



International Journal of Research Publication and Reviews

Journal Homepage: www.ijrpr.com ISSN 2582-7421

A Study on Ethical and Regulatory Challenges of AI Adoption in Healthcare

Ch. Swathi ¹, Dr. Y. Venkata Ramana ²

¹Research Scholar, Department of Business Administration, KL Business School, Koneru Lakshmaiah Education Foundation, Guntur, AP, India.

Email: choddisettiswathi@gmail.com

² Senior Assistant Professor, Department of MBA, KL Business School, Koneru Lakshmaiah Education Foundation vaddeswaram.

E-Mail: dryvr2224@gmail.com

ABSTRACT

Artificial Intelligence (AI) can improve diagnostics, predictive analytics, and clinical decision support, yet its adoption in healthcare is constrained by unresolved ethical, legal, and regulatory challenges, including algorithmic bias, limited transparency, unclear accountability, and weak data-privacy safeguards. This study examines how ethical challenges, regulatory challenges, and governance factors relate to responsible AI adoption in healthcare. A descriptive, quantitative survey was conducted using a structured questionnaire; after removing repeated submissions, 22 valid responses from students and professionals in healthcare, technology/IT, and other fields were analysed using descriptive statistics, Cronbach's alpha, multiple linear regression, a chi-square test, and one-way ANOVA.

Respondents reported uniformly favourable views (construct means between 4.49 and 4.67 on a 5-point scale). The regression model was statistically significant ($F(3, 18) = 57.62, p < .001$, adjusted $R^2 = .890$). Ethical challenges ($\beta = .40, p = .001$) and governance factors ($\beta = .41, p = .027$) were significant positive predictors of responsible AI adoption, whereas regulatory challenges were not ($p = .211$), partly because they overlap strongly with governance factors. Responsible AI adoption did not differ significantly by respondent status ($\chi^2(2) = 0.041, p = .980$) or by field group ($F(2, 19) = 0.41, p = .667$). Most respondents named data privacy as the main challenge (68.2%) and attributed responsibility for AI errors mainly to AI developers (63.6%).

Because the sample is small, convenience-based, and predominantly young students, the findings are exploratory. The study nevertheless contributes an initial governance-oriented picture of what stakeholders expect from responsible AI adoption and identifies priorities for healthcare managers, policymakers, and regulators.

Keywords: Artificial Intelligence, Healthcare, Ethical Challenges, Regulatory Compliance, AI Governance, Patient Trust, Responsible AI Adoption

1. Introduction

Artificial Intelligence (AI) is rapidly transforming the global healthcare landscape, offering unprecedented capabilities in diagnostics, treatment planning, predictive analytics, drug discovery, patient monitoring, and administrative efficiency. Technologies such as machine learning algorithms, natural language processing, and AI-driven decision-support systems are increasingly being integrated into clinical workflows, enabling healthcare providers to detect diseases earlier, personalize treatment, and optimize resource allocation. The promise of AI in healthcare is substantial: improved diagnostic accuracy, reduced clinical errors, streamlined administrative processes, and ultimately, better patient outcomes at lower costs.

Despite this potential, the adoption of AI in healthcare has not progressed as rapidly or as smoothly as in other industries. A central reason for this lag is the persistence of unresolved ethical, legal, and regulatory challenges surrounding AI's use in clinical settings. Unlike many other applications of AI, healthcare decisions directly affect human life and wellbeing, making questions of accuracy, fairness, transparency, and accountability especially critical. Concerns around algorithmic bias, where AI systems trained on non-representative data may produce inequitable outcomes for certain patient groups, raise serious ethical questions about fairness and justice in care delivery. Similarly, the "black box" nature of many AI models creates challenges around transparency and explainability, making it difficult for clinicians and patients to understand how a given AI-driven recommendation was reached.

Compounding these ethical concerns is a fragmented and, in many jurisdictions, underdeveloped regulatory environment. Existing legal frameworks were largely designed before the emergence of AI and often fail to adequately address issues such as liability when an AI system contributes to a clinical error, data privacy and security in the handling of sensitive patient information, and the certification standards required to validate AI tools for clinical use. The absence of standardized, universally accepted frameworks for ensuring AI safety, transparency, accountability, and data protection leaves healthcare organizations to navigate a complex and inconsistent landscape largely on their own.

