as direct predictors of behavioural intention to adopt AI. At the organisational level, three TOE dimensions are proposed as direct predictors of organisational AI adoption. Perceived AI bias, operationalised as a second-order construct reflecting the six sub-dimensions of the Cross et al. (2024) taxonomy, is proposed to affect behavioural intention directly and also to affect perceived trust, which in turn affects intention. Adoption (treated as the realised outcome of individual intention and organisational decision) is proposed to influence all three TBL sustainability outcomes. Finally, perceived bias is proposed to moderate the effect of adoption on social sustainability specifically—capturing the argument that AI deployed without attention to fairness can deliver economic and environmental returns while undermining equity.

The framework’s architecture is depicted graphically in Figure 1. It is organised around four zones: (i) an individual zone containing the seven UTAUT2 predictors and the behavioural-intention outcome; (ii) an organisational zone containing the three TOE predictors and the organisational-adoption outcome; (iii) an inhibitor zone containing perceived AI bias and perceived trust, positioned so that bias acts on both intention directly and on trust, with trust in turn feeding into intention; and (iv) an outcome zone containing economic, social, and environmental sustainability, predicted by adoption, with the bias × adoption interaction entering the pathway to social sustainability.

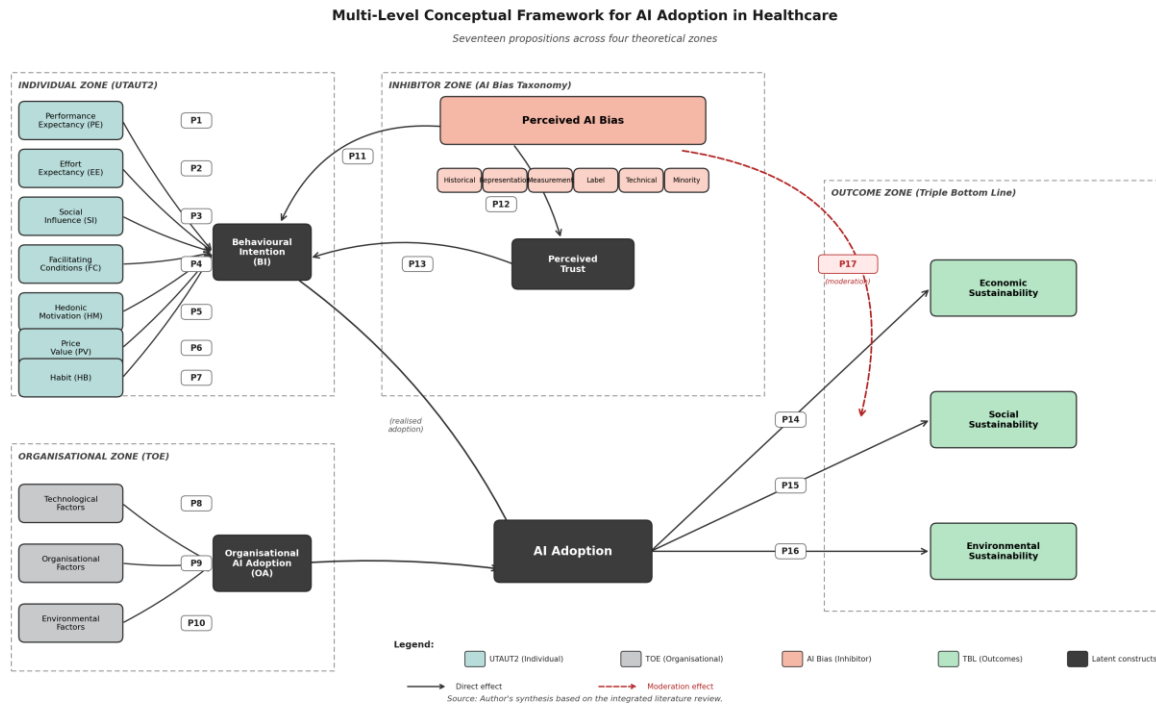


Figure 1: Multi-Level Conceptual Framework for AI Adoption in Healthcare. *Source: Author’s synthesis based on the integrated literature review.*

6.2 The seventeen propositions

The seventeen propositions are organised around the framework's four zones and are stated in directional form to facilitate future empirical testing. They are summarised in Table 4 and briefly justified in the text.

