



SPECIAL CONTRIBUTION

The Practice of Emergency Medicine



Check for updates

An Evidence-Based Framework for Patient-Facing Artificial Intelligence Integration in the Emergency Department

Tehreem Rehman MD, MPH¹ , Philip Jarrett MD, MBA², Joshua Lesko MD³, James Augustine MD⁴, Bradley D. Shy MD⁵, Rohit B. Sangal MD, MBA⁶, Nicholas Genes MD, PhD⁷, Donald U. Apakama MD, MS^{1,8}, Ethan E. Abbott DO, MSCR¹, Abhi Mehrotra MD, MBA⁹, Richard Andrew Taylor MD, MHS¹⁰

¹Department of Emergency Medicine, Icahn School of Medicine at Mount Sinai, New York, New York, USA

²Department of Emergency Medicine, University of Texas Southwestern Medical Center, Dallas, Texas, USA

³Eastern Virginia Medical School at ODU Emergency Medicine, Norfolk, Virginia, USA

⁴Department of Emergency Medicine, Wright State University, Dayton, Ohio, USA

⁵Department of Emergency Medicine, University of Colorado School of Medicine, Aurora, Colorado, USA

⁶Department of Emergency Medicine, Yale University, New Haven, Connecticut, USA

⁷Ronald O. Perelman Department of Emergency Medicine, NYU Grossman School of Medicine, New York, New York, USA

⁸Artificial Intelligence and Human Health, Icahn School of Medicine at Mount Sinai, New York, New York, USA

⁹Department of Emergency Medicine, University of North Carolina, Chapel Hill, North Carolina, USA

¹⁰Department of Emergency Medicine, University of Virginia, Charlottesville, Virginia, USA

Correspondence

Tehreem Rehman, MD, MPH, 1468 Madison Ave, MC level, Room 220 New York, NY 10029, USA. Email: tehreem.rehman@mountsinai.org

Received: January 5, 2026

Revised: April 26, 2026

Accepted: June 11, 2026

<https://doi.org/10.1016/j.acepjo.2026.100466>

Abstract

The emergency department is a high-need setting for patient-facing artificial intelligence (PF-AI) because it functions as both a health care safety net and, increasingly, a digital front door for unscheduled acute care. We define PF-AI as tools that interact directly with patients to collect information, deliver tailored communication, and support decision making under clinician governance. This article presents a constrained framework for where PF-AI may add value across the emergency care journey—including structured intake, patient education, multilingual communication, and postdischarge follow-up—and where strict boundaries are required given risks of undertriage, confabulated content, workflow burden, privacy harm, inequitable access, cost escalation, and medico-legal ambiguity. We map operational touchpoints from prearrival to follow-up, specify information-flow and accountability checkpoints, and propose a human oversight model with defined escalation roles. The

abstract continues

Supervising Editor: Adam Landman, MD, MS

© 2026 The Authors. Published by Elsevier Inc. on behalf of American College of Emergency Physicians. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Abstract (continued)

framework emphasizes equity, interoperability, usability, privacy, safety verification, and auditability, recognizing that PF-AI should augment rather than replace clinician judgment or existing non-AI workflows.

Keywords: *patient-facing artificial intelligence, emergency medicine, large language models, human oversight, health equity, natural language processing, patient safety*

1 BACKGROUND

For millions of patients each year, the emergency department (ED) serves as the first point of entry into the health care system.¹ Many arrive with a fragmented medical record, inconsistent ambulatory care, or limited engagement with digital health tools.^{2–4} As a result, the ED functions both as a communication bottleneck and a primary entry point for clinical data. Key history, medications, and contact information are captured and routed into longitudinal records, under pressure and with incomplete context. At the same time, the ED visit contains multiple periods of unavoidable waiting: before triage, while tests are pending, and prior to discharge. These pauses, often experienced as idle or anxiety-provoking, can instead be reframed as opportunities for information capture, education, and patient engagement.

In this manuscript, patient-facing artificial intelligence (PF-AI) is an operational term for patient-interactive digital systems that perform three functions: (1) data transformation, converting patient-reported information into structured records through natural language processing; (2) workflow automation, including asynchronous triage-support and follow-up coordination; and (3) multimodal communication across text, voice, and image modalities. Current evidence for these capabilities is heterogeneous and context-dependent.^{5–10}

We present PF-AI as a constrained implementation framework, not a universal solution. In many circumstances, simpler alternatives (for example, structured forms, nurse-led protocols, or nonartificial-intelligence digital tools) may be safer or more practical. The key question is not whether artificial intelligence is available, but whether it adds measurable value over existing workflows for specific ED populations and tasks. Recent commercialization of patient-facing health AI platforms and growing consumer wearable signal streams further underscore the need for ED-specific governance and evidence standards.^{11–16}

Risks are substantial and cluster in three domains. Technical risks include unreliable outputs^{17–19} and model bias.²⁰ Implementation risks include fragmented integration into ED workflows²¹ and clinician cognitive overload from parallel data streams.²² Structural risks include digital- and health-literacy barriers,²³ the environmental cost of compute-intensive training,²⁴ and unresolved liability when recommendations contribute to harm.^{25,26}

This framework is written for emergency clinicians, health-system executives, and clinical informatics teams

designing PF-AI programs in emergency care. The evaluative domains used in this framework (equity, interoperability, usability, privacy, safety, and accountability) are grounded in the World Health Organization's ethics and governance guidance for AI in health,²⁷ clinical reporting and audit standards, and health-system governance recommendations. We organize applications across six phases of the ED visit (previsit, triage, preclinician evaluation, postevaluation communication, discharge, follow-up) with explicit human oversight expectations at each step.^{27–30}

PF-AI tools must account for the realities of emergency care: time pressure, emotional intensity, and variable digital literacy. Patients in pain, acute anxiety, or cognitive distress process information differently than ambulatory populations, and effective design requires accommodating these states directly.³¹ PF-AI should always be paired with an equivalent staff-assisted fallback, such as interpreter-supported intake, bedside staff engagement, or standard nondigital workflows, so patients who cannot or do not wish to engage digitally are not excluded.^{32,33}

PF-AI is most useful when translating complex medical information into accessible formats. Systems should avoid vague verbal probability terms and instead use consistent numerical denominators and absolute risk differences. Visual risk displays should use part-to-whole formats such as icon arrays, which outperform bar charts and numeric probability alone for conveying risk magnitude to patients with variable numeracy.³⁴ Information delivery should also be multimodal: audiovisual media improve comprehension for patients who struggle with dense text, particularly when explaining treatment options or discharge instructions.^{35,36}

