

ReveIQ

OPEN FRAMEWORK · FREE TO USE

DIPA

Describe. Interpret. Predict. Act.

A clinical reasoning framework for making the clinical story
explicit while it can still change the outcome

ReveIQ Edition · Version 3.0 · September 2026 · Working draft for review and comment

Matthew Ayinde, MBA, BSN, RN

Emergency Nurse · Clinical Documentation Integrity · AI Workflow and Prompt Architecture

Published openly. No license or fee. Attribution requested for derivative or published work. The name and mark
DIPA: Describe. Interpret. Predict. Act. are retained by the author. Comments and corrections:
dipa@reveiqai.com



DIPA Framework

About this version

Version 3.0 adds the mechanism that makes the framework operational.

Predict now has two levers. The primary lever is clinical intuition, the clinician's own judgment about trajectory. The secondary lever is AI predictive support, which surfaces patient-specific trends for the clinician to consider but never authors the prediction.

A closed-loop capability connects Predict to Act. When a clinician documents an escalation trigger, that trigger becomes a monitorable condition. If the condition is met and no response is documented, the system notifies. This converts the paired-statement rule from an exhortation into a mechanism, and it directly addresses failure to rescue: the most common failure is not that risk went unrecognized, but that recognized risk was not acted on in time.

Section 6 is substantially expanded. A system that records unmet triggers creates discovery exposure that must be designed for, not discovered later. Section 6.4 proposes structuring the aggregate monitoring layer as patient safety work product.

Section 8.4 addresses regulatory positioning. The design constraint that keeps this capability out of medical device regulation is the same constraint that keeps it clinically sound: the system tracks execution of a clinician-authored plan. It does not detect deterioration.

Version 2.0 scoped Predict exclusively to clinical trajectory, excluding financial and workflow risk. That scoping is unchanged.

1. Executive summary

Healthcare organizations hold abundant clinical, operational, and financial data, yet often lack insight early enough to change the outcome. Relevant information sits across notes, orders, flowsheets, laboratory results, imaging, utilization records, and authorization files without being assembled at the moment the care team can still act.

DIPA is a **clinical reasoning and communication framework**. It complements, and does not replace, SOAP, APSO, assessment-and-plan, progress-note, and discharge-documentation workflows. It makes four elements of a complete clinical story explicit:

	Question	What it captures
D: Describe	What is happening with the patient now?	Findings, objective data, change from baseline
I: Interpret	What does it mean clinically, and why is it happening?	Diagnosis, etiology, acuity, severity, relationships among conditions
P: Predict	What is the expected clinical course, and what would signal deterioration?	Trajectory, deterioration signals, monitoring needs, escalation triggers, disposition barriers
A: Act	What is the care team doing, monitoring, and escalating?	Treatment, reassessment, escalation plan, coordination, disposition requirements

DIPA is not a coding checklist, an automated diagnosis engine, or a revenue-maximization tool.

The central principle:

Insight is seeing risk early enough to change the outcome.

Recognition alone is insufficient. A framework that makes risk visible but does not ensure it is acted on has solved the easier half of the problem.

2. The problem

Clinicians document extensively. In busy environments, the record often fails to connect five things:

- The patient's findings and change from baseline
- The clinician's interpretation of those findings
- The underlying cause, severity, acuity, and relationships among conditions
- The anticipated course, and what would signal that treatment is failing
- The specific monitoring, reassessment, escalation, and disposition plan


There is a sixth failure, which is not a documentation failure at all: **recognized risk that is not acted on**. A clinician identifies a deterioration signal, states the threshold, and then the patient is handed off three times, the threshold is crossed at 2:40 a.m. by a nurse who was not present for the original reasoning, and the escalation does not occur. This is the failure-to-rescue pattern, and it is a quality problem before it is a documentation problem.

