

# Translating AI Ethics into Hospital Operations: A PPTO Framework for Evidence-Based Governance

Vernette Grant<sup>1\*</sup>, Robert E. Levasseur<sup>2</sup>

<sup>1</sup>Independent Researcher, Aventura, FL, USA

<sup>2</sup>College of Management, Walden University, Minneapolis, MN, USA

Email: \*vernetgrant1@gmail.com, robert.levasseur@mail.waldenu.edu

**How to cite this paper:** Grant, V., & Levasseur, R. E. (2026). Translating AI Ethics into Hospital Operations: A PPTO Framework for Evidence-Based Governance. *Open Journal of Business and Management*, 14, 967-979.

<https://doi.org/10.4236/ojbm.2026.142056>

**Received:** February 3, 2026

**Accepted:** March 2, 2026

**Published:** March 5, 2026

Copyright © 2026 by author(s) and Scientific Research Publishing Inc. This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

<http://creativecommons.org/licenses/by/4.0/>



Open Access

---

## Abstract

Healthcare organizations face mounting pressure to adopt artificial intelligence (AI) responsibly despite ethical risks that include algorithmic bias, opacity, and inequitable care delivery. While high-level AI ethics principles abound, hospitals lack practical frameworks to operationalize these principles into governance structures with measurable outcomes. This paper advances the People-Process-Technology-Operations (PPTO) framework specifically for ethical AI governance in hospital settings, extending prior organizational applications to address unique challenges of translating abstract ethical principles into concrete operational practices. Drawing on recent governance case studies, implementation protocols, and scoping reviews of AI ethics frameworks in healthcare, this work provides hospital leaders with actionable guidance including role definitions, lifecycle review workflows, monitoring mechanisms, and performance metrics aligned with ethical commitments. The framework emphasizes scalability across varying organizational maturity levels and resource constraints while integrating with existing clinical governance structures. By systematically connecting governance activities to operational outcomes, including equity indicators, safety metrics, and stakeholder trust measures, PPTO transforms ethical AI from compliance burden into strategic advantage. Hospital executives can use this framework as both a diagnostic tool for assessing governance readiness and a roadmap for building capabilities that accelerate responsible AI adoption while mitigating legal, ethical, and reputational risks.

## Keywords

AI Governance, Ethical AI, Healthcare Management, Bias Mitigation, Organizational Frameworks, Responsible AI

---

## 1. Introduction

Artificial intelligence is rapidly transforming healthcare delivery through applications spanning clinical diagnostics, predictive analytics, and operational optimization. However, this transformation brings substantial ethical risks. Algorithmic bias can perpetuate health disparities, lack of transparency undermines clinical trust, and inequitable deployment may widen existing care gaps (Ahadian et al., 2026; Dankwa-Mullan et al., 2021). High-level ethical principles—including beneficence, non-maleficence, autonomy, and justice, increasingly guide AI ethics discourse, yet healthcare organizations struggle to embed these abstract values into operational decisions and governance structures (Chan et al., 2025; Freeman et al., 2025).

A recent scoping review examining a decade of AI ethics framework implementation found that while organizations frequently established ethics committees and review processes, evidence of measurable impact on outcomes, such as bias reduction or equity improvement, remained sparse (Chan et al., 2025). Three critical barriers impede translation of ethical principles into practice: organizational silos among clinicians, data scientists, and compliance teams; regulatory uncertainty about adequate governance; and resource constraints, particularly in smaller hospitals (Freeman et al., 2025; Kim et al., 2025).

This paper addresses these barriers by adapting the People-Process-Technology-Operations (PPTO) framework specifically for ethical AI governance in hospitals. While PPTO has been applied to organizational AI governance generally (Grant & Levasseur, 2025; Kim et al., 2025), this work makes three distinct contributions: explicitly connecting PPTO pillars to core ethical principles, synthesizing recent empirical work into actionable guidance for hospitals at varying maturity levels, and providing operational specificity through example roles, workflows, and metrics that executives can implement (Freeman et al., 2025; Wells et al., 2025).

