OPINION/PERSPECTIVE

Beyond the algorithm: a six-phase framework for clinical workflow design in the deployment of artificial intelligence within health systems

Hasan Nawaz Tahir[1]

Author affiliations

[1]FCPS (Community Medicine), Assistant Professor of Community Medicine, Shaqra University, Saudi Arabia

To cite: Tahir HN. Beyond the algorithm: a six-phase framework for clinical work-flow design in the deployment of artificial intelligence within health systems. Health Innov J Eur. 2026; 1:1-8.

Received [24 September 2026]

Accepted [30 October 2026]

© Author(s) (or their employer(s)) 2026. Re- use permitted under CC BY- NC. Published by Health Innovation Journal Europe.

Correspondence to

Dr Hasan Nawaz Tahir;
hasan.na-waz@su.edu.sa
ORCID iD: 0000-0002-8710-3314

ABSTRACT

Background: The widespread adoption of artificial intelligence (AI) in clinical settings has been accompanied by a paradox: technically sound models are often deployed into health systems and underperform or are abandoned. In the author's experience, the failure is more often structural—rooted in the absence of deliberate clinical workflow design—than purely algorithmic.

Objective: To propose a six-phase framework—Discover, Map, AI Fit, Redesign, Govern, and Deploy—for the systematic design of clinical workflows in AI deployment contexts. The framework is grounded in the author's direct experience across WHO-funded projects, government health-system implementations, and digital decision-support deployments in low- and middle-income settings, and is offered as a field-tested, practical approach rather than a formally validated instrument.

Methods: This perspective synthesizes practitioner experience from multiple AI and digital-health deployments, including a WHO Alliance-funded digital decision support system (DDSS) embedded within Pakistan's largest telemedicine platform, GIS-based vaccination coverage mapping, and hospital digitization projects in South Asia and Saudi Arabia. The framework was developed through iterative reflection on field notes, co-design sessions, post-deployment debriefs, and a selective reading of the clinical AI implementation and human-factors literature.

Conclusion: Clinical AI deployment is primarily a change-management challenge that involves technology. The six-phase framework presented here offers a structured, field-tested approach for clinicians and health-system leaders tasked with introducing AI into real-world care settings.

Keywords: clinical workflow design; artificial intelligence deployment; digital health; health systems; change management; clinical informatics; low- and middle-income countries

1. INTRODUCTION

Models have achieved or surpassed human-level performance in diabetic retinopathy screening,1 sepsis treatment strategy selection,2 and, more broadly, across imaging and clinical decision domains as re-viewed by Topol.3 Colleagues and I have also documented performance of AI for Helicobacter pylori detection from upper gastrointestinal endoscopy im-ages.4 The evidence base continues to grow.

Yet the deployment gap persists. AI tools that per-form well in validation studies routinely underperform or are abandoned in clinical practice. Implementation-science literature attributes this to work-flow disruption, clinician resistance, inadequate governance, absent accountability structures, and mis-match between training data and live clinical envi-ronments.5 In a critical appraisal I co-authored, we found that one of the more dangerous failure modes was model confidence in the presence of inaccuracy.6 Clinicians who trust a system without understanding its failure modes can be more vulnerable than those with no AI support at all.

Over six years I have deployed AI and digital-health tools inside health systems—from digitization of patient registration at a tertiary hospital in Karachi, to a WHO-funded digital decision support system (DDSS) within Pakistan's largest telemedicine plat-form, to GIS-based vaccination coverage mapping for high-risk urban communities. In each case the technical component was the smaller challenge. The dominant challenges were human and organizational: understanding the actual workflow, earning clinical trust, building accountability structures that protect patients when the tool is wrong, and sustaining adoption after go-live.

This perspective proposes a six-phase framework for clinical workflow design in AI deployment, drawn from those deployments and from the implementation-science and clinical-informatics literature. It is an attempt to systematize lessons—including failures—from settings in which the consequences of design error fall on vulnerable patients. The frame-work is offered as field-tested practitioner guidance, not as a formally validated instrument.

The hardest problem in clinical AI is rarely the model; it is the deployment. In the author's experience the model typically accounts for a minority of the overall effort; the majority is clinical, organizational, and human.