Beyond ethical and regulatory concerns, the successful adoption of AI in healthcare also depends heavily on organizational governance and managerial factors. Leadership commitment, clear internal policies, risk-management structures, staff training, and change-management practices all play a decisive role in determining whether an organization can implement AI responsibly rather than merely rapidly. Yet much of the existing literature addresses AI ethics and regulation at a conceptual or policy level, with comparatively little empirical research examining how these managerial and governance factors function in real-world healthcare settings. Furthermore, patient trust, a critical determinant of whether AI-assisted care is accepted and effective, remains insufficiently understood, particularly regarding how transparency and accountability mechanisms shape patients' willingness to trust AI-supported clinical decisions.

Given these gaps, there is a clear need for research that moves beyond identifying ethical and regulatory challenges in the abstract and instead examines the managerial and governance mechanisms that can enable healthcare organizations to adopt AI responsibly and effectively. This study addresses that need by exploring how ethical challenges, regulatory and legal gaps, and organizational governance factors relate to responsible AI adoption in healthcare, and by considering the role of AI safety, transparency, and patient trust in this process. In doing so, it aims to contribute to the academic literature on AI governance in healthcare and to practical guidance for healthcare managers, policymakers, and regulators.

The remainder of this paper is organized as follows: Section 2 reviews the literature; Section 3 presents the research gap, conceptual framework, objectives, and hypotheses; Section 4 describes the research methodology; Section 5 reports the data analysis and results; Section 6 discusses the findings and their implications; and Section 7 concludes with limitations and directions for future research.

2. Review of Literature

Artificial Intelligence in Healthcare Delivery (2024)

This study maps the main application types of AI used in healthcare delivery, including medical imaging analysis, predictive analytics, charting and chatbot tools, and AI-driven robotics, and identifies the implementation challenges associated with each, chiefly data quality, algorithmic bias, interpretability, regulatory uncertainty, and cost.

Artificial Intelligence in Healthcare (2021)

This paper describes the stages through which AI systems are built and deployed in healthcare (design and development, evaluation and validation, diffusion and scaling, and ongoing monitoring and maintenance) and maps these stages against major use-case domains such as connected and augmented care, precision diagnostics, precision therapeutics, and precision medicine, across short-, medium-, and long-term adoption timelines.

Three Epochs of Artificial Intelligence in Health Care (2024)

This study classifies the evolution of AI in healthcare into three epochs (symbolic AI, deep learning, and foundation models), distinguishing them by core functionality and training method, and illustrates the differing performance capabilities and risks associated with each epoch in clinical use.

Artificial Intelligence in Healthcare: Past, Present and Future (2017)

This review surveys the major types of healthcare data used to train AI systems (imaging, genetic, electrophysiological, and electronic medical record data) and categorizes AI devices into machine learning and natural language processing approaches. It further reviews specific algorithms, including support vector machines, neural networks, and deep learning/convolutional neural networks, as applied to cancer, neurology, cardiovascular disease, and stroke.

Artificial Intelligence in Healthcare (2025)

This paper clarifies the hierarchical relationship between artificial intelligence, machine learning, and deep learning, and surveys the breadth of AI applications in healthcare, spanning medical imaging, drug discovery, personalized medicine, predictive analytics, patient monitoring, robotic surgery, disease-outbreak prediction, electronic health record management, telemedicine, clinical decision support, and chatbots, while reporting associated accuracy, sensitivity, and specificity outcomes.

Synthesis

Taken together, the reviewed literature shows that AI's technical capability in healthcare, across imaging, diagnostics, predictive analytics, and clinical decision support, is well ahead of the ethical, legal, and organizational infrastructure needed to deploy it responsibly. Each of the studies reviewed identifies pieces of this gap (data bias, interpretability, regulatory uncertainty, cost, and evolving model risk) but stops short of empirically testing how ethical challenges, regulatory gaps, and organizational governance factors jointly determine whether an organization adopts AI responsibly. This is the gap the present study addresses.