Design must account for the diverse cognitive profiles of the ED population. For aging patients or those with limited health literacy, interfaces should prioritize universal symbols and simplified navigation structures that require minimal steps.^{23,37,38} Content should be tailored not just to the patient's medical history but to their status in the ED journey, providing timely updates that reduce anxiety without causing information overload.^{32,33} For patients with cognitive impairment or those overwhelmed by the ED environment, AI tools can act as "scaffolding," breaking down tasks into manageable components and offering step-by-step guidance.²³

Patients must understand they are interacting with an automated system, not a clinician. PF-AI interfaces should use clear disclosure and nonanthropomorphic design cues to distinguish automated support from clinician judgment. Trust

also depends on transparent privacy practices, data provenance, and performance monitoring across demographic groups.³⁹ Ransomware attacks that have disrupted emergency care at adjacent hospitals have made health care cybersecurity risks visible to patients themselves.^{40,41} Because PF-AI tools handle sensitive patient-reported data, explicit safeguards are a precondition for trust, not a downstream assurance.

By codesign, we refer to sustained, iterative partnership with patients and caregivers across problem definition, workflow design, prototype testing, and governance decisions, rather than one-time feedback sessions. Participatory design should intentionally include groups at highest risk of exclusion, so that PF-AI systems do not widen disparities while attempting to improve access.^{32,36,42}

A central concern is that PF-AI could perpetuate or amplify existing health disparities. When training data reflect historical biases, algorithms perform less accurately for underrepresented groups, including racial minorities and patients with lower socioeconomic status.^{20,43} Digital technology adoption itself varies by age, race, and ethnicity.⁴⁴ These gaps are most acute for non-English-speaking populations, who face compounded barriers of limited broadband access, lower digital literacy, and a lack of linguistically or culturally tailored interfaces.^{9,44}

Mitigating these risks requires ongoing equity audits, transparent reporting of model performance across demographic subgroups, and active participation of patients and community representatives in both design and evaluation.^{28,45} Codesign surfaces biases early and keeps PF-AI aligned with

patient-centered needs rather than structural assumptions.^{32,46} Combined with robust privacy protections, explainability, and continuous monitoring for bias and error,^{35,47} these practices allow PF-AI to strengthen communication and safety without displacing the human connection that defines emergency medicine.

2 THE PATIENT-CARE JOURNEY: INTEGRATING PATIENT-FACING AI

Patient-facing AI most clearly adds value when it reduces friction in predictable, guideline-driven touchpoints while preserving clinician authority at decision points. We organize applications across six phases of emergency care—previsit, triage, preclinician evaluation, postevaluation communication, discharge, and follow-up—and apply a uniform test at each: does PF-AI add measurable value over existing workflows, and what human oversight structure is required to use it safely? [Table 1](#) summarizes evidence-based applications across these phases.

2.1 Previsit

Before patients reach the triage desk, PF-AI can involve them directly in information gathering. Deployment channels include patient portals, mobile apps, and tablets in ambulances or waiting areas, each allowing patients to enter or verify information in real time. Pre-triage tools can analyze these inputs to guide patients toward the most appropriate

TABLE 1. Patient-facing AI applications across the emergency care journey.

Phase	Representative PF-AI applications	Required human oversight
Previsit	Patient-portal and mobile symptom intake; multilingual pre-triage narrative capture; on-route ambulance or waiting-area data gathering	Licensed triage nurse retains acuity authority; staff verify any fields that could imply acuity
ED greeting and triage	Structured patient-entered data linked to evidence-based risk scores; patient goal and concern elicitation; dynamic queue and wait-time updates labeled as estimates	Narrative capture separated from acuity assignment; licensed triage nurse retains final triage decision
Posttriage and preclinician evaluation	LLM-enabled intake mapping spoken narrative to structured fields; personalized preventive screening and education micro-modules; voluntary virtual lower-acuity pathways	Role-stratified review (nurse or quality coordinator for low-risk; physician for diagnostic-adjacent outputs); EMTALA-compliant reversible pathways
Postevaluation communication	Condition-specific education reinforcement for new or incidental findings; closed-loop linkage to social-needs referrals	Clinician triggers and validates patient-facing content; social work or care management confirms referral receipt
Transition to hospitalization	Plain-language explanation of boarding steps; asynchronous patient Q&A while awaiting bed placement; expectation-setting for admission	Nursing sign-off on clinical content; clear escalation pathway for deterioration-related questions
Discharge	Literacy-leveled discharge instruction generation; medication counseling; red-flag symptom reinforcement	Clinician review and sign-off of all patient-facing text before release; no full automation of discharge communication
Postdischarge follow-up	Automated personalized check-ins; symptom escalation triage; medication-adherence and return-precaution reinforcement	Each outbound check-in mapped to a defined backend action owner (nurse coordinator, pharmacist, social worker)

ED, Emergency Department; EMTALA, Emergency Medical Treatment and Active Labor Act; LLM, large language model; PF-AI, patient-facing artificial intelligence; Q&A, question and answer.

care setting, potentially reducing unnecessary ED volume.⁴⁸ Patients can also review key elements of their prehospital story so that handoffs to clinicians are both clinically accurate and patient-validated.⁴⁹

Natural language processing pipelines convert patient-generated narratives into structured data for electronic health record integration.^{50–53} A secondary benefit is confidentiality: patients can document sensitive information at their own pace rather than recounting it in a public triage area. A critical constraint applies. PF-AI should not delay safety-critical nursing triage or emergency stabilization. The practical model is to collect optional supplemental information after minimum triage safeguards are completed, then feed only validated, role-appropriate summaries into clinician workflows.⁵⁴

2.2 ED Greeting and Triage

Triage sets the stage for patient trust and departmental flow.⁵⁵ For patients able to participate, PF-AI may enhance the objectivity of triage by leveraging patient-entered data.⁴⁹ These tools can capture information often missed before arrival, addressing limits on prehospital teams who may lack interoperable systems, reliable connectivity, or time for detailed documentation.^{56,57} The information can prepopulate portions of the Emergency Severity Index assessment to reduce repetitive questioning and ease nursing workload.⁵⁸

Beyond data capture, PF-AI can transform triage into a more communicative and goal-oriented experience. Before clinician evaluation, these tools can elicit why the patient believes they are there, what worries them most, what outcomes they hope to achieve, and what values should guide their care. Capturing these priorities early allows clinicians to tailor discussions and align care plans with patient expectations, so that triage itself sets the stage for downstream flow and patient trust.

