ASPEN POLICY ACADEMY
aspen institute

Model Framework

TRACE: A Monitoring Framework for AI-Enabled Clinical Decision Support Tools

By Arfa Rehman

January 2026

This Model Framework was developed by an Aspen Policy Academy fellow while participating in the Science and Technology Policy Fellowship. This framework provides guidance on to monitor artificial intelligence tools used in medical clinic decisionmaking. The full project, including a policy brief explaining the fellow's core recommendations, is available here. Please note that the author's opinions published here are their own. This publication does not reflect the views of the Aspen Policy Academy or the Aspen Institute.

Executive Summary

Post-deployment monitoring is a critical yet often overlooked step in the life cycle of AI-enabled clinical decision support (AI-CDS) tool implementation. As the clinical use of AI continues to expand, robust monitoring will be critical to ensure that these technologies deliver on their promise without compromising patient safety. The TRACE monitoring framework offers a practical approach for evaluating real-world performance across 5 core domains:

1. **Technical Integrity:** Is the tool functioning reliably and consistently across time and settings?
2. **Real-World Use:** How are clinicians actually using the tool, and is that use meeting expectations?
3. **Alignment and Accuracy:** Are the tool's outputs clinically sound and consistent with current guidelines?
4. **Clinical Fairness:** Does the tool work equally well across patient subgroups, or are disparities emerging?
5. **Escalation and Safety Response:** When problems arise, is there a process in place to detect, investigate, and respond?

The TRACE framework includes suggested metrics, data requirements, and governance considerations for each domain. It is designed to be flexible, scalable, and to integrate

smoothly into existing quality improvement and oversight processes. It may be incorporated into vendor requirements, procurement guidelines, funding requirements, accreditation standards, or internal governance policies. Adoption of TRACE can help health AI developers and implementers ensure that AI-CDS tools remain safe, effective, and equitable.

The TRACE Framework: Core Monitoring Domains

The TRACE framework outlines 5 domains that represent the essential elements of a post deployment monitoring strategy for AI-CDS tools. Each domain addresses a distinct dimension of real-world tool behavior and implementation risk and includes best practices for implementation and sample metrics to monitor.

T – Technical Integrity

Technical integrity refers to the consistency, transparency, and auditability of the system in which the AI model and tool operate. Monitoring in this domain ensures that model outputs are stable, logs are preserved, and the system infrastructure supports traceability and reproducibility. Strong technical integrity allows organizations to verify the source of outputs, evaluate performance trends over time, and conduct meaningful investigations when reviewing safety events or quality concerns.

Best Practices

- Maintain comprehensive version control for all deployed models, including documentation of model updates, retraining events, and configuration changes.
- Log all model inputs and outputs with associated metadata (e.g., time stamps, clinician identifiers, patient encounter context).
- Conduct stability testing for identical inputs over time, particularly for systems based on large language models (LLMs), which may produce different outputs even when given the same input.
- Track system uptime, latency, and availability within clinical workflows.
- Implement automated alerts for configuration or infrastructure anomalies.

Sample Metrics

- Percentage of identical-input stability tests that yield output variation beyond an acceptable threshold.
- Number of model or configuration changes fully documented in version control.
- Average and maximum latency for model responses during clinical use.

- Percentage of complete and retrievable log records for all tool interactions.
- System uptime percentage during clinical operating hours.

R – Real-World Use

Understanding how clinicians engage with AI-CDS tools in practice provides essential insight into both safety and utility. Real-world use monitoring reveals whether the tool is being used as intended, how it fits within existing workflows, and the extent to which it influences clinical decisions. This information helps identify opportunities to improve usability, build clinician trust, and ensure that the tool is delivering value in everyday clinical contexts.

Best Practices

- Track utilization patterns across departments, clinician types, and time periods to assess reach and consistency of use.
- Monitor override and dismissal rates, along with contextual factors that may explain these actions (e.g., time of day, patient complexity).
- Identify patterns of selective, inconsistent, or excessive use that may signal workflow misalignment, insufficient training, or over-reliance by clinicians who may be overworked or overly confident in the tool's capabilities.
- Where feasible, capture measures such as decision time, engagement duration, or workflow interruptions to assess impact on efficiency.
- Segment usage and override data by role and location to identify where adoption is strongest and where additional support may be required.
- Incorporate findings into regular feedback loops with clinical teams to refine deployment and training strategies.

Sample Metrics

- Percentage of eligible patient encounters in which the tool is accessed.
- Override or dismissal rate, segmented by clinician role, specialty, or site.
- Median time spent engaging with the tool during a clinical encounter.
- Rate of workflow interruptions attributed to tool use (e.g., increased task completion time, task abandonment or detours observed in audit logs, clinician-reported disruptions via surveys or incident reports).
- Change in utilization rate following workflow adjustments, training interventions, or system updates.

A – Alignment and Accuracy

Alignment and accuracy monitoring evaluates whether the AI-CDS tool continues to produce clinically valid, relevant, and guideline-aligned outputs after deployment. Tracking this domain ensures that recommendations remain appropriate as clinical guidelines evolve, patient populations change, or the tool's underlying model is updated. Consistent alignment with evidence-based practice supports clinical trust, safeguards patient outcomes, and ensures that the tool continues to deliver meaningful value in its intended use case.