3. Research Gap, Conceptual Framework, Objectives and Hypotheses

3.1 Research Gap

Although artificial intelligence has significant potential to improve healthcare, its adoption is limited by unresolved ethical, legal, and regulatory challenges. There is also a lack of standardized frameworks for ensuring AI safety, transparency, accountability, and data privacy in clinical practice. Limited real-world implementation studies and insufficient understanding of patient trust further hinder successful adoption. Most existing literature treats these ethical, regulatory, and governance dimensions descriptively or in isolation, without empirically testing how they jointly influence responsible AI adoption. Therefore, this study aims to explore the managerial and governance factors, alongside ethical and regulatory challenges, that can support the responsible and effective implementation of AI in healthcare.

3.2 Conceptual Framework

Drawing on the ethical, regulatory, and governance dimensions identified above, the study proposes a framework in which three independent variables (ethical challenges, regulatory challenges, and governance factors) influence the dependent variable, responsible AI adoption, both directly and indirectly through a mediating variable, AI safety and patient trust.




Figure 1. Conceptual framework showing independent, mediating, and dependent variables

- **Independent variables:** Ethical Challenges (bias, consent, accountability); Regulatory Challenges (privacy laws, compliance gaps); Governance Factors (leadership, policy, training).
- **Mediating variable:** AI Safety & Patient Trust (transparency and accountability mechanisms).
- **Dependent variable:** Responsible AI Adoption (compliance, outcomes, patient trust).
- A direct effect of the independent variables on the dependent variable is also modeled alongside the mediated path, consistent with a partial-mediation specification.

3.3 Research Objectives

General objective. To explore the managerial and governance factors that support the responsible and effective implementation of AI in healthcare, addressing the ethical, legal, and regulatory challenges that currently limit its adoption.

Specific objectives:

1. To identify the key ethical challenges (bias, informed consent, clinical accountability) affecting AI adoption in healthcare organizations.
2. To examine the legal and regulatory gaps, including the absence of standardized frameworks for AI safety, transparency, accountability, and data privacy, that hinder responsible AI adoption.
3. To assess the managerial and governance practices (leadership support, organizational policies, risk-management structures) that enable or constrain responsible AI implementation.

4. To evaluate the extent of patient trust in AI-assisted healthcare and how it influences adoption outcomes.
5. To examine whether AI adoption readiness differs across respondent status (student, professional, other) and field group (healthcare, technology/IT, other).
6. To recommend a governance-based framework that healthcare organizations can use to adopt AI responsibly and effectively.

3.4 Research Hypotheses

H1: Ethical Challenges → Responsible AI Adoption

Null hypothesis (H0): Ethical challenges have no significant effect on responsible AI adoption in healthcare.

Alternative hypothesis (Ha): Ethical challenges have a significant effect on responsible AI adoption in healthcare.

H2: Regulatory Challenges → Responsible AI Adoption

Null hypothesis (H0): Regulatory challenges have no significant effect on responsible AI adoption in healthcare.

Alternative hypothesis (Ha): Regulatory challenges have a significant effect on responsible AI adoption in healthcare.

H3: Governance Factors → Responsible AI Adoption

Null hypothesis (H0): Governance factors have no significant effect on responsible AI adoption in healthcare.

Alternative hypothesis (Ha): Governance factors have a significant effect on responsible AI adoption in healthcare.

H4: Respondent Status × AI Adoption Readiness Level (Chi-Square)

Null hypothesis (H0): There is no significant association between respondent status (student / professional / other) and AI adoption readiness level (low / medium / high).

Alternative hypothesis (Ha): There is a significant association between respondent status (student / professional / other) and AI adoption readiness level (low / medium / high).

H5: Field Group and Responsible AI Adoption (ANOVA)

Null hypothesis (H0): There is no significant difference in mean responsible AI adoption scores across field groups (healthcare, technology/IT, other).

Alternative hypothesis (Ha): At least one field group differs significantly in mean responsible AI adoption score.

4. Research Methodology

4.1 Research Design

The study adopts a descriptive and quantitative research design based on a structured questionnaire survey, aimed at empirically testing the relationships specified in the conceptual framework between ethical challenges, regulatory challenges, governance factors, AI safety and patient trust, and responsible AI adoption.