## 2. Background: The Principles-to-Practice Gap

Contemporary AI ethics frameworks emphasize core principles adapted to healthcare contexts: beneficence and non-maleficence (AI should improve outcomes and avoid harm), justice and equity (AI should reduce disparities), autonomy and human oversight (clinicians should retain meaningful control), and transparency and accountability (stakeholders should understand how systems work and who is responsible) (Ahadian et al., 2026; Dankwa-Mullan et al., 2021). While these principles enjoy broad endorsement, translating them into operational governance remains challenging.

Chan and colleagues' systematic review revealed that organizations commonly implemented ethics processes—oversight committees, review checklists, documentation requirements—but rarely demonstrated outcome-level improvements in equity, safety, or deployment reliability (Chan et al., 2025). Their review identified patterns contributing to this gap: procedural compliance without outcome tracking, limited stakeholder engagement, insufficient lifecycle integration, and

weak connection to existing clinical governance structures.

Research on AI leadership in healthcare organizations reveals that successful AI transformation requires executive sponsorship, clear accountability structures, cultural readiness, and systematic capacity building (Grant & Levasseur, 2025; Sriharan et al., 2024). These organizational factors create the foundation upon which effective governance frameworks can be built, emphasizing that governance is fundamentally an organizational capability challenge rather than purely a technical or policy exercise.

### 3. The PPTO Framework for Ethical AI Governance

The PPTO framework structures organizational capabilities into four interdependent pillars: **People** (defined roles, competencies, and leadership structures ensuring ethical stewardship), **Process** (standardized workflows embedding ethical considerations across the AI lifecycle), **Technology** (tools enabling transparency, monitoring, and auditability), and **Operations** (metrics linking governance to measurable performance in safety, equity, and trust).

These pillars are mutually reinforcing. Governance committees (People) execute review workflows (Process) using documentation tools (Technology) while tracking equity metrics (Operations). The framework draws on sociotechnical systems theory, which recognizes that effective governance requires alignment between social elements (people, culture, roles) and technical elements (tools, infrastructure), with processes and operations mediating their interaction (Grant & Levasseur, 2025).

#### 3.1. Governance Scope

The PPTO framework applies to four primary categories of hospital AI: vendor-supplied clinical decision support tools (e.g., sepsis prediction, radiology interpretation aids), internally developed predictive models, generative AI tools used in clinical or administrative workflows, and operational analytics systems affecting resource allocation or staffing. These categories represent the AI applications most commonly deployed in hospital settings and carrying the greatest potential for patient harm, inequitable impact, or regulatory exposure (Freeman et al., 2025; Wells et al., 2025).

Out-of-scope applications include basic rule-based automation, standard statistical reporting tools, and administrative software without predictive or generative components, as these lack the opacity and bias risks that necessitate AI-specific governance (Kim et al., 2025). Establishing this boundary prevents governance overreach while ensuring that systems capable of perpetuating health disparities or undermining clinical judgment receive appropriate oversight (Ahadian et al., 2026; Dankwa-Mullan et al., 2021).

#### 3.2. Operationalizing Ethical Principles

**Table 1** illustrates how PPTO operationalizes ethical principles, demonstrating

how abstract ethical commitments convert into concrete organizational capabilities across all four pillars.