For participating patients, structured patient-entered data may improve consistency of history-taking, though any recommendations should be treated as decision support rather than autonomous triage decisions.^{55,58–60} Operationally, this requires separating patient narrative capture from acuity assignment. Patients can report goals, concerns, and barriers early, while licensed staff retain final authority over triage level, immediate testing priorities, and escalation pathways.

Transparent updates—such as current queue status and next steps—can reduce uncertainty, but messaging must avoid false precision. Wait-time estimates should, therefore, be clearly labeled as dynamic projections, and any high-risk symptom change should trigger immediate nurse reassessment rather than automated reassurance.

2.3 Posttriage and Preclinician Evaluation

Large Language Model (LLM)-enabled intake can collect richer patient narratives than generic symptom checkers by converting spoken accounts into structured fields mapped to evidence-based risk scores for clinician review.^{61–64} LLMs can adapt questioning in real time based on health literacy or language preference,^{65,66}

and organized outputs can be presented to the clinician to complete formal risk assessment and decision making.⁶⁷

Additionally, PF-AI can support preventive screenings and educational engagement, for example through personalized micro-modules on chronic disease management.^{68–70} Emergency clinicians are necessarily focused on acute complaints, stabilization, and disposition, leaving limited opportunity for consistent preventive counseling. Under clinical oversight, PF-AI can automate parts of this work and shift it to waiting periods, adding value without delaying evaluation and helping ensure that health education and screening occur equitably across patient populations.⁷¹

Reliability safeguards are mandatory in high-volume settings. Recommended controls include constrained prompts, evidence-linked decision trees, schema-bound outputs, and physician-approved escalation logic before data enter actionable pathways. Because emergency physicians cannot reasonably review all generated content, oversight should be role-stratified: frontline nurses or quality coordinators review low-risk outputs, physicians adjudicate high-risk or discordant recommendations, and medical directors oversee periodic quality audits and model updates.

For lower-acuity pathways, virtual alternatives should remain voluntary, reversible, and compliant with Emergency Medical Treatment and Labor Act requirements. The expected benefit is reduced left-without-being-seen rates. However, systems must track whether digital diversion shifts burden onto patients with fewer resources.

2.4 Postevaluation Communication

PF-AI can also support structured, patient-centered communication about new and incidental findings. When diagnosing a new condition, clinicians can trigger PF-AI-assisted modules that reinforce key counseling points, expected treatment course, and follow-up needs. For incidental findings, PF-AI can deliver consistent, calm, and guideline-concordant explanations that emphasize next steps without alarmism, potentially reducing reliance on inconsistent web searches while keeping clinicians in control of the message.⁷² These same mechanisms can be extended to consent-related education for high-stakes diagnostics or procedures, where digital modules and comprehension checks complement, but do not replace, direct physician conversation.⁷³

PF-AI can also connect clinical findings to social needs identified during intake.⁷⁴ For example, when diagnosing an illness worsened by food insecurity, the clinician can refer the patient to a social work consultation and provide interactive information about local food banks. Linking medical and social interventions within the same workflow keeps the care plan actionable and the clinician central to coordination.

2.5 Transition to Hospitalization

The decision to hospitalize, and the subsequent wait for a bed, is often confusing and stressful for patients. PF-AI can support

this transition through plain-language explanations of next steps, asynchronous question-and-answer channels for patients to raise concerns before the final discussion, and educational materials previewing what to expect during a hospital stay.^{75,76} When transfer to another facility is considered, the system can present information about the receiving hospital's resources and transportation options so that patients can make informed decisions.^{77,78} The objective at this stage is informed understanding of care pathways, not replacement of clinician counseling.^{79,80}

2.6 Discharge

Discharge is a high-risk transition in which operational pressure routinely compromises clear communication.¹⁰ Discharge instructions commonly use medical language above most patients' reading levels.^{7,10} Generative AI can rewrite standard instructions in plain language and generate patient-specific audiovisual supplements (for example, short narrated videos or translated handouts), and recent studies show substantial gains in readability and comprehension.^{7,9,10}

However, important safety limitations persist. Earlier LLMs produced discharge materials in which up to 18% of outputs contained potentially harmful inaccuracies. Newer models have reduced these risks through improved factual grounding, retrieval-augmented generation, and tighter integration with structured electronic health record data.^{7,8,81} Under clinician oversight, such models can now support a broader set of discharge tasks: context-aware tailoring of educational content, follow-up reinforcement, medication-adherence prompts, and connections to community resources.^{82,83}

Full automation at discharge remains neither feasible nor advisable because discharge communication depends on clinician judgment, clinical reasoning, and interpretation of social context. The safest model is human-in-the-loop: generative AI drafts patient instructions, and clinicians review, edit, and contextualize content before release.^{21,46,84}

2.7 Postdischarge Follow-Up

The period after discharge is one of the most vulnerable transitions in care, as patients move from a structured hospital environment to self-management at home.⁸⁵ Breakdowns in coordination during this phase often lead to medication errors, missed follow-ups, and preventable readmissions.⁸⁶ Common barriers include difficulties filling prescriptions, uncertainty about care instructions, and challenges with scheduling appointments or accessing ambulatory services due to limited transportation.^{87–90}

PF-AI may bridge these gaps through personalized, automated check-ins. AI can augment existing text message-based or telephonic follow-up interventions, which decrease ED revisits and hospital readmissions^{6,91} but can be slow, labor-intensive, and difficult to scale. PF-AI can also collect brief updates on symptoms, recovery progress, and daily function, giving clinicians a clearer view of how patients are doing after they leave the

ED.^{92,93} Educational prompts can reinforce self-management, and reports on pain, fatigue, or wound healing can be summarized for care teams to identify patients who would benefit from further intervention. By maintaining open communication, PF-AI converts passive postdischarge follow-up into an ongoing partnership to support patient adherence, safety, and confidence during recovery.

Postdischarge automation is useful only when paired with a staffed response system. Every outbound check-in should map to a defined backend action owner (for example, nurse coordinator, pharmacist, social worker, or covering clinician) and response-time targets, so identified barriers trigger real intervention rather than unresolved alerts.

2.8 Operational Governance and Information Flow

Responsible PF-AI governance rests on three interdependent requirements: preserving information integrity, assigning accountability to specific roles, and verifying performance on an ongoing basis. Information flow across every touchpoint should follow a closed-loop design in which patient input is converted into a structured summary, risk-stratified, routed for role-based review, confirmed by a licensed clinician, and documented for downstream follow-through. Governance policies should specify who may release messages to patients, who is authorized to escalate safety concerns, what must be clinician-reviewed before any patient-facing output is released, and how every override is logged for subsequent quality review.