When clinical reasoning is incomplete, dispersed, or recognized-but-unexecuted, the consequences reach well beyond coding:

- Handoffs carry findings without the reasoning behind them, and escalation criteria are implicit
- Case management and utilization review lack an explicit rationale for continued level of care
- CDI teams pursue clarification after the best documentation opportunity has passed
- Coders encounter unresolved ambiguity after discharge


- Payers challenge medical necessity, clinical validity, or severity
- Revenue cycle sees the problem only when it becomes a denial
- Patients experience avoidable delays, confusion, and fragmented transitions

MOST ORGANIZATIONS



The problem becomes visible only after it is expensive.

WITH DIPA



The problem becomes visible while it can still be changed.

The shift DIPA asks for. The same information, surfaced at a point where it can still change the care rather than only explain it.

3. The framework


3.1 Elements



The four elements. Predict and Act are bracketed because the framework does not permit one without the other.

Element	Clinician question	Documentation output	Where technology may assist
Describe	What is happening now?	Symptoms, change from baseline, examination, vital signs, laboratory and imaging findings, treatment response, relevant history, support needs	Surface relevant data trends, baseline comparisons, pending results, and objective evidence distributed across the record
Interpret	What does this mean clinically, and why is it happening?	Diagnosis, differential where applicable, cause-and-effect relationships, acuity, severity, acute/chronic status, clinical significance	Present relevant clinical indicators and relationships for clinician review; never independently establish a diagnosis
Predict	What is the expected course, and what would tell us it is not going well?	Expected trajectory, deterioration and treatment-failure signals, monitoring requirements, escalation triggers, anticipated disposition barriers	Surface candidate risks and patient-specific trends for clinician consideration; never author the prediction (see §3.2)
Act	What are we doing, monitoring, and escalating?	Treatment plan, consults, monitoring parameters, reassessment interval, escalation criteria, coordination, disposition requirements	Monitor clinician-authored triggers and notify when a trigger is met without documented response (see §8.4)

3.2 Predict: two levers



Authorship hierarchy. The constraint that keeps the framework clinically sound, medico-legally defensible, and outside medical device regulation is the same constraint.

Predict operates through two levers, in a fixed order of authority.

Lever one: clinical intuition (primary)

The clinician's own judgment about where this patient is heading. This is formed from pattern recognition, experience, examination, and knowledge of the patient that is not fully represented in structured data. It is the primary lever and it is never subordinate to the second.

Clinical intuition is also frequently the earliest signal available. Experienced clinicians often recognize that a patient is deteriorating before objective criteria are met. A framework that privileges model output over that judgment discards the earliest warning the organization has.

Lever two: AI predictive support (secondary and advisory)

Technology may surface patient-specific trends, trajectory patterns, and candidate risks the clinician may not have weighted, a slow oxygen requirement creep across shifts, a lactate that is falling more slowly than expected, a creatinine trend crossing a threshold overnight.

The second lever operates under a hard constraint:

The AI never authors a Predict statement. It presents candidate risks. The clinician accepts, modifies, or rejects. Once accepted, the prediction is the clinician's own, in the clinician's own words.

This constraint preserves clinical primacy, satisfies the non-leading assistance principle, keeps the clinician as the author of record, and has regulatory consequence (see §8.5).

How the levers combine

	Lever one	Lever two
Source	Clinician judgment	Patient-specific data trends
Authority	Primary	Advisory
Output	The Predict statement	Candidate risks for consideration
Failure mode if over-weighted	Missed signals in complex or high-volume settings	Automation bias; clinician defers to model; alert fatigue

The levers are complementary because they fail differently. Intuition misses the slow trend across shifts. Models miss the patient who looks wrong for reasons not yet in the data.

3.3 What Predict is and is not

Predict includes: 1. Expected clinical course with current treatment 2. Deterioration and treatment-failure signals 3. Monitoring requirements that follow from that risk 4. Escalation triggers, stated as conditions 5. Anticipated disposition or discharge barriers

Predict excludes: - Risk scores, probabilities, or model output pasted into a note as documentation - Denial likelihood, coding rework risk, authorization exposure, or any financial forecast - Any prediction the clinician did not independently make or adopt

3.4 Required design rule: Predict and Act are inseparable

A prediction without a corresponding action is not DIPA documentation.