4.2 Population, Sampling and Data Cleaning

The intended population comprised people with a stake in healthcare AI: healthcare professionals, technology/IT staff, and students preparing for these fields. Respondents were reached through convenience sampling, and responses were collected on 28–29 September 2026. The exported data file contained 85 rows. Sixty-three of these were exact repeats of earlier submissions (identical timestamps and identical answers) and were removed, leaving 22 unique valid responses ($n = 22$), on which all analyses are based. Because the sample is small and convenience-based, results are treated as exploratory.

4.3 Data Collection Tool and Coding

Data were collected with a structured questionnaire containing four demographic items (age group, gender, education, and occupation) and 21 substantive items. Seventeen of the substantive items were converted to a common 5-point scale (1 = lowest agreement, 5 = highest agreement) for analysis: rating items and agreement statements were used as recorded (Strongly disagree = 1 to Strongly agree = 5), and categorical answers were recoded as shown in Table 1. Four items (two multi-select items and two single-choice items) are not ordinal and are reported as frequencies only.

Table 1. Coding scheme for categorical responses

Score	Response category	Example responses
5	Unconditional agreement	"Yes"; "Yes, always"; "Strongly agree"
4	Agreement with a condition	"Yes, in most cases"; "Only if properly designed and tested"; "Only under human supervision"; "Agree"
3	Midpoint / uncertain	"Partially"; "Not sure"; "Neutral"
2	Disagreement	"Disagree"
1	Rejection	"No"; "Strongly disagree"

Note. A categorical "No" was treated as the lowest anchor (1). This is an analytic assumption.

4.4 Variables and Measurement

Scale scores were computed as the mean of the items assigned to each construct (Table 2), so that every construct is expressed on the original 1–5 metric.

Table 2. Measurement of study constructs

Construct	Role	Items	No. of items
Ethical Challenges (EC)	Independent	Q9, Q10, Q13	3
Regulatory Challenges (RC)	Independent	Q14, Q15, Q17, Q18, Q19	5
Governance Factors (GF)	Independent	Q6, Q21	2
AI Safety & Patient Trust (ST)	Mediating	Q11, Q12	2
Responsible AI Adoption (RA)	Dependent	Q1, Q2, Q4, Q5, Q7	5

[AUTHOR NOTE – delete before submission: the source file had no header row, so this item-to-construct grouping follows the order of items in the questionnaire. Please check it against the original questionnaire and replace the item codes with the actual item wording. If the grouping changes, all results below must be re-run, and the Discussion wording that describes each construct (e.g. Governance as human oversight, Ethical Challenges as consent/accountability) must be revised too.]

4.5 Statistical Tools

- Descriptive statistics (frequency, percentage, mean, standard deviation) to profile respondents and summarize each construct.
- Cronbach's alpha to assess the internal consistency of each construct scale, and Pearson correlations to describe relationships among constructs.
- Multiple linear regression to test the effect of ethical challenges, regulatory challenges, and governance factors on responsible AI adoption (H1–H3), with variance inflation factors (VIF) to check multicollinearity.
- Chi-square test of independence to test the association between respondent status and AI adoption readiness level (H4). Readiness level was defined by equal-width bands of the responsible AI adoption score: Low (1.00–2.33), Medium (2.34–3.67), and High (3.68–5.00).
- One-way ANOVA to test differences in mean responsible AI adoption scores across field groups (H5), with a Kruskal–Wallis test as a robustness check given the skewed scores and small groups.

Calculations were carried out in Microsoft Excel using worksheet formulas and cross-checked with independent calculations in Python. The significance level was set at $\alpha = .05$.

Note on variables: the original design specified organization type (public/private) and job role (clinician/administrator/IT staff) as grouping variables. The questionnaire did not collect organization type, and its occupation item distinguishes students from professionals in healthcare and technology/IT. H4 and H5 were therefore tested using respondent status (student, professional, other) and field group (healthcare, technology/IT, other), respectively.

5. Data Analysis and Results

5.1 Profile of Respondents

Table 3 summarizes the 22 respondents. The sample is young and student-heavy: 90.9% were aged 18–25, and 54.5% were students. 13 (59.1%) were female and 9 (40.9%) male. Most held an undergraduate qualification (72.7%). By field, 9 (40.9%) were healthcare-related, 8 (36.4%) technology/IT-related, and 5 (22.7%) from other backgrounds.