2. THE PROBLEM: WHY AI FAILS AFTER THE MODEL SUCCEEDS

A common misconception frames deployment as the second phase of a two-phase process whose first phase is model development and validation. Under that framing, deployment is treated as implementation—getting a validated model into the clinical set-ting. The premise of the present framework is different: deployment is itself a clinical design problem, and the absence of deliberate workflow design is a primary driver of failure.

This view is consistent with accumulating evidence. Inadequate workflow integration has been identified as a leading cause of adoption failure for health in-formation technology.7 In our own DDSS deployment within the Sehat Kahani telemedicine platform (a WHO Alliance AHPSR-funded study), the tool performed as intended in usability testing yet required substantial workflow redesign to function in the live environment.8 The standard operating procedure for Lady Health Worker (LHW) consultations and the actual practice of those consultations differed meaningfully. Designing only for the written proto-col would have produced a tool that clinicians routed around. Three structural problems characterize many clinical AI deployments that fail after the model succeeds.

First, the existing workflow is not adequately mapped before the AI is introduced. The tool is added to a process that has not been fully understood, generating friction, workarounds, and clinicians who bypass the AI in ways invisible to the deployment team.

Second, every decision node is treated as an AI opportunity. Not every clinical decision is appropriate for augmentation or automation. Criteria for appropriate insertion—high volume, structured data avail-ability, recoverable error consequences, existing performance inconsistency, and acceptable latency—are rarely applied rigorously.

Third, governance is treated as a compliance exercise rather than a design problem. Accountability (who is responsible when the AI contributes to harm), monitoring (how drift will be detected), and failure-mode response protocols are not defined in advance. When problems surface, the response is reactive and trust is damaged. The six-phase framework addresses these problems through sequential clinical workflow design.

3. THE SIX-PHASE FRAMEWORK

The framework is sequential. Each phase produces outputs required by the next. Skipping a phase does not save time; it produces problems that surface un-der operational pressure and are costlier to remedy.

Table 1 Overview of the six phases

Phase	Primary focus	Key inputs	Key outputs	Typical roles involved	Approximate duration (focused deployment)
1. Discover	Understand the actual current work-flow	Direct observation, arte-facts, time-volume data, pain-point interviews	Observation notes, artefact library, time-volume baseline, prioritized pain points	External facilitator + front-line clinicians of every role	Days to 1–2 weeks
2. Map	Make the work-flow visible	Outputs of Phase 1	Current-state swim lane diagram with decision nodes marked	Clinical team (primary authors) + facilitator	1–3 workshops
3. AI Fit	Decide where AI belongs	Current-state map + five fit criteria	Selected decision nodes, autonomy level, latency specification	Clinical leads + technical leads	1–2 workshops
4. Redesign	Build the future-state workflow	Selected nodes + design principles	Future-state swim lane (AI as its own lane), paper proto-types, case walkthroughs	Clinical team + design facilitator	1–2 weeks
5. Govern	Define accountability before go-live	Future-state map	Accountability matrix, monitoring thresholds, failure-mode catalogue, regulatory pathway	Clinical leadership + governance/ethics	Parallel with redesign
6. Deploy	Earn and sustain clinical adoption	All prior outputs	Controlled pilot, champion activation, training that includes failure modes, weekly case review	Clinical champions + full team	Pilot ≥ weeks; active support ≥ 90 days

Phase 1: Discover — Understanding the Workflow That Actually Exists. Direct observation of the clinical environment pre-cedes any design work. The difference between writ-ten standard operating procedures and what staff actually do is often the difference between a tool that is used and one that is routed around.

In the Sehat Kahani DDSS deployment, observation of LHW consultations showed that workers managed cognitive load by applying abbreviated decision heuristics rather than the full referral criteria in their training manuals—not because of inadequate training, but because consultation volume and time constraints made full protocol adherence impractical. A DDSS designed solely around the written protocol would have increased cognitive load; one designed around observed heuristics could reduce it.

Four activities characterize this phase: (1) direct observation of every role that touches the workflow; (2) collection of every form, checklist, referral template and screen actually used; (3) a simple time-and-volume study; and (4) structured interviews that ask "What is the hardest part of your day?" rather than "What would you like AI to do?" The latter question tends to elicit answers shaped by popular conceptions of AI rather than by real constraints.