Anticipated risk must be paired with the monitoring, reassessment, and escalation plan that answers it. They are one move, not two. This rule exists for clinical reasons and for medico-legal ones.

Anticipated risk must be paired with the monitoring, reassessment, and escalation plan that responds to it. They are one move, not two.

Under §8.4, this rule becomes structural rather than advisory: a documented escalation trigger generates a monitored condition, and the Act content defines what closes it.

3.5 Escalation triggers must be stated as conditions

For a trigger to be monitorable, and more importantly for it to be actionable by a covering clinician who was not present for the original reasoning, it must be stated as a condition with a threshold and a response.

Monitorable: "Escalate to ICU for MAP below 65 sustained over 30 minutes, lactate rising above 4.0, oxygen requirement above 6 L/min, or new confusion."

Not monitorable: "Watch closely. Escalate if worsening."

This is a documentation discipline, not a technology requirement. It improves handoff quality whether or not any system is monitoring. But it is a prerequisite for §8.4.

3.6 Origin of the sequence

DIPA's sequence is adapted from a familiar analytics maturity progression:

Analytics maturity	DIPA element	Practical question
Descriptive	Describe	What is happening?
Diagnostic	Interpret	Why is it happening?
Predictive	Predict	What is likely to happen next?
Prescriptive	Act	What should we do now?

One clarification matters. The analogy describes the *sequence of reasoning*, not the actor. In analytics, the model performs each step. In DIPA, the clinician does, with optional model support at Predict. Organizations reading this table as a technology roadmap rather than a cognitive sequence will build the wrong thing.

4. Relationship to existing documentation and to AI tools

4.1 Existing documentation methods

Existing approach	What it captures well	DIPA contribution
SOAP / APSO	Symptoms, objective data, assessment, plan	Makes expected trajectory, deterioration signals, monitoring needs, and escalation triggers explicit
Assessment and plan	Diagnoses and treatment decisions	Connects clinical indicators to interpretation, anticipated course, and action
Daily progress note	Interval events and current management	Provides concise structure for reassessment and next-step clarity
Discharge summary	Hospital course and final plan	Synthesizes what occurred, why, residual risk, and follow-up needs
CDI review	Clarification of incomplete or ambiguous documentation	Enables earlier case prioritization and supports clinically complete documentation the first time

4.2 AI documentation tools

Ambient documentation tools now generate notes at very large scale. They are highly capable at capture and transcription. They are not designed to assess whether the resulting note contains coherent clinical reasoning, and they do not track whether a documented plan was executed.

DIPA is **a standard the output can be measured against**, not a competitor to those tools. An organization deploying ambient documentation can use DIPA to answer a question those tools do not: does this note explain what is happening, why, what may happen next, and what we are doing about it, and did what we said we would do actually happen?

4.3 Early warning systems

INPATIENT

Deterioration without escalation

The patient is in front of you. The criterion was written, the condition occurred, nobody acted.

Mechanism: escalation criteria as observable conditions.

AMBULATORY

The open loop

The patient goes home. A result is unreviewed, a referral uncompleted, a follow-up never booked.

Mechanism: return precautions plus owned loops.

Where DIPA applies. The four elements hold in both settings. What changes is the failure being guarded against, and therefore the mechanism.

DIPA's monitoring capability is not an early warning system and should not be procured, governed, or evaluated as one.

	Early warning systems (MEWS, NEWS, deterioration indices)	DIPA trigger monitoring
Trigger source	Population-derived model	The treating clinician
Question asked	Does this patient look like deteriorating patients?	Did the condition this clinician specified occur without response?
Scope	All patients on the unit	Only patients with a documented trigger
Clinician relationship	Unsolicited second opinion	The clinician's own plan, returned at the relevant moment
Primary failure mode	Alert fatigue from population-level false positives	Trigger vagueness; clinicians writing looser criteria to avoid notification

The two are complementary. An early warning system can inform lever two. It cannot replace lever one, and its output should never be documented as the clinician's prediction.