Table 3. Sample profile ($n = 22$)

Variable	Category	n	%
Age group	Below 18	1	4.5%
	18–25	20	90.9%
	26–35	1	4.5%
Gender	Female	13	59.1%
	Male	9	40.9%
Education	High school	1	4.5%
	Diploma	3	13.6%
	Undergraduate	16	72.7%
	Doctorate	2	9.1%
Occupation	Healthcare student	5	22.7%
	Healthcare professional	4	18.2%
	Technology/IT student	7	31.8%
	Technology/IT professional	1	4.5%
	Other	5	22.7%

Note. Percentages are of the 22 valid responses and may not sum to 100 because of rounding.

5.2 Descriptive Statistics and Reliability

All constructs were rated highly (Table 4), with means between 4.49 and 4.67. Scores were bunched near the top of the scale (a ceiling effect), so variability is limited. Internal consistency was good for Ethical Challenges ($\alpha = .81$) and Responsible AI Adoption ($\alpha = .88$), acceptable for Regulatory Challenges ($\alpha = .75$), questionable for AI Safety & Patient Trust ($\alpha = .61$), and poor for Governance Factors ($\alpha = .38$). The two-item Governance and Safety/Trust scales should therefore be interpreted with caution.

Table 4. Descriptive statistics and reliability of constructs

Construct	Items	Mean	SD	Range	Cronbach's α
Ethical Challenges	3	4.67	0.58	3.33–5.00	.81
Regulatory Challenges	5	4.62	0.57	3.00–5.00	.75
Governance Factors	2	4.61	0.84	2.00–5.00	.38

Construct	Items	Mean	SD	Range	Cronbach's α
AI Safety & Patient Trust	2	4.64	0.64	2.50–5.00	.61
Responsible AI Adoption	5	4.49	0.83	2.40–5.00	.88

Note. Scale: 1 = lowest, 5 = highest agreement.

5.3 Accountability and Main Challenge

Two single-choice items address Objectives 1 and 2 directly. Asked who should bear responsibility when AI contributes to an error, 63.6% of respondents ($n = 14$) pointed to the AI developer or company, 22.7% to shared responsibility, 9.1% to the hospital or healthcare organization, and 4.5% to the doctor or healthcare professional. Asked to name the main challenge for AI in healthcare, 68.2% ($n = 15$) chose data privacy, well ahead of cybersecurity (9.1%) and lack of regulation (9.1%); algorithmic bias, lack of professional training, and other reasons were each chosen by one respondent.

5.4 Correlations

All constructs were positively and strongly correlated (Table 5). Responsible AI Adoption correlated most strongly with Governance Factors ($r = .89$) and Regulatory Challenges ($r = .88$). Regulatory Challenges and Governance Factors were themselves very highly correlated ($r = .90$), which has consequences for the regression below.

Table 5. Pearson correlations among constructs ($n = 22$)

	EC	RC	GF	ST	RA
EC	–				
RC	.71	–			
GF	.69	.90	–		
ST	.68	.77	.76	–	
RA	.85	.88	.89	.70	–

Note. EC = Ethical Challenges; RC = Regulatory Challenges; GF = Governance Factors; ST = AI Safety & Patient Trust; RA = Responsible AI Adoption.

5.5 Regression Analysis (H1–H3)

Responsible AI Adoption was regressed on Ethical Challenges, Regulatory Challenges, and Governance Factors (Table 6). The model was statistically significant, $F(3, 18) = 57.62$, $p < .001$, and explained 90.6% of the variance ($R^2 = .906$; adjusted $R^2 = .890$).

- **H1 (supported).** Ethical Challenges had a significant positive effect ($B = 0.569$, $\beta = .40$, $t = 3.86$, $p = .001$), so the null hypothesis is rejected.
- **H2 (not supported).** Regulatory Challenges were not a significant predictor ($B = 0.325$, $\beta = .23$, $t = 1.30$, $p = .211$); the null hypothesis is not rejected.
- **H3 (supported).** Governance Factors had a significant positive effect ($B = 0.402$, $\beta = .41$, $t = 2.40$, $p = .027$), so the null hypothesis is rejected.

Regulatory Challenges and Governance Factors were closely related ($VIF = 5.77$ and 5.60 , below the conventional cut-off of 10 but above 5), so the non-significant result for regulation reflects overlap with governance and should not be read as evidence that regulation is unimportant.