Table 4: Proposed Integrated Conceptual Framework — Seventeen Testable Propositions

P#	Proposition	Zone
P1	Performance expectancy is positively associated with behavioural intention to adopt AI	Individual
P2	Effort expectancy is positively associated with behavioural intention	Individual
P3	Social influence is positively associated with behavioural intention	Individual
P4	Facilitating conditions are positively associated with behavioural intention	Individual
P5	Hedonic motivation is positively associated with behavioural intention	Individual
P6	Price value is positively associated with behavioural intention	Individual
P7	Habit is positively associated with behavioural intention	Individual
P8	Technological factors positively influence organisational AI adoption	Organisational
P9	Organisational factors positively influence organisational AI adoption	Organisational
P10	Environmental factors positively influence organisational AI adoption	Organisational
P11	Perceived AI bias negatively affects behavioural intention to adopt AI	Inhibitor
P12	Perceived AI bias negatively affects perceived trust in AI systems	Inhibitor
P13	Perceived trust partially mediates the relationship between perceived AI bias and behavioural intention	Inhibitor
P14	AI adoption positively influences economic sustainability outcomes	Outcome
P15	AI adoption positively influences social sustainability outcomes	Outcome
P16	AI adoption positively influences environmental sustainability outcomes	Outcome
P17	Perceived AI bias moderates (weakens) the positive relationship between AI adoption and social sustainability outcomes	Outcome

Source: Developed by the author based on the integrated literature review.

6.3 Justification of the less-obvious propositions

Propositions P1 through P10 restate, in integrated form, relationships that each source framework already supports. Their inclusion is not itself novel; the novelty is their simultaneous specification within a single testable model (addressing Gap 1).

Propositions P11 through P13 position bias as a first-class adoption construct (addressing Gap 2). P11 and P12 are direct effects consistent with the converging evidence reviewed in Section 4.3. P13—the mediation proposition—reflects the theoretical argument that bias operates in significant part through the psychological mechanism of trust. Partial rather than full mediation is proposed because bias plausibly operates through additional mechanisms (such as perceived professional risk or ethical concern) that trust does not fully capture.

Propositions P14 through P16 extend the TBL framework into adoption outcome measurement (addressing Gap 3). They assert positive effects in all three dimensions, consistent with the emerging empirical evidence that AI adoption is, on net, beneficial. Whether these benefits are distributed equitably is a separate question—which P17 specifically addresses.

Proposition P17 is the most distinctive element of the framework. It states that perceived bias acts as a boundary condition on the relationship between adoption and social sustainability: the greater the perceived bias, the weaker the equity benefits of adoption. This captures a policy-relevant intuition—that AI deployed without attention to fairness can generate economic returns while leaving social-equity outcomes unimproved or even eroded—and turns it into a testable claim. Empirical tests of this specific moderation relationship have not, to the author's knowledge, been reported in the healthcare AI literature, and the framework is designed to make such tests tractable.

7. Discussion And Implications

The framework is designed to serve three audiences. For the research community, it provides a scaffold for empirical work that is more complete than single-theory designs allow. For institutional leaders, it provides a vocabulary for thinking about AI adoption that does not collapse the distinct concerns of clinicians, administrators, and governance bodies into a single concept. For policy-makers, it offers a way to connect adoption decisions to the distributional and environmental outcomes that regulators increasingly care about.

7.1 Theoretical implications

The framework's principal theoretical contribution is to make bias visible as an adoption construct. Existing adoption theory has handled bias awkwardly—sometimes absorbing it into risk, sometimes into trust, sometimes omitting it—because the adoption frameworks were developed before algorithmic fairness became a central concern. The framework proposed here treats bias as a distinct multi-dimensional construct, specifies its direct and indirect effects on intention, and further specifies its role as a boundary condition on equity outcomes. Empirical validation of this structure would mark a substantive advance over single-framework designs.