5. What DIPA is and is not

DIPA is	DIPA is not
A clinical reasoning and communication framework	A coding checklist
An overlay complementing SOAP, APSO, assessment/plan, and progress notes	A replacement for clinician judgment
A structure making findings, interpretation, anticipated course, and action explicit	An automated diagnosis engine
A standard against which documentation, including AI-generated documentation, can be evaluated	A deterioration detection or early warning system
A mechanism for ensuring documented plans are executed	A query factory or documentation generator
A support to care coordination, quality, utilization, CDI, coding, and revenue integrity	A mechanism to encourage unsupported diagnoses or reimbursement-driven documentation
A prospective model for recognizing and acting on risk during the episode of care	A substitute for compliant clinical review, coding, CDI, or physician-advisor oversight

6. Medico-legal considerations

This section must be addressed before implementation. DIPA asks clinicians to document anticipated clinical course, and §8.4 creates a record of whether documented triggers were acted on. Both have consequences that must be designed for.

6.1 The documentation concern

Documented anticipation of deterioration establishes foreseeability. If a clinician documents that a patient is at risk for a specific deterioration, that deterioration occurs, and the record shows no corresponding response, the documentation supports a failure-to-rescue theory.

Clinicians will raise this. An implementation with no answer will fail at the first education session.

6.2 Why the paired-statement rule is the answer

The risk arises from **unpaired prediction**: anticipated risk with no documented response. It does not arise from prediction as such.

Documented anticipated risk, paired with monitoring parameters, reassessment intervals, and escalation criteria, is contemporaneous evidence that the clinician recognized the risk and planned for it. That is the record that defends clinical care. Its absence is what leaves a clinician explaining, months later and from memory, what they were thinking.

6.3 The monitoring log creates a distinct and larger exposure

This is the more serious issue, and it is new in version 3.0.

A system that monitors clinician-authored triggers necessarily generates a record of every instance in which a trigger was met. Where no response was documented, that record constitutes timestamped evidence of documented foreseeability, automated notification, and documented non-response.

That is a materially stronger artifact than the note alone. Organizations must decide deliberately what is created, where it lives, how long it is retained, and what is privileged, before deployment, not after the first request for production.

6.4 Structuring the aggregate layer as patient safety work product

LAYER 1 · CLINICAL RECORD

Belongs in the chart

- The trigger and the data that met it
- The clinician or nurse response
- Documented reassessment and rationale

Discoverable, and usually protective.

LAYER 2 · AGGREGATE ANALYSIS

Quality improvement work

- Close rates and response time patterns
- Unit, shift, and routing performance
- Trigger quality and override analysis

Structure as patient safety work product.

Separate them architecturally, not by policy. Storing quality analysis inside the chart makes improvement work discoverable. Relying on aggregate analysis as the record leaves the chart incomplete.

Two artifacts, governed differently. Confusing the layers is the failure mode.

The proposed structure separates two distinct artifacts.

The clinical record. What belongs in the chart is the clinical fact: the trigger was met, and the clinician responded, either by intervening or by documented reassessment concluding the plan was unchanged and why. This is patient care information. It is part of the medical record, it is discoverable, and it should be. It is also, in the ordinary case, protective: it documents that the clinician reassessed.

The aggregate monitoring layer. Trigger closure rates, response time distributions, unit- and service-level patterns, override analysis, and system performance data are quality improvement artifacts. They exist to make the organization safer, not to document an individual patient's care.

The Patient Safety and Quality Improvement Act of 2005 established Patient Safety Organizations and created federal privilege and confidentiality protections for **patient safety work product**: information assembled or developed by a provider within a patient safety evaluation system for reporting to a listed PSO. The framework exists precisely so that organizations can examine near-misses and system failures without generating evidence against themselves, which is the condition under which honest examination is possible at all.

Proposed structure:

1. Define a patient safety evaluation system that encompasses the aggregate monitoring layer, documented before deployment
2. Report that layer to a listed PSO under the organization's existing PSO relationship where one exists
3. Maintain clear separation between the medical record content (§6.4, "clinical record") and the PSES content, architecturally, not merely by policy
4. Establish retention and disposition policy for both layers at design time
5. Document the quality improvement purpose of the aggregate layer contemporaneously

Necessary cautions. This structure must be designed and confirmed by counsel; it is not self-executing.