Table 6. Multiple regression predicting Responsible AI Adoption ($n = 22$)

Predictor	B	SE	β	t	p	95% CI for B	VIF
Intercept	-1.523	0.601	–	-2.53	.021	[-2.79, -0.26]	–
Ethical Challenges	0.569	0.148	.40	3.86	.001	[0.26, 0.88]	2.07

Predictor	B	SE	β	t	p	95% CI for B	VIF
Regulatory Challenges	0.325	0.251	.23	1.30	.211	[-0.20, 0.85]	5.77
Governance Factors	0.402	0.167	.41	2.40	.027	[0.05, 0.75]	5.60

Note. $R^2 = .906$, adjusted $R^2 = .890$, $F(3, 18) = 57.62$, $p < .001$; standard error of estimate = 0.274. Dependent variable: Responsible AI Adoption.

5.6 Chi-Square Test (H_4)

Table 7 cross-tabulates respondent status with AI adoption readiness level. No respondent fell in the Low band, and 4 of 22 were in the Medium band. The association between status and readiness level was not significant, $\chi^2(2, N = 22) = 0.041$, $p = .980$, Cramér's $V = .04$. The null hypothesis is therefore not rejected. This test should be treated as indicative only: 5 of the 6 usable cells had expected counts below 5, so the chi-square approximation is weak. A permutation test led to the same conclusion.

Table 7. Respondent status by AI adoption readiness level, observed (expected) counts

Status	Low	Medium	High	Total
Student	0 (0.0)	2 (2.2)	10 (9.8)	12
Professional	0 (0.0)	1 (0.9)	4 (4.1)	5
Other	0 (0.0)	1 (0.9)	4 (4.1)	5
Total	0	4	18	22

5.7 One-Way ANOVA (H_5)

Mean Responsible AI Adoption scores were highest among healthcare-related respondents ($M = 4.69$, $SD = 0.86$), followed by the Other group ($M = 4.36$, $SD = 1.04$) and technology/IT-related respondents ($M = 4.35$, $SD = 0.70$). These differences were not statistically significant, $F(2, 19) = 0.41$, $p = .667$, $\eta^2 = .04$, so the null hypothesis is not rejected. A Kruskal–Wallis test gave the same result, $H(2) = 1.30$, $p = .523$. With only 5–9 respondents per group, the test could detect only large differences.

Table 8. Responsible AI Adoption by field group, and ANOVA summary

Field group	n	Mean	SD
Healthcare	9	4.69	0.86
Technology/IT	8	4.35	0.70
Other	5	4.36	1.04
Total	22	4.49	0.83

Source	SS	df	MS	F	p
Between groups	0.597	2	0.299	0.414	.667
Within groups	13.701	19	0.721		
Total	14.298	21			

5.8 Summary of Hypothesis Tests

Table 9. Summary of hypothesis tests ($\alpha = .05$)

Hypothesis	Test	Result	Decision
H1: Ethical Challenges $\rightarrow$ RA	Regression	$\beta = .40, p = .001$	Supported
H2: Regulatory Challenges $\rightarrow$ RA	Regression	$\beta = .23, p = .211$	Not supported
H3: Governance Factors $\rightarrow$ RA	Regression	$\beta = .41, p = .027$	Supported
H4: Status $\times$ readiness level	Chi-square	$\chi^2(2) = 0.041, p = .980$	Not supported
H5: Field group differences in RA	One-way ANOVA	$F(2, 19) = 0.41, p = .667$	Not supported

Note. RA = Responsible AI Adoption. "Supported" means the null hypothesis was rejected.

6. Discussion and Implications

6.1 Discussion of Findings

The results point to a sample that is broadly positive about AI in healthcare but attaches strong conditions to it. The significant, positive link between Ethical Challenges and Responsible AI Adoption (H1) should be read as a link between attitudes rather than as a barrier: respondents who placed more weight on consent, accountability, and related ethical safeguards also scored higher on awareness, perceived value, and acceptance of AI. This fits the idea that support for AI is conditional on safeguards being in place, and it is consistent with the literature reviewed in Section 2, which treats ethical and organizational infrastructure as the missing precondition for responsible deployment.