Phase 2: Map — Making the Workflow Visible. Observations are converted into a swim lane diagram in which each horizontal lane represents one role, rectangles represent actions, diamonds represent decisions, and arrows that cross lanes represent handoffs. Handoffs are a primary locus of error in complex clinical workflows.[9]

In the vaccination-coverage mapping project for high-risk communities in Karachi, the mapping phase revealed that immunization data moved from paper records kept by vaccinators, through transcription by supervisors, to digital entry by administrative staff—three handoffs before analysis. The GIS tool assumed clean, timely digital data; the map showed that data-collection redesign was required first.

The critical principle is co-production: the map is produced *with* the clinical team, not *for* them. A map the team authors is a map the team owns; ownership is a precondition for adoption. Every decision node is marked as a potential AI-insertion candidate so that the next phase can evaluate it systematically.

Phase 3: AI Fit — Deciding Where AI Belongs. This phase applies explicit criteria to determine which decision nodes are appropriate candidates for AI augmentation or automation. Five conditions are proposed as a practical filter (they are not claimed to be exhaustive or universally necessary):

- The decision is high-volume and repetitive, with sufficient structural similarity for a model to generalize.
- The required data exist in structured form at the point the decision is made.
- The consequence of an AI error is recoverable by a human clinician before irreversible harm.
- The decision is currently made with meaningful inter-clinician variation, so a reliable signal can improve rather than merely replicate performance.
- The workflow can accommodate the system's response latency.

The fifth condition is frequently overlooked. In emergency triage a model that requires several seconds may be clinically unusable regardless of accuracy.

Autonomy level must also be specified. Three levels are distinguished:

- Level 1: the AI surfaces relevant information (guidelines, risk scores, prior results) without a recommendation.
- Level 2: the AI produces a recommendation that the clinician reviews, accepts or overrides; the clinician remains accountable.

- Level 3: the AI acts with human review of logs rather than of each individual decision.

Almost all clinical AI deployments should begin at Level 1 or 2. Level 3 is appropriate only for low-risk, high-volume, highly routine tasks after demonstrated reliability at Level 2.

Automation bias—the tendency to accept AI recommendations without critical evaluation—is a recognized risk.10 The design response is to make the over-ride pathway as carefully engineered as the recommendation pathway. If overriding takes longer or re-quires more effort than accepting, bias is structurally built into the deployment.

Phase 4: Redesign — Building the Future-State Workflow. The future-state workflow is documented as a new swim lane in which the AI system appears as its own lane. Data inputs, outputs, the clinician who receives the output, and the actions that follow are made explicit. Every handoff to and from the AI, and every accountability assignment, is documented.

Three design principles guide this phase:

1. The future-state workflow must be faster or less cognitively demanding than the current state for the roles affected. In the DDSS work we used measured consultation time as a hard ceiling.
2. The future-state must be paper-prototyped before code is written. In a hospital registration digitization project in Korangi, Karachi, paper prototyping revealed that a proposed digital form required 14 fields the paper form did not; redesign reduced the required fields to six. The change, discovered in a two-hour session, would have taken weeks to correct after go-live.
3. The future-state map is walked through with the clinical team using real historical cases, including edge cases. This is both quality assurance and the single most important early trust-building activity.

Phase 5: Govern — Defining Accountability Before Go-Live. Governance is treated as a design activity that answers what happens when the AI is wrong. Four components are specified:

1. An accountability matrix: for every AI decision node a named human clinician is designated as clinically accountable for the outcome, regardless of whether the recommendation was accepted or overridden.
2. A performance-monitoring framework with pre-specified metrics (e.g., referral rate, override rate, consultation time, selected clinical outcomes) and thresholds at which deployment will be paused for clinical review. Defining thresholds while conditions are stable avoids decision-making under cri-sis pressure.
3. A failure-mode catalogue written in clinical language, covering at minimum high-confidence error, distribution shift, edge-case presentation, and (for generative systems) hallucination, each with a detection mechanism and response protocol.
4. Identification of the applicable regulatory and ethics pathway before design is finalized. In the DDSS project this included Aga Khan University Ethics Review Committee approval and relevant provincial health-authority clearance. In Saudi Arabia the Saudi Food and Drug Authority (SFDA) is the primary regulator for AI-enabled software as a medical device.