**Table 1.** Mapping ethical principles to PPTO governance pillars.

<b>Ethical Principle</b>	<b>People</b>	<b>Process</b>	<b>Technology</b>	<b>Operations</b>
<b>Justice &amp; Equity</b>	Equity champions on governance committees	Mandatory fairness assessments	Bias measurement dashboards	Equity metrics by subgroup
	Diverse stakeholder representation	Subgroup performance evaluation	Demographic stratification tools	Disparity Reduction targets
	Training on health disparities	Equity impact reviews	Disparity alert systems	Distributional impact tracking
<b>Transparency &amp; Explainability</b>	Documentation training for staff	Standardized model documentation	Model cards and fact sheets	Transparency perception Surveys
	Communication specialists	Decision logging requirements	Audit trails and logs	Documentation compliance rates
	Patient engagement roles	Explainability assessments	User-friendly reporting interfaces	Stakeholder accessibility metrics
<b>Autonomy &amp; Human Oversight</b>	Clinical Champions Preserve judgment	Human in-the-loop requirements	Override Functionality in systems	Clinician Autonomy Satisfaction
	Frontline staff in governance	Override Mechanism protocols	Alert fatigue monitoring tools	Override rate tracking by system
	Override authority clearly defined	Workflow integration reviews	Clinical decision support safeguards	Automation bias indicators
<b>Beneficence &amp; Safety</b>	Safety officers on committees	Adverse event protocols	Safety performance dashboards	Clinical outcome metrics
	Clinical quality expertise	Performance monitoring workflows	Drift detection and alerts	Near-miss and incident rates
	Patient safety champions	Safety review checkpoints	Incident tracking systems	Safety threshold compliance

**Continued**

	Clear decision rights defined	Documented approval authority	Traceability systems	Governance coverage metrics
<b>Accountability</b>	Escalation pathways established	Audit and review procedures	Decision documentation tools	Audit completion rates
	Executive sponsors assigned	Responsibility matrices	Compliance tracking databases	Accountability gap analyses

For justice and equity: equity champions serve on governance committees (People), mandatory fairness assessments are required (Process), bias measurement dashboards are deployed (Technology), and equity metrics are tracked by subgroups (Operations). For transparency: staff receive training on documentation (People), standardized model documentation is required (Process), model cards and audit trails are maintained (Technology), and transparency perception surveys are conducted (Operations).

#### 4. People: Governance Roles and Structures

Effective AI governance begins with clearly defined roles bridging clinical, technical, ethical, and administrative domains. A cross-functional AI Governance Committee serves as the primary decision-making body, including clinical leadership, data science and IT representatives, compliance and legal counsel, quality and patient safety officers, equity champions, and operations representatives (Freeman et al., 2025; Kim et al., 2025). The committee operates under a formal charter specifying decision authority, meeting cadence, and reporting obligations to executive leadership.

Beyond the committee, specialized roles support effective oversight. An AI Ethics Champion (0.5 - 1.0 Full Time Equivalent (FTE)) leads ethics review processes and develops review criteria and training materials. A Clinical AI Champion (0.25 - 0.5 FTE) translates between clinical workflows and technical design while engaging frontline clinicians. An AI Governance Officer (1.0 FTE in mature programs) coordinates committee activities, maintains documentation, and manages project intake workflows. A Training Lead (0.25 - 0.5 FTE) develops governance training for stakeholders at all levels (Reid & Levasseur, 2025).

Systematic training programs ensure stakeholders understand governance responsibilities. Executive leadership receives 2 - 4 hours annually on AI fundamentals and strategic opportunities. Governance committee members receive 8 - 12 hours initially on ethics principles, bias mitigation, and governance processes. AI project teams receive 4 - 6 hours on documentation requirements and equity assessment methods. Frontline clinicians receive 1 - 2 hours on appropriate use of specific AI tools as they are deployed (Sriharan et al., 2024).

## 5. Process: Lifecycle-Integrated Governance Workflows

Process pillars translate ethical principles into standardized workflows ensuring consistent oversight across the AI lifecycle. Governance begins with project intake using a standardized form capturing intended purpose, affected populations, data requirements, anticipated risks, and resource needs. This enables early risk triage, categorizing projects as low, medium, or high risk based on patient impact and complexity (Freeman et al., 2025; Wells et al., 2025).