A second theoretical contribution is the reframing of outcomes. The framework makes explicit what the adoption literature has long implied but rarely stated: that adoption is not an end in itself but an intermediate step toward outcomes that can be evaluated across multiple dimensions. By committing to the TBL outcome space, the framework forces analyses to ask not only "does adoption happen?" but also "does it produce what was intended?"—and to answer the second question along economic, social, and environmental axes simultaneously.

7.2 Practical implications for institutional leaders

For hospital administrators and health-system executives, the framework suggests that AI-adoption decisions should be organised around four questions rather than one. The first is whether clinicians within the institution are likely to accept the tool; this is a UTAUT2 question and requires investment in the facilitating conditions that the reviewed literature consistently identifies as decisive. The second is whether the institution itself is ready; this is a TOE question and requires honest assessment of technological fit, organisational capacity, and environmental pressure. The third is whether the tool presents unresolved bias concerns; this is a fairness question and requires transparent vendor disclosures, demographic subgroup performance data, and, where available, independent validation (through platforms

such as BODH). The fourth is whether the tool is likely to deliver benefits across all three TBL dimensions; this is a sustainability question and requires institutional commitment to measuring what it has thus far avoided measuring.

These are not four separate workflows. They are four questions that a single procurement and governance process should be able to address simultaneously. The framework's value for institutional leaders lies precisely in preventing the functional separation that has become common in practice—where procurement asks the first question, IT asks the second, a fairness committee (if one exists) asks the third, and the fourth is not asked at all.

7.3 Policy implications

For policy-makers, the framework's most directly relevant component is P17, the moderation proposition. If empirical testing confirms that perceived bias materially weakens the social-equity benefits of AI adoption, the policy implication is that bias-audit and fairness-validation requirements are not regulatory impositions on top of adoption policy—they are integral to it. Policy frameworks that promote adoption without also requiring fairness demonstration may produce economic gains that are distributed inequitably, a pattern that the literature consistently identifies as socially and politically unsustainable. Initiatives such as SAHI and the BODH benchmarking platform in India are exactly the kind of infrastructure that the framework suggests will matter. Whether they matter in the magnitudes predicted is an empirical question the framework makes tractable.

A second policy implication concerns monitoring. The TBL outcome space implies that adoption monitoring should include equity and environmental metrics alongside efficiency metrics. Many healthcare systems already monitor equity outcomes in other contexts (through, for instance, quality-adjusted life-year analyses or disparities reporting); extending this monitoring to AI deployment is technically feasible and, on the framework's logic, normatively required.

8. Limitations of the Review

Three limitations of the review are worth stating plainly. First, it is a narrative rather than a systematic review in the PRISMA sense. The selection of studies involved judgement at several stages, and a fully systematic protocol would have permitted stronger claims about the completeness of coverage. Second, the review was conducted in English and relied on four databases; studies published in other languages or indexed in regional databases may have been missed. Third, the temporal cutoff of March 2026 means that any subsequent developments—and given the pace of the field, there will be many—are not reflected. The framework, however, is designed to be generative rather than final: it is meant to be tested, refined, and eventually superseded by better-specified models as the evidence accumulates.

9. Directions for Future Empirical Research

The framework generates a research agenda that is broader than any single empirical study could fulfil. The following directions seem most immediately tractable.

9.1 Large-sample structural equation tests of the full framework

The most direct empirical test is a cross-sectional survey of healthcare professionals with a sample large enough to estimate all seventeen propositions within a single SEM. A minimum sample of around 500 is advisable given the complexity of the model; larger samples would permit multi-group analysis across institution types and geographies. The Indian healthcare system, with its heterogeneity of institutional settings and its strategic importance, is a particularly rich context for this test.

9.2 Longitudinal designs

Cross-sectional tests establish structural relationships but cannot settle questions of causal direction. Longitudinal panel studies that track the same respondents across two or more time points would permit stronger causal inference, particularly for the trust-mediation pathway (P13) and the bias-moderation pathway (P17), both of which imply temporal dynamics that cross-sectional data can only approximate.

9.3 Application-specific refinement

The framework treats AI as a single broad category. Healthcare AI is not a single technology—it spans diagnostic imaging, clinical decision support, generative AI for documentation, administrative automation, and robotic surgery, each with its own adoption dynamics and bias profiles. Refining the framework for specific application classes would yield more actionable findings and would help identify where the general structure holds and where application-specific adaptations are required.