- PSWP protection does not extend to the original medical record, and it cannot be used to shield information a provider is otherwise required to collect or maintain. Information does not become privileged by being copied into a PSES.
- The scope of the privilege has been litigated with varying results, and some courts have construed it narrowly, particularly where information is required to be maintained under state law.
- Some courts have required evidence of actual reporting to a PSO, not merely an intent to report.
- State peer review and quality improvement privileges may apply in parallel and vary substantially.

The realistic goal is not absolute protection. It is a defensible, deliberate, documented structure that distinguishes patient care information from quality improvement analysis, which is materially better than an undifferentiated log created without having thought about it.

6.5 Grounding in established practice

DIPA does not introduce anticipatory clinical reasoning. Failure to rescue, deterioration that was recognizable but not acted on in time, is a long-standing quality construct, and early warning systems exist because anticipating trajectory changes outcomes. Clinicians already perform this reasoning continuously.

DIPA makes it visible in the record rather than leaving it in the clinician's head, where it is unavailable to the next team, the covering physician, the night nurse, or anyone reviewing the case later. The monitoring capability addresses the specific failure in which the reasoning was documented and simply not executed.

6.6 Required before implementation

- Review with risk management, malpractice counsel, and health information management
- Written guidance on documenting uncertainty and differential trajectory

- Explicit instruction that predictions are clinical judgments, revisable as the patient changes
 - Confirmation that paired-statement discipline is reflected in templates and education
 - **Defined architecture separating medical record content from the aggregate monitoring layer**
 - **PSES definition, PSO reporting arrangement, and retention policy established before go-live**
 - A defined process for what happens when a documented escalation trigger is met
 - Regulatory counsel review of the monitoring capability under §8.5
-

7. Worked clinical example

Pneumonia with sepsis, acute hypoxic respiratory failure, and acute kidney injury

Describe. Patient admitted with three days of productive cough, fever, worsening dyspnea, and progressive weakness. Temperature 39.1°C, HR 118, RR 28, BP 88/54, oxygen saturation 86% on room air. Chest radiograph demonstrates right lower-lobe infiltrate. WBC 18.4, lactate 3.8 mmol/L, creatinine 2.0 from baseline 0.9.

Interpret. Sepsis due to community-acquired right lower-lobe pneumonia, with hypotension and lactic acidosis. Acute hypoxic respiratory failure due to pneumonia. Acute kidney injury likely secondary to sepsis-associated hypotension and volume depletion.

Predict. At risk for progression to septic shock and worsening hypoxemia if hemodynamics, oxygen requirement, lactate, and mental status do not improve with resuscitation and antibiotics. Renal function may worsen before improving. Requires monitoring of hemodynamics, oxygenation, urine output, renal function, and lactate clearance.

Act. IV fluid resuscitation, blood and sputum cultures, IV ceftriaxone plus azithromycin, supplemental oxygen to maintain saturation above 92%, serial lactate and BMP at 6 hours, strict intake and output, reassessment every 4 hours. **Escalate to ICU for MAP below 65 sustained over 30 minutes, lactate rising above 4.0, oxygen requirement above 6 L/min, or new confusion.** Discharge only after hemodynamic stability, improving oxygenation, tolerance of oral medication, and a safe follow-up plan.

Note three things. Every anticipated risk in Predict has a corresponding monitoring parameter or escalation trigger in Act. The escalation criteria are stated as conditions with thresholds, not intentions. And those conditions are specific enough that a covering clinician at 3 a.m. , or a monitoring system, can determine unambiguously whether they have been met.

8. The role of technology

Technology may support DIPA only as a transparent clinical-insight assistant, never as an autonomous clinician, coder, or revenue optimizer.