PF-AI is most accurately framed as a high-risk socio-technical intervention rather than a discrete application deployment. Successful implementation, therefore, depends on transparent governance, explicit workflow ownership, and sustained institutional investment in safety monitoring. A recent national American College of Emergency Physicians membership survey documented meaningful institutional AI adoption among emergency medicine physicians (52%) alongside substantial uncertainty about bias (38%) and Health Insurance Portability and Accountability Act compliance (16%), underscoring the need for professional-society-backed governance norms.⁹⁴ Strategic questions for emergency clinicians to pose to health-system leaders when evaluating, governing, and sustaining PF-AI deployments are provided in [Table 2](#).

2.9 Challenges and Limitations

Workflow integration remains the single largest operational barrier. Fragmented tools generate duplicate work and contradictory documentation, whereas interoperable systems demand substantial build costs, ongoing maintenance, and standards alignment across vendors and health information exchanges. Strategic due-diligence questions for emergency clinicians to pose to PF-AI vendors are summarized in [Table 3](#).^{95,96}

TABLE 2. Strategic questions for emergency physicians to ask health-system leaders.

Category	Strategic question
Protecting the bedside	How is this PF-AI deployment being sequenced relative to bedside investments (nurse staffing, interpreter services, social work)? What guarantees ensure that PF-AI spend does not displace direct patient-care capacity?
Physician wellness	What physician-wellness metrics (documentation burden, inbox volume, after-hours chart time, burnout inventory) will be tracked pre- and postdeployment, and who is accountable for reversing course if those metrics worsen?
Opt-out pathway	What is the opt-out pathway for patients (and for clinicians) who decline PF-AI mediation, and how is that pathway operationalized without penalizing access to care or clinician time?
Continuous oversight	Which multidisciplinary governance body owns ongoing oversight (clinical quality, equity, privacy, patient experience), how often does it review performance data, and what authority does it have to pause or terminate deployment?
Kill switch	What is the documented kill-switch procedure: who has authority to immediately disable PF-AI organization-wide, under what triggers, and what is the fallback workflow while the system is paused?

PF-AI, patient-facing artificial intelligence.

Data protection obligations extend well beyond transport-layer security to full lifecycle governance: explicit consent language, granular access controls, defined retention limits, and user-facing controls that allow patients to disconnect data sources. Publicly available generative AI tools such as ChatGPT, Gemini, Llama, and Grok introduce additional Health Insurance Portability and Accountability Act exposure when used outside an institutional Business Associate Agreement; ED

clinicians should verify platform coverage and consult privacy officers before entering any protected health information.⁹⁷ Institutional trust also depends on rapid, transparent incident response when privacy events occur.

Liability remains unresolved when automated recommendations contribute to clinical action.²⁶ Institutions should predefine accountability matrices spanning vendors, health systems, supervising clinicians, and operational review teams,

TABLE 3. Strategic questions for emergency physicians to ask AI vendors.

Category	Strategic question
Confabulation guardrails	What are the specific technical guardrails against confabulation (formerly “hallucination”) in patient-facing outputs? Please describe the retrieval-augmented generation architecture, schema-bound output constraints, evidence-linked decision trees, and human-review requirements that prevent fabricated clinical content from reaching patients.
Xenomorph design	How does the patient-facing interface communicate that patients are interacting with an automated system rather than a clinician? What nonanthropomorphic design cues, disclosure language, and identity disambiguation mechanisms are embedded, and how were these tested with patients from low digital-literacy backgrounds?
Medical incapacity	Under what clinical conditions is the system designed to explicitly disengage and escalate? Please list the medical-incapacity boundaries (eg, unstable vitals, altered mental status, active psychiatric crisis, pediatric high-acuity, pregnancy, and anaphylaxis) at which patient-facing outputs are suppressed and a licensed clinician is notified.
Click fatigue test	What is the mean number of additional clicks, fields, or screens an ED clinician must interact with per patient encounter when your PF-AI module is active, compared with baseline workflow? Please share any time-motion data from pilot sites.
Human-in-the-loop mechanics	Which roles (nurse coordinator, pharmacist, social worker, physician, medical director) receive which outputs at which decision points? What are the defined response-time targets for escalated outputs, and what happens when no human responds within the target window?
Algorithmic audits	What subgroup performance data—by race/ethnicity, language, age, disability status, and insurance—are generated and externally audited on a recurring cadence? Can institutions receive site-specific audit reports, and under what confidentiality terms?
Medico-legal accountability	Where does liability reside when a patient is harmed following a PF-AI output? Please describe the contractual allocation of liability among vendor, health system, and clinician, your indemnification terms, audit-log preservation policy, and cooperation commitments during litigation or regulatory inquiry.

ED, Emergency Department; PF-AI, patient-facing artificial intelligence.

TABLE 4. Key barriers to PF-AI deployment and corresponding safeguards.

Barrier	Risk	Safeguard
Workflow integration	Fragmented tools create duplicate documentation, contradictory records, and added clinician time burden	Interoperable EHR-integrated deployment, pre-launch workflow mapping, and ongoing documentation-burden monitoring
Data protection and consent	Lifecycle data reuse may exceed patient expectations; insufficient access or retention controls	Explicit consent language, defined access controls, retention limits, and user-facing disconnect or delete options
Liability and accountability	Ambiguity when automated recommendations influence clinical action; responsibility diffuses across vendors, clinicians, and health systems	Predefined accountability matrix delineating responsibility across vendors, health systems, supervising clinicians, and end users
Safety boundary	Autonomous determination of triage acuity, diagnosis, treatment plan, or discharge of potentially unstable patients	PF-AI must not independently determine emergency triage level, final diagnosis, definitive treatment, or discharge of potentially unstable patients; licensed clinician retains final authority
Performance monitoring and equity	Model drift, subgroup performance gaps, unreported failures	Learning health-system surveillance of engagement, recommendation accuracy, override rates, clinician time burden, LWBS, and return visits, disaggregated by demographic subgroup

EHR, electronic health record; LWBS, left without being seen; PF-AI, patient-facing artificial intelligence.

with auditable logs for every high-risk recommendation and override decision. For high-risk recommendations—those affecting triage acuity, differential diagnosis in potentially unstable patients, or discharge of patients at risk for clinical deterioration—physician adjudication should be the designated endpoint of accountability.