### 5.1. Risk-Tiering Rubric for Project Intake

The intake risk-tiering rubric evaluates five criteria drawn from the healthcare AI governance literature: patient harm potential (likelihood and severity of adverse clinical outcomes), autonomy impact (degree to which the system influences or displaces clinical judgment), scale of deployment (number of patients, clinicians, or facilities affected), regulatory exposure (FDA oversight, billing implications, or liability considerations), and equity sensitivity (whether affected populations include historically marginalized or vulnerable groups) (Dankwa-Mullan et al., 2021; Freeman et al., 2025; Wells et al., 2025).

Projects scoring high on any single criterion or moderate on multiple criteria are categorized as follows. Low-risk systems score low across all five criteria and proceed through streamlined documentation review by the Governance Officer without full committee deliberation. Medium-risk systems exhibit moderate scores on one or more criteria and require bias and fairness assessment, model documentation review, and committee approval before deployment. High-risk systems score high on patient harm potential, autonomy impact, or equity sensitivity and require the full multi-stakeholder review process, including mandatory fairness assessment, external expert consultation where appropriate, and documented committee approval with executive notification (Freeman et al., 2025; Kim et al., 2025). This tiering approach mirrors the risk-stratified governance structure validated in practice by Wells and colleagues' FAIR-AI framework (Wells et al., 2025).

### 5.2. Equity and Fairness Assessment Requirements

For medium- and high-risk systems, the minimum required equity and fairness assessment package includes mandatory subgroup stratification across race and ethnicity, sex, age group, insurance status or payer type, and primary language, as these dimensions reflect the demographic axes along which health disparities are most consistently documented in the United States healthcare system (Dankwa-Mullan et al., 2021). Where data availability permits, stratification should additionally include disability status and rural versus urban geography.

Acceptable fairness metrics include at minimum: demographic parity (equal positive prediction rates across subgroups), equalized odds (equal true positive and false positive rates), and predictive parity (equal positive predictive value across subgroups). Project teams must report which metrics were assessed, which

were not feasible given data constraints, and the rationale for any metric not meeting threshold (Wells et al., 2025). Systems failing to achieve equity thresholds—defined as a performance disparity exceeding ten percentage points across any mandatory subgroup on a primary metric—require documented remediation plans prior to deployment approval (Dankwa-Mullan et al., 2021; Freeman et al., 2025). This specificity transforms fairness assessment from an interpretive exercise into an implementable requirement, directly addressing the evidence gap Chan and colleagues identified in which ethics processes rarely demonstrated measurable equity outcomes (Chan et al., 2025).

### 5.3. Accountability and Decision-Rights Model

The accountability pillar is operationalized through the following decision-rights structure across five governance functions:

- **Deployment Approval:** The AI Governance Committee is Responsible and Accountable; Clinical AI Champion and Data Science lead are Consulted; IT and Compliance are Informed.
- **Ongoing Monitoring Ownership:** Data Science and IT are Responsible; the AI Governance Officer is Accountable; the Ethics Champion is Consulted; Clinical Leadership is Informed.
- **Escalation Triggers:** The AI Governance Officer is Responsible for identifying and escalating; the Committee Chair is Accountable for convening review; Clinical, Safety, and Compliance leads are Consulted; Executive Leadership is Informed.
- **Corrective Action (Modification or Suspension):** The AI Governance Committee is Responsible and Accountable; Data Science, IT, and Clinical AI Champion are Consulted; Operations and Legal are Informed.
- **Stop or Decommission Authority:** Executive Sponsor holds final Accountable authority; the AI Governance Committee is Responsible for the recommendation; Legal, Compliance, and Clinical Leadership are Consulted; all affected operational units are Informed.

This structure reflects the clear accountability architecture that Sriharan and colleagues identified as essential for successful AI transformation in healthcare organizations, where ambiguity in decision rights frequently stalls governance action (Sriharan et al., 2024). It also directly addresses the responsibility matrices included in the Accountability row of **Table 1** by giving those matrices concrete role-by-function assignments (Freeman et al., 2025; Kim et al., 2025).