Best Practices

- Establish processes to routinely review changes to relevant clinical guidelines and reevaluate the tool's outputs for alignment, including documenting when guidelines change and assessing whether retraining, rule adjustments, or workflow updates are needed.
- Evaluating the correctness and clinical appropriateness of the tool's outputs over time.
- Using regular clinician feedback, audit logs, or incident reviews to determine whether the tool is contributing to adverse events, unnecessary testing, or inappropriate care pathways.
- Where feasible, link model-guided decisions to clinical outcomes (e.g., diagnostic accuracy, treatment effectiveness, patient trajectory) to assess impact.

Sample Metrics

- Percentage of outputs concordant with current clinical guidelines or expert consensus.
- Frequency of clinician-reported disagreements with tool recommendations.
- Change in diagnostic accuracy or treatment appropriateness when tool recommendations are followed.
- Frequency of adverse events or inappropriate care associated with tool-guided decisions.
- Time between guideline changes and documented tool reassessment or update.

C – Clinical Fairness

Clinical fairness monitoring evaluates whether the AI-CDS tool performs consistently across different patient populations and does not introduce or reinforce bias. Stratifying performance and usage data by demographic, socioeconomic, and geographic factors

helps identify disparities in accuracy, access, and outcomes. Proactive clinical fairness monitoring supports inclusive care, reduces the risk of exacerbating existing health inequities, and strengthens institutional accountability to patients and communities.

Best Practices

- Stratify key performance metrics (e.g., accuracy, false positive/negative rates) by demographic variables such as race, ethnicity, gender, age, insurance status, language preference, and geography.
- Analyze usage patterns and override rates across subgroups to identify potential access or trust gaps.
- Investigate differences in outcomes, including missed detections or false reassurance, for specific patient subgroups.
- Ensure demographic data is captured in a standardized and complete manner to enable meaningful analysis.
- Incorporate patient and community representatives into governance structures to inform monitoring priorities and interpretation of findings.
- Develop and document action plans for addressing identified disparities, including targeted retraining or workflow changes.

Sample Metrics

- Difference in accuracy rates between demographic subgroups for a defined clinical task.
- Variation in tool utilization rates across patient populations.
- Subgroup-specific override rates compared to overall averages.
- Percentage of patient encounters with complete demographic data available for analysis.
- Number of fairness-related corrective actions implemented over a defined period.

E – Escalation and Safety Response

Escalation and safety response monitoring ensures that findings from all other TRACE domains lead to timely and appropriate action. Clear protocols for investigating deviations, addressing safety concerns, and implementing corrective measures are essential for maintaining trust, minimizing risk, and supporting continuous improvement. This domain links monitoring activities to governance processes, ensuring that observed issues are not only identified but also resolved in a systematic and accountable manner.

Best Practices

- Define thresholds or trigger conditions for escalation (e.g., sudden performance decline, clinical fairness flags, safety event reports).
- Establish a multidisciplinary review body – such as a clinical AI oversight committee or patient safety board – to evaluate escalated cases.
- Integrate monitoring data into existing institutional quality improvement and incident reporting systems.
- Develop predefined workflows for model retraining, suspension, or removal from clinical use when necessary.
- Document all escalation events, investigations, and resolutions to create an institutional knowledge base for future risk mitigation.
- Communicate outcomes of escalation processes to all relevant stakeholders, including clinical teams, governance bodies, and those responsible for maintaining or improving the tool.

Sample Metrics

- Number of escalation events triggered by monitoring thresholds in a given period.
- Average time from detection of an issue to initiation of an investigation.
- Percentage of escalation events resolved within established time frames.
- Frequency of model retraining, suspension, or removal actions following escalations.
- Number of process or policy changes implemented as a direct result of escalation outcomes.

Implementation Considerations

Effective application of the TRACE framework depends on embedding monitoring within formal governance processes. A well-designed governance structure ensures that monitoring activities are not only performed but also reviewed, acted on, and documented in a consistent manner.

Key enablers include:

- Integration with existing governance bodies (e.g., quality and safety committees, clinical AI oversight boards, institutional review panels) to ensure that findings are evaluated in a structured forum.
- Clear assignment of responsibilities for data collection, analysis, decision-making, and implementation of corrective actions, mapped across clinical, technical, and operational roles.

- Defined escalation pathways that specify thresholds for action, investigation procedures, and timelines for resolution.
- Embedded logging and telemetry at the point of deployment to ensure that monitoring of data is complete, timely, and accessible for governance review.
- Data pipelines and analytic dashboards aligned with each TRACE domain to support transparent reporting and decisionmaking.
- Alignment with institutional policies and accreditation requirements so that monitoring outcomes contribute directly to compliance and continuous improvement objectives.

Monitoring strategies should be calibrated to tool risk level, scale of deployment, and institutional capacity. While full implementation of TRACE may not be feasible for all tools, partial adoption can provide valuable safeguards and serve as a stepping stone to more comprehensive oversight.

If you'd like to learn more, see the full project, including a policy brief explaining the fellow's core recommendations, at aspenpolicyacademy.org/project/ai-cds-monitoring-2025.

Stoica Adrian's Images via Canva

Sezeryadigar from Getty Images Signature via Canva

ASPEN POLICY ACADEMY
aspen institute

aspen institute

About the Aspen Policy Academy

The Aspen Institute's Policy Academy helps community leaders and experts across the political spectrum elevate their voices, influence key decisions, and strengthen democracy from the ground up. Our innovative training programs and resources equip people across sectors – from tech to the environment, science to civic engagement – with the skills to shape critical policy efforts. Learn more at aspenpolicyacademy.org.