A practical safety boundary is that PF-AI should not independently determine emergency triage level, final diagnosis, definitive treatment plans, or discharge of potentially unstable patients. Its role is to structure information, support communication, and surface concerns for licensed clinician decision making. This boundary is especially important given documented pre-existing disparities in human-assigned ED triage and disposition decisions, which any AI tool would

otherwise inherit or amplify.^{98,99} Physician adjudication should be required for high-risk recommendations and safety overrides.

As part of a learning health system, programs should monitor process and outcome metrics: engagement, recommendation accuracy, override rates, clinician time burden, left-without-being-seen, return visits, follow-up completion, patient complaints, adverse events, and subgroup performance by language, literacy, race/ethnicity, insurance status, age, and digital access.^{28–30,100}

Table 4 organizes these challenges as paired barrier-risk-safeguard triads spanning workflow integration, data protection, liability, safety boundaries, and performance monitoring. Equitable deployment, however, requires more than internal

TABLE 5. Strategic questions for co-designing PF-AI with patients and community-based organizations.

Category	Strategic question
Phase 1 — authentic partnerships	Which patient communities most affected by this PF-AI tool have been identified, and how were they identified? Which community-based organizations have been engaged as partners before design began, and what is the compensation structure for their time? How will power asymmetries between the health system and patient partners be surfaced and mitigated throughout the engagement?
Phase 2 — participatory codesign workgroups	What mechanisms allow patient and community partners to shape tool behavior, user interface, and escalation logic — not just provide reactive feedback? How are dissenting patient perspectives documented and escalated when they conflict with clinical or operational priorities? What is the cadence and structure of codesign sessions across the deployment lifecycle (pre-pilot, pilot, go-live, postdeployment)?
Phase 3 — shared accountability metrics	Which equity-relevant outcomes (access, concordance, disparities in accuracy) are monitored, and who sees the dashboards? What is the formal process for patient/community partners to trigger mitigation, pause, or sunset of a PF-AI tool? How are results of shared accountability monitoring reported back to participating communities in plain language?

PF-AI, patient-facing artificial intelligence.

safeguards—it requires sustained codesign with the communities PF-AI will most affect. Drawing on community-based participatory research methodology, we frame this as a three-phase commitment. First, authentic partnerships are established before design begins: institutions identify patient communities most affected by the tool, engage community-based organizations as compensated partners rather than reactive reviewers, and surface the power asymmetries inherent in health-system-led development. Second, participatory codesign workgroups empower patient and community partners to shape tool behavior, interface design, and escalation logic across prepilot, pilot, go-live, and postdeployment phases, with explicit mechanisms to document and escalate dissenting perspectives when they conflict with clinical or operational priorities. Third, shared accountability metrics track equity-relevant outcomes—access, concordance, disparities in accuracy—on dashboards visible to community partners, who retain formal authority to trigger mitigation, pause, or sunset of a deployed tool and who receive plain-language reports on monitoring results. [Table 5](#) provides strategic questions to operationalize each phase.

3 CONCLUSION

The ED is an important but operationally demanding setting for responsible PF-AI. The strongest near-term opportunities lie in structured data transformation, multimodal communication support, and workflow augmentation—each implemented under explicit human oversight.

Adoption should proceed cautiously, with continued emphasis on reliability verification, equity protection, interoperability, cost sustainability, and legal accountability. Under this framework, PF-AI functions as an instrument for safer communication and coordination, although clinicians and care teams remain accountable for clinical decisions.

FUNDING AND SUPPORT

By *JACEP Open* policy, all authors are required to disclose any and all commercial, financial, and other relationships in any way related to the subject of this article as per ICMJE conflict of interest guidelines (see www.icmje.org). The authors have stated that no such relationships exist.

DECLARATION OF GENERATIVE AI AND AI-ASSISTED TECHNOLOGIES IN THE MANUSCRIPT PREPARATION PROCESS

During the preparation of this work the authors used Anthropic Claude (Sonnet 4.5) and Google Gemini large language model tools, together with Model Context Protocol connectors to PubMed and [ClinicalTrials.gov](https://clinicaltrials.gov), in order to retrieve 2025 to 2026 peer-reviewed literature aligned to reviewer comments, draft structured candidate insertion text for iterative first-author review, assist in preparing tracked-change revisions, and help structure the response-to-reviewers document. After using these tools/services, the

authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

ACKNOWLEDGMENTS

The authors acknowledge the American College of Emergency Physicians Artificial Intelligence Task Force for their support and contributions to this work.

ORCID

Tehreem Rehman MD, MPH  <https://orcid.org/0000-0002-3134-1313>

REFERENCES

1. Cairns CKK. National hospital ambulatory medical care survey: 2022 emergency department summary tables. 2022. Accessed July 1, 2026. https://www.cdc.gov/nchs/data/nhamcs/web_tables/2022-nhamcs-ed-web-tables.pdf
2. Bledsoe B, Wasden C, Johnson L. Electronic prehospital records are often unavailable for emergency department medical decision making. *West J Emerg Med*. 2013;14:482-488.
3. Damen DJ, Schoonman GG, Maat B, Habibović M, Krahmer E, Pauws S. Patients managing their medical data in personal electronic health records: scoping review. *J Med Internet Res*. 2022;24(12):e37783.
4. Kern LM, Safford MM, Slavin MJ, et al. Patients' and providers' views on causes and consequences of healthcare fragmentation in the ambulatory setting: a qualitative study. *J Gen Intern Med*. 2019;34:899-907.
5. Naved BA, Ravishankar S, Colbert GE, Johnston A, Slott QM, Luo Y. LLM enabled classification of patient self-reported symptoms and needs in health systems across the USA. *NPJ Digit Med*. 2025;8(1):390.
6. Lee G, Bruce C, Nisar T, Holderread B, Pletcher S, Nguyen NA. Implementation and evaluation of a virtual transitional care intervention using automated text messaging and virtual visits after emergency department discharges: retrospective cohort study. *JMIR Mhealth Uhealth*. 2025;13:e77973.
7. Zaretsky J, Kim JM, Baskharoun S, et al. Generative artificial intelligence to transform inpatient discharge summaries to patient-friendly language and format. *JAMA Netw Open*. 2024;7(3):e240357.
8. Stanceski K, Zhong S, Zhang X, et al. The quality and safety of using generative AI to produce patient-centred discharge instructions. *NPJ Digit Med*. 2024;7(1):329.
9. Ray M, Kats DJ, Moorkens J, et al. Evaluating a large language model in translating patient instructions to Spanish using a standardized framework. *JAMA Pediatrics*. 2025;179(9):1026-1033.
10. Huang T, Safranek C, Socrates V, et al. Patient-representing population's perceptions of gpt-generated versus standard emergency department discharge instructions: randomized blind survey assessment. *J Med Internet Res*. 2024;26:e60336.
11. OpenAI. Introducing ChatGPT health. Published January 7, 2026. Accessed March 17, 2026. <https://openai.com/index/introducing-chatgpt-health/>
12. Microsoft AI. Introducing copilot health. Published March 12, 2026. Accessed March 17, 2026. <https://microsoft.ai/news/introducing-copilot-health/>
13. Ramaswamy A, Tyagi A, Hugo H, et al. ChatGPT health performance in a structured test of triage recommendations. *Nat Med*. 2026;32(5):1671-1675. <http://doi.org/10.1038/s41591-026-04297-7>
14. Amazon Web Services. Introducing Amazon connect health, agentic AI built for healthcare. Published March 2026. Accessed March 17,