### 5.4. Patient and Community Stakeholder Engagement

For all high-risk AI applications—particularly those affecting patient access, triage, diagnosis, or treatment recommendation—a mandatory patient and community review touchpoint is required prior to deployment approval. This touchpoint takes the form of a structured review session conducted with patient advocates, community advisory board members, or representatives from populations most

likely to be affected by the system. Reviewers receive and evaluate three standardized artifacts: the model fact sheet (documenting intended use, known limitations, and performance metrics), the equity assessment summary (presenting subgroup performance results and any identified disparities), and the monitoring plan (specifying how performance will be tracked post-deployment and how concerns can be reported) (Freeman et al., 2025; Wells et al., 2025).

Feedback from this review is documented and presented to the AI Governance Committee as part of the deployment decision record. The committee must formally acknowledge community input and document how it was addressed or, where it was not incorporated, provide written justification. This requirement responds directly to the limitation Chan and colleagues identified—that existing AI ethics frameworks exhibit limited stakeholder engagement—and reflects Ahadian and colleagues' finding that true equity in AI governance requires engaging affected communities rather than relying solely on technical fairness metrics (Ahadian et al., 2026; Chan et al., 2025).

### **5.5. Development, Validation, and Deployment**

During development and validation, governance processes require bias and fairness assessment (analyzing training data representativeness and evaluating performance across demographic subgroups), model documentation (capturing intended use, limitations, performance metrics, and monitoring plans), transparency mechanisms (ensuring appropriate explainability), and human factors assessment (addressing workflow integration and automation bias) (Dankwa-Mullan et al., 2021; Wells et al., 2025).

Before deployment, high- and medium-risk applications undergo multi-stakeholder review assessing documentation completeness, performance adequacy, equity alignment, workflow readiness, and monitoring plans. The committee may approve, require modifications, or decline deployment, with decisions documented for auditability (Kim et al., 2025).

Post-deployment monitoring tracks model performance continuously, detects drift, monitors subgroup performance for emerging disparities, and tracks clinical outcomes, utilization patterns, and adverse events. Pre-specified thresholds trigger escalation for review, retraining, or decommissioning. To avoid parallel bureaucracies, AI governance integrates with existing quality improvement, patient safety, clinical protocol review, and IT governance processes (Freeman et al., 2025; Sriharan et al., 2024).

## **6. Technology: Enabling Infrastructure**

Technology infrastructure supports governance by making AI systems visible, documentable, and monitorable. A centralized AI Project Registry tracks all projects under development or deployed, with project owners, risk categories, compliance status, and deployment locations. The registry enables executives to understand their AI portfolio and identify governance gaps (Kim et al., 2025; Wells

et al., 2025).

Standardized model documentation templates ensure consistent capture of model purpose, data sources, performance metrics, limitations, and monitoring plans. These can be implemented through document libraries initially, progressing to structured databases as maturity increases. Bias and fairness assessment tools support demographic data analysis, fairness metric calculation, performance visualization across subgroups, and drift detection (Dankwa-Mullan et al., 2021).

Monitoring infrastructure includes performance dashboards, automated alerts when metrics drop below thresholds, audit logs recording predictions and user interactions, and integration with clinical systems for outcome tracking. Technology investments should match organizational maturity: early-stage organizations use spreadsheet-based registries and manual assessments; growing programs implement database-backed systems and semi-automated tools; mature programs deploy enterprise platforms with real-time monitoring and automated alerting (Grant & Levasseur, 2025; Kim et al., 2025).

## 7. Operations: Linking Governance to Outcomes

The Operations pillar ensures governance delivers measurable value through performance indicators across process, technical, and organizational domains. Process metrics assess whether governance workflows function as intended: percentage of AI projects undergoing required reviews (target 100% for medium/high-risk), documentation completion rates, training completion rates, median time from intake to decision, and stakeholder satisfaction with governance processes (Chan et al., 2025; Freeman et al., 2025).