9.4 Pre- and post-BODH comparative studies

The launch of the BODH benchmarking platform in February 2026 creates an unusual opportunity for naturalistic comparison. Studies that compare adoption dynamics, trust levels, and perceived bias among clinicians using BODH-validated tools against those using unvalidated tools would provide direct empirical evidence on whether independent fairness benchmarking changes the adoption equation. If it does, the policy logic behind such platforms is confirmed; if it does not, alternative governance strategies will need to be considered.

9.5 Patient-centred extensions

The framework as presented is clinician- and institution-centred. Patients are an obvious missing constituency. Extending the framework to include patient trust, patient-perceived fairness, and patient acceptance of AI involvement in their care would close an important gap, particularly as healthcare AI becomes increasingly visible to patients themselves (through, for example, AI-assisted triage in telemedicine platforms and patient-facing chatbots).

9.6 Environmental-sustainability deep-dives

The environmental dimension of the TBL outcome space is, in the reviewed literature, the most weakly measured. Dedicated studies that quantify energy consumption, carbon emissions, and electronic-waste profiles associated with healthcare AI deployments would strengthen the empirical basis of P16 and would connect the framework to the broader green-AI agenda concerned with environmentally responsible model development and deployment.

10. Conclusion

The literature on artificial intelligence adoption in healthcare has produced considerable insight, but its yields have been limited by a tendency to study the phenomenon one lens at a time. Acceptance theorists have mapped individual intention without reliably connecting it to organisational context. Organisational scholars have described institutional readiness without systematically incorporating individual-level dynamics. Fairness researchers have documented algorithmic bias without translating their findings into adoption theory. Sustainability scholars have raised distributive and environmental concerns without embedding them in adoption models. The consequence is that questions that cut across these perspectives—how bias awareness affects clinician intention, whether adoption produces equitable outcomes, how the three dimensions of sustainability depend on one another and on bias—remain under-researched.

This review has synthesised the four literatures and proposed a consolidated framework of seventeen testable propositions that aims to close the integration gap. The framework's most distinctive feature is the positioning of algorithmic bias as a first-class adoption construct with direct, trust-mediated, and moderation effects. Its second distinctive feature is the commitment to TBL outcomes, which forces analyses to ask not only whether AI is adopted but whether adoption produces what it is supposed to produce, across economic, social, and environmental dimensions. Neither feature is, on its own, new to the individual literatures from which it is drawn. Their simultaneous incorporation into a single testable framework is novel.

The framework is offered not as a final account but as a scaffold for empirical work. Its value will be determined by how well it survives confrontation with data, and by whether the empirical tests it invites produce findings that are more actionable than those currently available from single-framework studies. The field is at a moment—the transition from AI pilots to enterprise deployment, the rollout of validation infrastructure such as BODH, the institutional uncertainty about how to govern adoption, where better conceptual tools have unusually high practical value. It is hoped that this contribution will be used, tested, and improved upon in the research and practice that follow.

DECLARATIONS

Ethical Approval: Not applicable. This article is a review of previously published literature and does not involve new research with human participants or animals. No primary data were collected.

Funding: This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Conflict of Interest: The author declares no conflict of interest.

Data Availability: All data underlying this review are drawn from publicly accessible peer-reviewed publications and policy documents cited in the reference list. No additional datasets were generated or analysed.