Phase 6: Deploy — Earning Clinical Adoption. Deployment begins at go-live and continues with active support for at least 90 days. It tests the design decisions of Phases 1–5 against operational reality and is the period in which clinical trust is earned or lost.

A controlled pilot (one ward or shift) runs with the previous process available as backup. The clinical champion—a respected senior clinician who participated in design and is willing to use the system first, visibly—is present. The parallel process is not re-moved until the clinical team itself requests its removal.

Adoption decisions are made primarily through

observation of trusted peers.13 In the Korangi registration project the department head was initially resistant after previous failed IT implementations. Engagement began by asking him to describe his most significant operational problem (the morning registration queue). A paper prototype demonstrated how the digital system would address it. A trusted staff member piloted the system for one week with the pa-per process in parallel; by day four she was defending the system to colleagues. The department head subsequently attended training voluntarily. Conversion occurred through observation of a trusted peer succeeding, not through technical salesmanship.

Training must include explicit teaching on failure modes. A clinician who understands where the tool can be wrong is a calibrated user; one who believes it is always right is a dangerous one. In the DDSS deployment, LHW training included case examples in which the recommendation would be incorrect and discussion of how to recognize and respond. Early override rates were higher than some expected; this was interpreted as evidence that the training was functioning rather than as system failure, although alternative interpretations (distrust, model limitations) remain possible without comparative data.

Weekly case-review sessions examine a small number of recent cases, including at least one override. Questions include: Why did the system recommend this? Was the override correct? What would need to change for this case type to be handled better? Answers feed both model-improvement requests and workflow refinement.

4. DISCUSSION

The components of the framework are not novel; elements appear in implementation science, clinical informatics guidance, and human-factors engineering. What is offered is their integration into a single sequential, clinician-facing sequence specifically for the forward deployment of clinical AI, grounded in practical experience in resource-constrained and clinically complex settings.

Relation to existing frameworks. Broader implementation frameworks such as NASSS (Non-adoption, Abandonment, Scale-up, Spread and Sustainability), CFIR (Consolidated Framework for Implementation Research), SEIPS 2.0, and the DECIDE-AI reporting guideline address technology implementation at a higher level of abstraction. WHO guidance on the ethics and governance of AI for health provides essential normative principles. The six-phase framework does not replace these; it operationalizes the pre-go-live workflow-design and local accountability gap that they leave relatively underspecified. It is intended to sit inside, rather than compete with, those larger frameworks.

Limitations. The framework was developed primarily from deployments in Pakistan and Saudi Arabia. Saudi Arabia is a high-income country; transferability to other high-income health systems (and the mapping of the Govern phase onto the EU AI Act, Medical Device Regulation and GDPR) requires further evaluation. The framework is relatively re-source-intensive and sequential; real deployments often iterate. It has not been validated for fully autonomous systems. Patient and community co-production, formal equity and bias testing, privacy and cybersecurity design, long-term post-90-day monitoring, version control, and decommissioning are acknowledged gaps and priorities for refinement. No comparative outcome data demonstrating superiority of the framework over alternative approaches are presented; the claims rest on practitioner experience and field notes.

Speed and rigor need not be opposed. Design phases (1–5) for a focused deployment can be completed in weeks; thoroughness early prevents the failures that most commonly delay or abort later stages. The controlled pilot of the Sehat Kahani DDSS ran for eight weeks before expansion; the paper-proto-typing session that reduced registration fields from 14 to 6 took two hours.

The single most important principle remains: clinical workflow design for AI cannot be done to clinical teams; it must be done with them. Every phase is a process of co-production. The external expert con-tributes structured methods and implementation experience; the clinical team contributes contextual knowledge, patient understanding, and the institutional credibility without which technology is not adopted. A workflow the team helped design is a workflow they will defend, use carefully, and report problems about honestly. Honest problem reporting is what makes a deployment safe; it is achieved only through genuine co-production.