2026. <https://aws.amazon.com/about-aws/whats-new/2026/03/amazon-connect-health-agent-ai-healthcare/>
15. Lambe R, Baldwin M, O'Grady B, Schumann M, Caulfield B, Doherty C. The accuracy of Apple watch measurements: a living systematic review and meta-analysis. *NPJ Digit Med.* 2026;9(1):63.
16. Kim H, Jo S, Lim MH, Choi DH. From non-agent large language models to multi-agent systems in emergency medicine: a scoping review. *Clin Exp Emerg Med.* 2026; <http://doi.org/10.15441/ceem.26.136>
17. Thirunavukarasu AJ, Ting DSJ, Elangovan K, Gutierrez L, Tan TF, Ting DSW. Large language models in medicine. *Nat Med.* 2023;29(8):1930-1940.
18. Nazi ZA, Peng W. Large language models in healthcare and medical domain: a review. *Informatics.* 2024;11(3):57.
19. Taylor RA, Sangal RB, Smith ME, et al. Leveraging artificial intelligence to reduce diagnostic errors in emergency medicine: challenges, opportunities, and future directions. *Acad Emerg Med.* 2025;32(3):327-339.
20. Omiye JA, Lester JC, Spichak S, Rotemberg V, Daneshjou R. Large language models propagate race-based medicine. *NPJ Digit Med.* 2023;6(1):195.
21. Silberlust J, Genes N. Clinical integration of local ai assistants into emergency department discharge documentation. *JAMA Netw Open.* 2025;8(10):e2538437.
22. Ancker JS, Edwards A, Nosal S, et al. Effects of workload, work complexity, and repeated alerts on alert fatigue in a clinical decision support system. *BMC Med Inform Decis Mak.* 2017;17(1):36.
23. Mejia A, Granof M, Mikkelsen J, et al. Bridging the digital divide with accessible digital literacy tools, training, and tests: sharing the DOORS program. *NEJM Catalyst.* 2025;6(12).
24. Luccioni S, Jernite Y, Strubell E. Power hungry processing: watts driving the cost of AI deployment. *Paper presented at: Proceedings of the 2024 ACM Conference on fairness, accountability, and transparency.* New York, NY: Association for Computing Machinery; 2024.
25. Matheny ME, Goldsack JC, Saria S, et al. Artificial intelligence in health and health care: priorities for action. *Health Aff (Millwood).* 2025;44(2):163-170.
26. Price WN 2nd, Gerke S, Cohen IG. Potential liability for physicians using artificial intelligence. *JAMA.* 2019;322(18):1765-1766.
27. WHO. Ethics and governance of artificial intelligence for health. Geneva, 2021. Accessed July 1, 2026. <https://www.who.int/publications/i/item/9789240029200>
28. Collins GS, Moons KGM, Dhiman P, et al. TRIPOD+AI statement: updated guidance for reporting clinical prediction models that use regression or machine learning methods. *BMJ.* 2024;385:e078378.
29. Liu X, Glocker B, McCradden MM, Ghassemi M, Denniston AK, Oakden-Rayner L. The medical algorithmic audit. *Lancet Digit Health.* 2022;4(5):e384-e397.
30. Gallifant J, Afshar M, Ameen S, et al. The TRIPOD-LLM reporting guideline for studies using large language models. *Nat Med.* 2025;31(1):60-69.
31. Kachman MM, Brennan I, Oskvarek JJ, Waseem T, Pines JM. How artificial intelligence could transform emergency care. *Am J Emerg Med.* 2024;81:40-46.
32. Seo W, Jain S, Le V, et al. Designing patient-centered interventions for emergency care: participatory design study. *JMIR Form Res.* 2025;9:e63610.
33. Ly S, Tsang R, Ho K. Patient perspectives on the digitization of personal health information in the emergency department: mixed methods study during the COVID-19 pandemic. *JMIR Med Inform.* 2022;10(1):e28981.
34. Zikmund-Fisher BJ, Thorpe A, Fagerlin A. How to communicate medical numbers. *JAMA.* 2025;334(16):1474-1475.
35. Hua Y, Xia W, Bates D, et al. Standardizing and scaffolding health care ai-chatbot evaluation: systematic review. *JMIR AI.* 2025;4:e69006.
36. Griffin AC, Javier SJ, Golding M, et al. Designing patient-friendly messages: tutorial on applying human-centered, self-determination theory with ai considerations. *J Med Internet Res.* 2025;27:e78173.
37. Khamaj A. AI-enhanced chatbot for improving healthcare usability and accessibility for older adults. *Alex Eng J.* 2025;116:202-213.
38. Park C, Park J, Kim S, et al. Explainable machine learning for assessing digital health literacy in older adults: validation and development of a two-stage model integrating performance-based and self-assessed indicators. *JMIR Med Inform.* 2026;14:e86171. <http://doi.org/10.2196/86171>
39. Milford SR, Herger E, Eichinger J, Elger BS. Promoting xenomorphic patient-facing AIs: the case against anthropomorphism in medical AIs. *NPJ Digit Med.* 2025;8(1).
40. Neprash HT, McGlave CC, Cross DA, et al. Trends in ransomware attacks on us hospitals, clinics, and other health care delivery organizations, 2016-2021. *JAMA Health Forum.* 2022;3(12):e224873. <http://doi.org/10.1001/jamahealthforum.2022.4873>
41. Dameff C, Tully J, Chan TC, et al. Ransomware attack associated with disruptions at adjacent emergency departments in the US. *JAMA Netw Open.* 2023;6(5):e2312270. <http://doi.org/10.1001/jamanetworkopen.2023.12270>
42. Ulizio CW, Dua D, Mukul NM, Areti S, Kostick-Quenet K, Rahimzadeh VN. Exploring the ethical and practical considerations of artificial intelligence in real-world health care settings: stakeholder focus group study. *JMIR AI.* 2026;5:e85163. <http://doi.org/10.2196/85163>
43. Johnson JA, Moore B, Hwang EK, Hickner A, Yeo H. The accuracy of race & ethnicity data in US based healthcare databases: a systematic review. *Am J Surg.* 2023;226(4):463-470.
44. Higashi RT, Thakur B, Repasky EC, et al. Digital health technology use among Spanish speakers in the US: a scoping review. *JAMA Netw Open.* 2025;8(5):e2510386.
45. Abbott EE, Rehman T, Rosania A, et al. Understanding and addressing bias in artificial intelligence systems: a primer for the emergency medicine physician. *J Am Coll Emerg Physicians Open.* 2025;7(1):100311. <http://doi.org/10.1016/j.acepjo.2025.100311>
46. Griot MF, Walker GA. A patient-in-the-loop approach to artificial intelligence in medicine. *JAMA Netw Open.* 2025;8(6):e2514460.
47. Kim C, Gadgil SU, Lee S-I. Transparency of medical artificial intelligence systems. *Nat Rev Bioeng.* 2025;4(1):11-29.
48. Nasser L, Morris E, Mathias I, Hall JN. Considerations for emergency department virtual triage. *Health Manage Forum.* 2025;38(2):108-113.
49. Wang P, Yu L, Li T, Zhou L, Ma X. Use of mobile technologies to streamline pretriage patient flow in the emergency department: observational usability study. *JMIR Mhealth Uhealth.* 2024;12:e54642.
50. Scheder-Bieschin J, Blümke B, De Buijzer E, et al. Improving emergency department patient-physician conversation through an artificial intelligence symptom-taking tool: mixed methods pilot observational study. *JMIR Form Res.* 2022;6(2):e28199.
51. Fraser HSF, Cohan G, Koehler C, et al. Evaluation of diagnostic and triage accuracy and usability of a symptom checker in an emergency department: observational study. *JMIR Mhealth Uhealth.* 2022;10(9):e38364.
52. Meer A, Rahm P, Schwendinger M, et al. A symptom-checker for adult patients visiting an interdisciplinary emergency care center and the safety of patient self-triage: real-life prospective evaluation. *J Med Internet Res.* 2024;26:e58157.
53. Brice JH, Friend KD, Delbridge TR. Accuracy of EMS-recorded patient demographic data. *Prehosp Emerg Care.* 2008;12:187-191.
54. Sim JA, Huang X, Horan MR, et al. Natural language processing with machine learning methods to analyze unstructured patient-reported outcomes derived from electronic health records: a systematic review. *Artif Intell Med.* 2023;146:102701.
55. Johnson KD, Lindsell CJ, Froehle C, Gillespie GL. Interruption of initial patient assessment in the emergency department and its effect on patient perception of care quality. *Int J Qual Health Care.* 2021;33(4):mzab146.