REFERENCES

1. Adithyan, R., et al. (2024). Knowledge and attitude towards artificial intelligence among healthcare professionals in a tertiary care hospital in India. *Journal of Clinical and Diagnostic Research*, 18(5), LC01–LC05.
2. Ammenwerth, E. (2019). Technology acceptance models in health informatics: TAM and UTAUT. *Studies in Health Technology and Informatics*, 263, 64–71.
3. Analytics Insight. (2026, February 18). India AI Impact Summit 2026: Centre sets AI healthcare guardrails with SAHI, BODH rollout. Analytics Insight.
4. Baker, J. (2012). The technology–organisation–environment framework. In Y. K. Dwivedi, M. R. Wade, & S. L. Schneberger (Eds.), *Information systems theory: Explaining and predicting our digital society*, Vol. 1 (pp. 231–245). Springer.
5. BCG (Boston Consulting Group). (2026). How digital and AI will reshape health care in 2026. BCG Publications.
6. Chen, R. J., Wang, J. J., Williamson, D. F. K., Chen, T. Y., Lipkova, J., Lu, M. Y., Sahai, S., & Mahmood, F. (2023). Algorithmic fairness in artificial intelligence for medicine and healthcare. *Nature Biomedical Engineering*, 7(6), 719–732. <https://doi.org/10.1038/s41551-023-01056-8>
7. Chettri, S. K., et al. (2025). Bridging the gap in the adoption of trustworthy AI in Indian healthcare: Challenges and opportunities. *AI (MDPI)*, 6(1), Article 10.
8. Chief Healthcare Executive. (2026, January 15). AI in health care: 26 leaders offer predictions for 2026. Chief Healthcare Executive Magazine.
9. Cobelli, N., & Blasioli, E. (2023). To be or not to be digital? A bibliometric analysis of adoption of eHealth services. *The TQM Journal*, 35(9), 299–331. <https://doi.org/10.1108/TQM-02-2023-0065>
10. Cross, J. L., Choma, M. A., & Onofrey, J. A. (2024). Bias in medical AI: Implications for clinical decision-making. *PLOS Digital Health*, 3(11), e0000651. <https://doi.org/10.1371/journal.pdig.0000651>
11. Dai, Q., Li, M., Yang, M., Shi, S., Wang, Z., Liao, J., Li, Z., E, W., Tao, L., & Tang, Y.-D. (2025). Attitudes, perceptions, and factors influencing the adoption of AI in health care among medical staff: Nationwide cross-sectional survey study. *Journal of Medical Internet Research*, 27, e75343. <https://doi.org/10.2196/75343>
12. Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, 13(3), 319–340. <https://doi.org/10.2307/249008>
13. Doximity. (2026). 2026 state of AI in medicine report. Doximity Inc.
14. Dwivedi, Y. K., et al. (2025). Generative AI organisational adoption: A mixed-methods study extending the TOE framework. *Information Systems Frontiers*, 27(2), 345–368.
15. Elkington, J. (1997). *Cannibals with forks: The triple bottom line of 21st century business*. Capstone Publishing.
16. Gichoya, J. W., Banerjee, I., Bhimireddy, A. R., Burns, J. L., Celi, L. A., Chen, L.-C., Correa, R., Dullerud, N., Ghassemi, M., Huang, S.-C., Kuo, P.-C., Lungren, M. P., Palmer, L. J., Price, B. J., Purkayastha, S., Pyrros, A. T., Oakden-Rayner, L., Okechukwu, C., Seyyed-Kalantari, L., ... Zhang, H. (2022). AI recognition of patient race in medical imaging: A modelling study. *The Lancet Digital Health*, 4(6), e406–e414. [https://doi.org/10.1016/S2589-7500\(22\)00063-2](https://doi.org/10.1016/S2589-7500(22)00063-2)
17. Grant, M. J., & Booth, A. (2009). A typology of reviews: An analysis of 14 review types and associated methodologies. *Health Information and Libraries Journal*, 26(2), 91–108. <https://doi.org/10.1111/j.1471-1842.2009.00848.x>
18. Hasanzadeh, F., et al. (2025). Bias in health AI: Sources, mitigation strategies, and ethical considerations. *AI and Ethics*, 5(1), 89–112.
19. Hassan, M. K., et al. (2026). AI adoption in healthcare organizations: Spheres of development and the virtue of visible value. *Technological Forecasting and Social Change*. Advance online publication. <https://doi.org/10.1016/j.techfore.2026.123967>
20. Kumar, A., et al. (2024). Representational bias in global health AI datasets: Implications for Indian healthcare. *BMJ Global Health*, 9(3), e014231.