Governance Factors were also a significant predictor (H3), which supports the study's managerial focus: people who favoured human oversight of AI decisions were also more ready to see AI adopted. Regulatory Challenges (H2) were not significant once governance and ethics were in the model. Because the two constructs are almost collinear ($r = .90$), the data cannot separate their effects, and a larger sample or better-differentiated scales would be needed to do so.

Readiness did not differ significantly by respondent status (H4) or by field group (H5). The most plausible reading is consensus: almost everyone scored high, leaving little between-group variation to detect, together with very low statistical power at $n = 22$. Healthcare-related respondents scored highest, but the difference is small ($\eta^2 = .04$) and not reliable.

The descriptive items add practical detail. Data privacy was overwhelmingly named as the main challenge, and most respondents placed responsibility for AI errors on developers rather than on clinicians or hospitals. The second finding matters for governance: if the people closest to the technology expect developers to carry the liability, healthcare organizations still need clear contractual and internal rules on who answers for what, because legal liability is unlikely to fall only on developers in practice.

6.2 Implications for Practice and Policy

Table 10 links the main findings to governance actions. These are suggestions derived from a small exploratory sample and from the literature; they need validation with clinicians, administrators, and patients before being adopted as a framework.

Table 10. Findings and proposed governance actions

Finding	Proposed governance action
Data privacy named as main challenge (68.2%)	Make privacy and cybersecurity controls a precondition for approving any clinical AI tool; require data-protection impact assessments.
Ethical Challenges strongest predictor ($\beta = .40$)	Embed informed-consent and disclosure rules, bias testing on local patient data, and clear escalation paths for clinicians who disagree with AI output.
Governance Factors significant ($\beta = .41$)	Require human oversight for AI-supported decisions, define leadership ownership, and set up a standing AI risk committee with staff training.

Finding	Proposed governance action
Responsibility placed mainly on developers (63.6%)	Allocate accountability in writing across developer, hospital, and clinician; align vendor contracts with internal policy.
Regulation not separable from governance in the data	Regulators should give risk-tiered, practical compliance standards that organizations can translate into internal governance.

7. Conclusion, Limitations and Future Research

7.1 Conclusion

This study examined how ethical challenges, regulatory challenges, and governance factors relate to responsible AI adoption in healthcare. In this exploratory sample, ethical concerns and governance factors were significant positive predictors of responsible adoption, regulatory challenges were not once the other two were controlled, and readiness did not differ by respondent status or field group. Respondents were highly favourable towards AI but wanted safeguards, especially for data privacy, and expected developers to be held accountable. The findings support the view that responsible AI adoption depends as much on organizational governance and ethical safeguards as on technical performance.

7.2 Limitations

- Sample size and representativeness. After removing repeated submissions, only 22 responses remained, gathered by convenience sampling. Most respondents were students aged 18–25, not the practising clinicians and administrators originally targeted, so results cannot be generalized. One respondent was under 18, which should be checked against the ethics and consent requirements of the study.
- Measurement. Some constructs rest on few items (Governance: 2, Safety & Trust: 2) with low reliability, item-to-construct assignment was made after data collection, and categorical answers were converted to a 1–5 scale by assumption.
- Ceiling effects and overlap. Most answers were 4 or 5, and several rating items were nearly identical across respondents, which limits variation and inflates R^2 and inter-construct correlations.
- Design. The data are cross-sectional and self-reported, so the results show associations, not causal effects. The mediating role of AI Safety & Patient Trust in the framework was not formally tested. Organization type was not collected.
- Statistical power. With $n = 22$ the tests can detect only large effects, and the chi-square test violated expected-count assumptions.

7.3 Future Research

Future work should (i) survey a larger, stratified sample of practising clinicians, hospital administrators, and IT/AI-governance staff across public and private organizations; (ii) use validated multi-item scales with at least three or four items per construct; (iii) test the mediating path of AI safety and patient trust using bootstrapped mediation or structural equation modelling; and (iv) include patients, whose trust is central to the framework but was not directly measured here.

References

[To be completed. The draft lists five studies in the Review of Literature by title and year only. Add the full author, journal, volume, and page details for those studies and any other works cited, in the citation style required by your department.]