56. Carter AJ, Davis KA, Evans LV, Cone DC. Information loss in emergency medical services handover of trauma patients. *Prehosp Emerg Care*. 2009;13(3):280-285.
57. Landman AB, Lee CH, Sasson C, Van Gelder CM, Curry LA. Prehospital electronic patient care report systems: early experiences from emergency medical services agency leaders. *PLoS One*. 2012;7(3):e32692.
58. Tyler S, Olis M, Aust N, et al. Use of artificial intelligence in triage in hospital emergency departments: a scoping review. *Cureus*. 2024;16(5):e59906.
59. Fraser H, Crossland D, Bacher I, Ranney M, Madsen T, Hilliard R. Comparison of diagnostic and triage accuracy of Ada Health and WebMD symptom checkers, ChatGPT, and physicians for patients in an emergency department: clinical data analysis study. *JMIR Mhealth Uhealth*. 2023;11:e49995.
60. Ilicki J, Edman S, Stalfors J, Molin CJ. Evaluating digital triage symptom checker with historical triage-related adverse events. *Scand J Prim Health Care*. 2025;44(1):1-14.
61. Ayers JW, Poliak A, Dredze M, et al. Comparing physician and artificial intelligence chatbot responses to patient questions posted to a public social media forum. *JAMA Intern Med*. 2023;183(6):589.
62. Rising KL, Padrez KA, O'Brien M, Hollander JE, Carr BG, Shea JA. Return visits to the emergency department: the patient perspective. *Ann Emerg Med*. 2015;65(4):377-386.e3.
63. Deleger L, Brodzinski H, Zhai H, et al. Developing and evaluating an automated appendicitis risk stratification algorithm for pediatric patients in the emergency department. *J Am Med Inform Assoc*. 2013;20(e2):e212-e220.
64. Sterling NW, Patzer RE, Di M, Schrager JD. Prediction of emergency department patient disposition based on natural language processing of triage notes. *Int J Med Inform*. 2019;129:184-188.
65. Lyu Q, Tan J, Zapadka ME, et al. Translating radiology reports into plain language using ChatGPT and GPT-4 with prompt learning: results, limitations, and potential. *Vis Comput Ind Biomed Art*. 2023;6(1):9.
66. Jeblick K, Schachtner B, Dext J, et al. ChatGPT makes medicine easy to swallow: an exploratory case study on simplified radiology reports. *Eur Radiol*. 2023;34(5):2817-2825.
67. Tierney AA, Gayre G, Hoberman B, et al. Ambient artificial intelligence scribes to alleviate the burden of clinical documentation. *NEJM Catalyst*. 2024;5(3).
68. Brown JL, Gause NK, Braun R, et al. Substance use and mental health screening within an emergency department-based HIV screening program: outcomes from 1 year of implementation. *Health Promot Pract*. 2024;25(6):1049-1057.
69. Kene M, Miller Rosales C, Wood S, Rauchwerger AS, Vinson DR, Sterling SA. Feasibility of expanded emergency department screening for behavioral health problems. *Am J Manag Care*. 2018;24(12):585-591.
70. Fein JA, Pailier ME, Barg FK, et al. Feasibility and effects of a web-based adolescent psychiatric assessment administered by clinical staff in the pediatric emergency department. *Arch Pediatr Adolesc Med*. 2010;164(12):1112-1117.
71. McDonald CE, Voutier C, Govil D, et al. Do health service waiting areas contribute to the health literacy of consumers? A scoping review. *Health Promot Int*. 2023;38(4):daad046.
72. Woo KMC, Simon GW, Akindutire O, et al. Evaluation of GPT-4 ability to identify and generate patient instructions for actionable incidental radiology findings. *J Am Med Inform Assoc*. 2024;31(9):1983-1993.
73. Edwards C, Abujarad F. Digital informed consent for older adults in emergency department research. *Appl Hum Factors Ergon Conf*. 2023;78. <http://doi.org/10.54941/ahfe1003449>
74. Kocielnik R, Agapie E, Argyle A, et al. HarborBot: a chatbot for social needs screening. *AMIA Annu Symp Proc*. 2019;2019:552-561.
75. Han CY, Lin CC, Goopy S, Hsiao YC, Barnard A, Wang LH. Waiting and hoping: a phenomenographic study of the experiences of boarded patients in the emergency department. *J Clin Nurs*. 2017;26(5-6):840-848.
76. Pilet C, Tandzi-Tonleu F, Lagarde E, Gil-Jardiné C, Galinski M, Lafont S. Feelings of patients admitted to the emergency department. *Healthcare (Basel)*. 2025;13(5):500.
77. Bai X, Wang S, Zhao Y, Feng M, Ma W, Liu X. Application of AI chatbot in responding to asynchronous text-based messages from patients with cancer: comparative study. *J Med Internet Res*. 2025;27:e67462.
78. Abbasgholizadeh Rahimi S, Cwintal M, Huang Y, et al. Application of artificial intelligence in shared decision making: scoping review. *JMIR Med Inform*. 2022;10(8):e36199.
79. Smalley CM, Simon EL, Meldon SW, et al. The impact of hospital boarding on the emergency department waiting room. *J Am Coll Emerg Physicians Open*. 2020;1(5):1052-1059.
80. Thomas ACR, Giroux EE, Soril LJJ, Sauro KM. Electronic health (eHealth) and artificial intelligence-based tools to optimize in-hospital patient flow: a scoping review. *J Patient Saf*. 2025;21(6):409-423.
81. Williams CYK, Bains J, Tang T, et al. Evaluating large language models for drafting emergency department encounter summaries. *PLOS Digit Health*. 2025;4(6):e0000899.
82. Amugongo LM, Mascheroni P, Brooks S, Doering S, Seidel J. Retrieval augmented generation for large language models in healthcare: a systematic review. *PLOS Digit Health*. 2025;4(6):e0000877.
83. Zakka C, Shad R, Chaurasia A, et al. Almanac — retrieval-augmented language models for clinical medicine. *NEJM AI*. 2024;1(2). <http://doi.org/10.1056/aioa2300068>
84. Song JW, Park J, Kim JH, You SC. Large language model assistant for emergency department discharge documentation. *JAMA Netw Open*. 2025;8(10):e2538427.
85. Forster AJ, Clark HD, Menard A, et al. Adverse events among medical patients after discharge from hospital. *CMAJ*. 2004;170(3):345-349.
86. Leppin AL, Gionfriddo MR, Kessler M, et al. Preventing 30-day hospital readmissions : a systematic review and meta-analysis of randomized trials. *JAMA Intern Med*. 2014;174(7):1095.
87. Arora VM, Prochaska ML, Farnan JM, et al. Problems after discharge and understanding of communication with their primary care physicians among hospitalized seniors: a mixed methods study. *J Hosp Med*. 2010;5(7):385-391.
88. Holland L, Nelson ML, Westrich K, Campbell PJ, Pickering MK. The patient's medication access journey: a conceptual framework focused beyond adherence. *J Manag Care Spec Pharm*. 2021;27(12):1627-1635.
89. Syed ST, Gerber BS, Sharp LK. Traveling towards disease: transportation barriers to health care access. *J Community Health*. 2013;38(5):976-993.
90. Virapongse A, Misky GJ. Self-identified social determinants of health during transitions of care in the medically underserved: a narrative review. *J Gen Intern Med*. 2018;33(11):1959-1967.
91. Harrison PL, Hara PA, Pope JE, Young MC, Rula EY. The impact of postdischarge telephonic follow-up on hospital readmissions. *Popul Health Manag*. 2011;14(1):27-32.
92. Tsang C, Lee KS, Richards H, Blazeby JM, Avery KNL. Electronic collection of patient-reported outcomes following discharge after surgery: systematic review. *BJS Open*. 2021;5(2):zraa072.
93. Ahn A, Morgan AU, Burke RE, et al. Postdischarge needs identified by an automated text messaging program: a mixed-methods study. *J Hosp Med*. 2024;19(12):1138-1146.
94. Shy BD, Baloesu C, Faustino IV, et al. Early insights among emergency medicine physicians on artificial intelligence: a national,