Technical metrics assess whether AI systems meet safety and equity standards: model performance metrics for each system, performance disparities across demographic subgroups, percentage of systems meeting equity thresholds (target >90%), AI-related adverse events, and override rates revealing misalignment with clinical judgment (Dankwa-Mullan et al., 2021; Wells et al., 2025).

Organizational metrics connect governance to strategic objectives: clinician and patient trust in AI tools, alignment of AI projects with equity goals, AI contribution to quality improvement targets, legal or regulatory actions (target zero), and return on governance investment measured as governance costs relative to total AI investment (Grant & Levasseur, 2025; Sriharan et al., 2024).

### 7.1. Monitoring-and-Response Loop: A Worked Example

The following example illustrates how the Operations pillar drives decisions rather than reporting alone, using a sepsis prediction model deployed across inpatient units.

Metric: Sensitivity of the sepsis alert model stratified by race, tracked monthly by the Data Science team.

Threshold: A drop in sensitivity exceeding five percentage points in any mandatory subgroup relative to the validated baseline, or a disparity exceeding ten

percentage points between the highest- and lowest-performing subgroups, triggers formal review.

**Alert:** The automated monitoring dashboard flags the threshold breach and generates a governance alert routed to the AI Governance Officer within 24 hours of the monthly performance calculation.

**Review Body:** The AI Governance Officer convenes the AI Governance Committee within ten business days. The Ethics Champion and Clinical AI Champion prepare a root cause analysis examining whether the disparity reflects data drift, population shift, workflow changes, or encoding bias in updated training data.

**Corrective Action:** Based on the root cause analysis, the committee selects from three pre-specified response pathways: (1) model retraining with updated and more representative data if data drift is confirmed; (2) workflow intervention and clinician alert adjustment if the issue reflects utilization patterns rather than model performance; or (3) suspension of the model pending full remediation review if neither pathway is feasible within 30 days.

**Documentation:** All committee deliberations, the root cause analysis, the selected corrective pathway, the responsible parties, and the target resolution date are recorded in the centralized AI Project Registry. A follow-up performance report is scheduled for the next monthly cycle, with results reported in the quarterly executive dashboard (Freeman et al., 2025; Kim et al., 2025; Wells et al., 2025).

This loop demonstrates how the metrics catalogued in Section 7 connect directly to governance decisions, transforming the Operations pillar from a reporting function into an active driver of responsible AI oversight—consistent with the outcome-oriented approach that distinguishes PPTO from the procedural compliance models Chan and colleagues found insufficient across a decade of healthcare AI ethics implementation (Chan et al., 2025).

## 7.2. Reporting and Continuous Improvement

Metrics inform regular reporting: monthly internal dashboards on process efficiency, quarterly executive reports on governance coverage and strategic alignment, and annual comprehensive reviews. Data drives continuous improvement through root cause analysis of failures, benchmarking against peer organizations, stakeholder feedback integration, and regular metric review as priorities evolve (Chan et al., 2025; Freeman et al., 2025).

## 8. Implementation Roadmap

Organizations should implement PPTO incrementally based on current maturity. Stage 1 (Months 1 - 6) establishes foundations: initial governance committee, designated governance officer (0.25 - 0.5 FTE), executive training, simple intake forms, risk categorization, and basic project tracking. Stage 2 (Months 7 - 12) standardizes processes: specialized roles added, formalized review workflows, bias assessment checklists, deployment criteria, and expanded metrics. Stage 3 (Months 13 - 24) integrates governance with clinical quality processes, establishes monitoring

protocols, deploys database-backed registries, and tracks full metric suites. Stage 4 (Months 24+) optimizes through advanced training, process refinement, enterprise platforms, real-time monitoring, and demonstrated outcome improvements (Kim et al., 2025; Wells et al., 2025).

This staged approach allows capability building while demonstrating value at each stage, facilitating sustained leadership support. Early-stage organizations with limited AI activity implement foundational governance with modest commitments (2 - 3 FTE-months initially). Growing programs standardize as project volumes justify investment. Mature programs optimize through automation and advanced analytics (Grant & Levasseur, 2025).