21. Lambert, S. I., Madi, M., Sopka, S., Lenes, A., Stange, H., Buszello, C.-P., & Stephan, A. (2023). An integrative review on the acceptance of artificial intelligence among healthcare professionals in hospitals. *npj Digital Medicine*, 6(1), 111. <https://doi.org/10.1038/s41746-023-00852-5>
22. NVIDIA. (2026). State of AI in healthcare and life sciences 2026. NVIDIA Corporation.
23. Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464), 447–453. <https://doi.org/10.1126/science.aax2342>
24. PIB (Press Information Bureau). (2026, February 17). Union Minister launches SAHI and BODH initiatives at India AI Impact Summit 2026. Government of India.
25. Pinto, A. S., Abreu, A., Pérez Cota, M., et al. (2025). A meta-analysis of TOE factors driving organizational adoption of artificial intelligence across industries. *Discover Artificial Intelligence*, 5, Article 36. <https://doi.org/10.1007/s44163-025-00747-2>
26. Pumplun, L., Fecho, M., Wahl, N., Peters, F., & Buxmann, P. (2021). Adoption of machine learning systems for medical diagnostics in clinics: Qualitative interview study. *Journal of Medical Internet Research*, 23(10), e29301. <https://doi.org/10.2196/29301>
27. Rajkomar, A., Hardt, M., Howell, M. D., Corrado, G., & Chin, M. H. (2018). Ensuring fairness in machine learning to advance health equity. *Annals of Internal Medicine*, 169(12), 866–872. <https://doi.org/10.7326/M18-1990>
28. Salah, O. H., & Ayyash, M. M. (2024). A systematic review of TOE framework in AI adoption research. *Information Technology & People*, 37(4), 1456–1489.
29. Schiavone, F., Mancini, D., Leone, D., & Lavorato, D. (2022). Digital business models and ridesharing for value co-creation in healthcare. *Technological Forecasting and Social Change*, 176, 121462.
30. Siddique, S., et al. (2024). Accuracy disparities of AI diagnostic models across minority populations. *Journal of Health Equity*, 8(2), 156–171.
31. Sjoding, M. W., Dickson, R. P., Iwashyna, T. J., Gay, S. E., & Valley, T. S. (2020). Racial bias in pulse oximetry measurement. *New England Journal of Medicine*, 383(25), 2477–2478. <https://doi.org/10.1056/NEJMc2029240>
32. Su, Y., et al. (2025). Factors influencing adoption of AI health assistants: An extended UTAUT model. *Scientific Reports*, 15(1), Article 2345.
33. Tamilmani, K., Rana, N. P., & Dwivedi, Y. K. (2021). Consumer acceptance and use of information technology: A meta-analytic evaluation of UTAUT2. *Information Systems Frontiers*, 23(4), 987–1005. <https://doi.org/10.1007/s10796-020-10007-6>
34. Tamilmani, K., Rana, N. P., Wamba, S. F., & Dwivedi, R. (2021). The extended Unified Theory of Acceptance and Use of Technology (UTAUT2): A systematic literature review and theory evaluation. *International Journal of Information Management*, 57, Article 102269. <https://doi.org/10.1016/j.ijinfomgt.2020.102269>
35. Thieme, A., et al. (2025). Surgeon adoption of AI surgical technology: A mixed methods study. *International Journal of Medical Informatics*, 189, Article 105467.
36. Tornatzky, L. G., Fleischer, M., & Chakrabarti, A. K. (1990). *The processes of technological innovation*. Lexington Books.
37. Ueda, D., et al. (2024). Environmental sustainability dimensions of AI in healthcare. *Sustainability*, 16(5), Article 2048.
38. Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. *MIS Quarterly*, 27(3), 425–478. <https://doi.org/10.2307/30036540>
39. Venkatesh, V., Thong, J. Y. L., & Xu, X. (2012). Consumer acceptance and use of information technology: Extending the unified theory of acceptance and use of technology. *MIS Quarterly*, 36(1), 157–178. <https://doi.org/10.2307/41410412>
40. Wolters Kluwer. (2026). 2026 healthcare AI trends: Insights from experts. Wolters Kluwer Health.
41. Xie, Y., et al. (2025). A bibliometric analysis of artificial intelligence in healthcare: Thirty years of research. *Artificial Intelligence in Medicine*, 148, Article 102757.