- convenience-sample survey of the American College Of Emergency Physicians. *J Am Coll Emerg Physicians Open*. 2025;7(1):100308. <http://doi.org/10.1016/j.acepjo.2025.100308>
95. Oskrochi Y, Roy-Highley E, Grimes K, Shah S. Digital health technology compliance with clinical safety standards in the national health service in England: national cross-sectional study. *J Med Internet Res*. 2025;27:e80076.
 96. Abdalnour R-EE, Gin B, Boscardin CK. Educational strategies for clinical supervision of artificial intelligence use. *New Engl J Med*. 2025;393(8):786-797.
 97. Schoolcraft D, Meltzer AC, Sangal R, et al. Health Insurance Portability and Accountability Act liability in the age of generative artificial intelligence. *J Am Coll Emerg Physicians Open*. 2026;7(2): 100317. <http://doi.org/10.1016/j.acepjo.2025.100317>
 98. Patel MD, Lin P, Cheng Q, et al. Patient sex, racial and ethnic disparities in emergency department triage: a multi-site retrospective study. *Am J Emerg Med*. 2023;76:29-35. <http://doi.org/10.1016/j.ajem.2023.11.008>
 99. Çalışkan YK, Başak F, Erdem O. Can AI safely choose antibiotics over the knife? A STROBE-guided benchmark of GPT-4, GPT-5, and Gemini for non-operative acute appendicitis management. *Int J Med Inform*. 2026;213:106389. <http://doi.org/10.1016/j.ijmedinf.2026.106389>
 100. Wilbourne P, Mirch-Kretschmann S, Walker D, Varghese M, Arnetoli R. AI-enabled, text-based health coaching and navigation for employees to support health outcomes: pre-post observational study. *JMIR Form Res*. 2025;9:e64553.

How to cite this article: Rehman T, Jarrett P, Lesko J, et al. An Evidence-Based Framework for Patient-Facing Artificial Intelligence Integration in the Emergency Department. *JACEP Open*. 2026;7:100466.

<https://doi.org/10.1016/j.acepjo.2026.100466>