## 9. Discussion

The core contribution of this work is demonstrating how PPTO provides a systematic bridge from abstract ethical principles to concrete organizational capabilities. While existing frameworks articulate important values, they often lack operational specificity needed for implementation (Chan et al., 2025; Freeman et al., 2025). PPTO addresses this by organizing governance into four domains that collectively determine whether ethical commitments translate into measurable outcomes.

PPTO explicitly addresses the evidence gap identified by Chan et al. (2025): lack of outcome tracking, insufficient lifecycle integration, limited stakeholder engagement, and weak integration with existing structures. By requiring measurable indicators across technical, organizational, and strategic domains, extending oversight across full AI lifecycles, mandating cross-functional committees, and emphasizing integration with clinical governance, PPTO creates conditions for demonstrable impact (Freeman et al., 2025; Wells et al., 2025).

A critical advantage is scalability across organizations with varying resources. The staged roadmap recognizes comprehensive governance requires sustained investment but value can be demonstrated incrementally. Resource-limited hospitals focus initial efforts on high-impact, low-cost activities: designating existing staff to governance roles, using document templates rather than specialized software, and leveraging open-source bias assessment tools (Kim et al., 2025; Sriharan et al., 2024).

The framework explicitly centers equity alongside safety and effectiveness, recognizing that historical health disparities may be perpetuated by AI trained on historical data and that vulnerable populations face compounded risks (Dankwa-Mullan et al., 2021). PPTO operationalizes equity through equity champions on committees, mandatory fairness assessments, bias dashboards, and disparity tracking with reduction targets. However, technical metrics alone are insufficient; true equity requires engaging affected communities in governance and prioritizing AI applications addressing rather than exacerbating disparities (Ahadian et al., 2026).

Effective governance ultimately depends on organizational culture. PPTO sup-

ports cultural transformation through clarity and transparency in expectations, psychological safety for raising concerns, visible leadership commitment, systematic capacity building, and accountability through metrics and performance expectations (Grant & Levasseur, 2025; Reid & Levasseur, 2025).

## 10. Limitations

This work has several important limitations. The paper is conceptual rather than empirical, synthesizing existing literature to propose an adapted PPTO model without presenting original implementation data. Recommendations require empirical validation through rigorous implementation studies. Generalizability is limited, as examples draw primarily from larger health systems in high-income countries; applicability to smaller hospitals, rural centers, and international settings may require substantial adaptation (Freeman et al., 2025; Kim et al., 2025).

Concrete evidence on whether PPTO implementation reduces algorithmic bias or health disparities in practice remains limited. The proposed equity mechanisms are theoretically sound but require empirical testing, particularly in settings serving highly marginalized populations (Dankwa-Mullan et al., 2021). Resource estimates are illustrative rather than empirically derived; actual needs will vary substantially. Technology tool maturity remains limited, and organizations may face implementation challenges. The rapidly evolving AI landscape means guidance current in 2025-2026 may become outdated as new capabilities, regulations, and governance approaches emerge (Ahadian et al., 2026; Wells et al., 2025).

## 11. Conclusion

The proliferation of AI in healthcare creates both extraordinary opportunities and significant ethical risks. While high-level ethical principles provide important guidance, healthcare organizations struggle to translate aspirations into operational reality. This paper addresses this implementation gap by adapting the PPTO framework specifically for ethical AI governance in hospitals, providing systematic structure for organizing governance capabilities across people, processes, technology, and operations.

The framework's contributions include explicitly connecting ethical principles to specific governance mechanisms, providing operational specificity through detailed roles and workflows, and emphasizing scalability across varying organizational resources while maintaining ethical rigor. Implementation enables hospitals to pursue AI innovation responsibly, building stakeholder trust, reducing deployment failures, and preventing costly ethical missteps (Grant & Levasseur, 2025; Sriharan et al., 2024).