5. CONCLUSION

Artificial intelligence offers genuine benefits to health systems, and the evidence continues to accumulate. The central remaining challenge is deployment. The six-phase framework—Discover, Map, AI Fit, Redesign, Govern, Deploy—provides a structured, field-tested approach to the clinical workflow design problems that most commonly cause technically sound AI to fail in operational settings. It is offered as a practitioner contribution to the growing body of implementation science for clinical AI—a field that needs voices from those doing the work on the ground as well as those evaluating it from the outside. The algorithm is rarely the hard part. The workflow is.

Funding No specific funding was received for the preparation of this perspective. The deployments referenced were funded by the WHO Alliance for Health Policy and Systems Research, the Bill & Melinda Gates Foundation, and intramural grants from the Aga Khan University and Shaqra University.

Competing interests The author is the founder of HealthRahber (Pakistan) and Tashafe (Saudi Arabia), AI health platforms. None of the specific deployments described in this paper (Sehat Kahani DDSS, Karachi registration digitization, GIS vaccination mapping) were commercial products of these ventures, and no commercial products are recommended herein. The author has no commercial interest in the framework itself.

Author contributions HNT conceived the framework, con-ducted or co-led the deployments referenced, reviewed the literature, and wrote the manuscript.

Ethics The DDSS deployment was approved by the Ethics Review Committee of the Aga Khan University (ERC reference 2023-8514-26533). Other observational and quality-improvement activities described were conducted under institutional permissions consistent with local requirements. No generative AI tools were used in the preparation of this manuscript.

Open access This is an open access article distributed in accordance with the Creative Commons Attribution Non Commercial (CC BY- NC 4.0) license, which permits others to dis-tribute, remix, adapt, build upon this work non- commercially, and license their derivative works on different terms, provided the original work is properly cited, appropriate credit is given, any changes made indicated, and the use is non- commercial. See: https://creativecommons.org/licenses/by-nc/4.0/.

REFERENCES

1. Gulshan V, Peng L, Coram M, et al. Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. JAMA 2016;316:2402–2410.
2. Komorowski M, Celi LA, Badawi O, et al. The Artificial Intelligence Clinician learns optimal treatment strategies for sepsis in intensive care. Nat Med 2018;24:1716–1720.
3. Topol EJ. High-performance medicine: the convergence of human and artificial intelligence. Nat Med 2019;25:44–56.
4. Parkash O, Lal A, Subash T, et al. Use of artificial in-telligence for the detection of Helicobacter pylori in-fection from upper gastrointestinal endoscopy images: an updated systematic review and meta-analysis. Ann Gastroenterol 2024;37:665–676.
5. Sendak MP, D'Arcy J, Kashyap S, et al. A path for translation of machine learning products into healthcare delivery. EMJ Innov 2020. doi:10.33590/emjinnov/19-00172
6. Tahir H, Khan A, Yousaf M, et al. When artificial in-telligence guides and misguides clinicians: A critical appraisal of AI recommendation correctness and diagnostic decision-making. Int J Med Inform 2026;209:106293.
7. Cresswell KM, Calvert M, Khan S, et al. Identifying and categorising the causes of IT-enabled healthcare quality improvement failures: a qualitative study of an NHS electronic prescribing system. BMJ Open 2014;4:e004050.
8. Tahir H, Habib S, Jabeen R, et al. Protocol of digital decision support system (DDSS) embedded within a tele-primary healthcare platform in Pakistan: an assessment of usability, acceptability, effectiveness and perceived quality of care in a resource-constrained set-

ting. BMJ Open. 2024;14:e082167. doi:10.1136/bmjopen-2023-082167
9. Arora VM, Manjarrez E, Dressler DD, et al. Hospitalist handoffs: a systematic review and task force recommendations. J Hosp Med 2009;4:433–440.
10. Skitka LJ, Mosier KL, Burdick M. Does automation bias decision-making? Int J Hum Comput Stud 1999;51:991–1006.
11. Tahir H, Khan A, Yousaf M, et al. Bridging performance and uncertainty: Cautionary notes on machine learning and large language models in TBI prognostication. Int J Med Inform 2026;210:106315.
12. Tahir H, Hassan A, Bibi S, et al. Critical appraisal of a multicenter clinical-radiological machine learning model for peripheral artery disease diagnosis. Int J Med Inform 2026;213:106382.
13. Rogers EM. Diffusion of Innovations. 5th ed. New York: Free Press; 2003.