However, realizing these benefits requires leadership commitment, adequate resources, cultural transformation, and sustained attention to equity. Organizations must resist treating governance as compliance checkboxes, instead viewing it as integral to responsible innovation. The research agenda ahead is substantial; rigorous empirical evaluation remains needed to transform PPTO from concep-

tually sound framework into evidence-based governance model (Chan et al., 2025; Freeman et al., 2025).

As AI capabilities advance, governance frameworks must evolve correspondingly. The principles underlying PPTO—systematic stakeholder engagement, lifecycle integration, transparency and accountability, outcome orientation, and equity centrality, provide enduring foundations even as specific tools and processes adapt. By translating principles into practice through aligned people, robust processes, enabling technologies, and measurable outcomes, hospitals can ensure AI serves patients and communities rather than subjecting them to avoidable risks (Ahadian et al., 2026; Wells et al., 2025).

## Conflicts of Interest

The authors declare no conflicts of interest regarding the publication of this paper.

## References

- Ahadian, P., Xu, W., Liu, D., & Guan, Q. (2026). Ethics of Trustworthy AI in Healthcare: Challenges, Principles, and Practical Pathways. *Neurocomputing*, 661, Article 131942. <https://doi.org/10.1016/j.neucom.2025.131942>
- Chan, A., Rahimi-Ardabili, H., Rogers, W. A., & Coiera, E. (2025). The Real-World Impact of Artificial Intelligence Ethics Frameworks across a Decade in Healthcare: A Scoping Review. *Journal of the American Medical Informatics Association*, 32, 1767-1777. <https://doi.org/10.1093/jamia/ocaf167>
- Dankwa-Mullan, I., Scheufele, E. L., Matheny, M. E., Quintana, Y., Chapman, W. W., Jackson, G. et al. (2021). A Proposed Framework on Integrating Health Equity and Racial Justice into the Artificial Intelligence Development Lifecycle. *Journal of Health Care for the Poor and Underserved*, 32, 300-317. <https://doi.org/10.1353/hpu.2021.0065>
- Freeman, S., Wang, A., Saraf, S., Potts, E., McKimm, A., Coiera, E. et al. (2025). Developing an AI Governance Framework for Safe and Responsible AI in Health Care Organizations: Protocol for a Multimethod Study. *JMIR Research Protocols*, 14, e75702. <https://doi.org/10.2196/75702>
- Grant, V., & Levasseur, R. E. (2025). Artificial Intelligence in Action: Navigating Advancements, Pitfalls, and Human Dynamics. *Open Journal of Business and Management*, 13, 3865-3874. <https://doi.org/10.4236/ojbm.2025.136210>
- Kim, J. Y., Hasan, A., Kueper, J., Tang, T., Hayes, C., Fine, B. et al. (2025). Establishing Organizational AI Governance in Healthcare: A Case Study in Canada. *npj Digital Medicine*, 8, Article No. 522. <https://doi.org/10.1038/s41746-025-01909-3>
- Reid, K. T., & Levasseur, R. E. (2025). The Ethical Use of Artificial Intelligence for Employee Supervision and Its Effect on Organizational Culture. *Open Journal of Business and Management*, 13, 3206-3220. <https://doi.org/10.4236/ojbm.2025.135169>
- Sriharan, A., Sekercioglu, N., Mitchell, C., Senkaiahliyan, S., Hertelendy, A., Porter, T. et al. (2024). Leadership for AI Transformation in Health Care Organization: Scoping Review. *Journal of Medical Internet Research*, 26, e54556. <https://doi.org/10.2196/54556>
- Wells, B. J., Nguyen, H. M., McWilliams, A., Pallini, M., Bovi, A., Kuzma, A. et al. (2025). A Practical Framework for Appropriate Implementation and Review of Artificial Intelligence (FAIR-AI) in Healthcare. *npj Digital Medicine*, 8, Article No. 514. <https://doi.org/10.1038/s41746-025-01900-y>