8.1 Appropriate functions

- Identify relevant objective evidence distributed across the record
- Display trends in vital signs, oxygen requirement, laboratory values, imaging, treatment response, and baseline status

- Surface candidate risks for clinician consideration at Predict (lever two, §3.2)
- Flag potential gaps in the documented clinical story for clinician review
- Monitor clinician-authored escalation triggers and notify when met without documented response (§8.4)
- Identify patients whose acuity or treatment intensity warrants concurrent CDI, physician-advisor, case-management, or coding review
- Route low-confidence, high-risk, conflicting, or complex cases to qualified humans


8.2 Prohibited without governance

- Establishing diagnoses independently
- **Authoring Predict content on the clinician's behalf.** Technology may surface a trend or a candidate risk; the documented prediction is the clinician's judgment
- **Generating escalation triggers the clinician did not set**
- Recommending diagnoses based on reimbursement potential
- Generating leading queries
- Making high-stakes clinical, denial, appeal, or billing decisions without defined human review
- Obscuring the patient-specific source evidence behind a recommendation
- Replacing professional coding, CDI, physician-advisor, or clinical judgment in complex cases

8.3 Required design principles

1. **Evidence traceability:** every prompt or recommendation identifies the source evidence used
2. **Human control:** users can accept, reject, edit, or defer any suggestion
3. **Confidence-based routing:** low-confidence, conflicting, or high-risk items route to the appropriate human expert
4. **Non-leading assistance:** prompts invite clinical assessment where evidence supports it; they do not direct a diagnosis
5. **Workflow integration:** support appears in the existing EHR workflow, not a separate portal
6. **Auditability:** the organization can review recommendation, evidence, user action, override, and outcome
7. **Clinical primacy:** care quality, patient safety, documentation clarity, and compliance take precedence over financial performance
8. **Clinician authorship:** predictions and triggers originate with the clinician; technology surfaces and monitors, it does not author

8.4 Trigger monitoring and closed-loop escalation



The system tracks execution of a plan a clinician wrote. It does not assess the patient.
Documented reassessment closes a trigger exactly as intervention does.

Trigger lifecycle. Every monitored condition traces back to a sentence a clinician wrote, and retires at the reassessment cycle unless reaffirmed.

This capability exists to address one failure: a documented escalation trigger is met, and no action follows.

What it does

1. A clinician documents an escalation trigger as a condition with a threshold (§3.5)
2. The trigger becomes a monitored condition, scoped to that patient and that admission
3. The system evaluates discrete clinical data against the condition
4. When the condition is met, the system checks whether a corresponding response has been documented
5. If not, it notifies: beginning with the care team member closest to the bedside
6. The trigger closes when the clinician acts or documents reassessment

What it explicitly does not do

- It does not detect deterioration. It checks conditions the clinician specified.
- It does not generate triggers. Every monitored condition was authored by a clinician.
- It does not interpret waveforms or continuous physiologic streams. It evaluates discrete, documented data: laboratory values, recorded vital signs, medication administration, order status.
- It does not direct treatment. It reports that a stated condition was met and no documented response exists.

The distinction is not cosmetic. It is the difference between plan-adherence tracking and deterioration detection, and it determines both clinical acceptability and regulatory status (§8.5).

Required design elements

Documented reassessment closes the trigger. The clinician must be able to close with "reassessed, plan unchanged" and a rationale. Without this path, the system forces either intervention or avoidance, and clinicians will choose avoidance. This path is also clinically and legally valuable: it is a documented reassessment where none might otherwise exist.

Tiered routing. First notification goes to the care team member closest to the bedside, ordinarily the nurse. Escalation to the provider, then to charge nurse or rapid response, occurs on non-response, not on trigger. Notifying the attending physician at 3 a.m. for every threshold crossing will end the pilot.

Triggers expire. A trigger set on hospital day one must not still fire on day six for a recovering patient. Bind expiry to the reassessment cycle so triggers are reaffirmed or retired at each DIPA update. This reinforces the documentation discipline rather than adding to it.

Trigger revision is expected. As the patient's trajectory changes, triggers change. The system must make revision easy and must log the revision.

Volume control. Configurable limits on concurrent active triggers per patient, with governance review of triggering patterns. A capability that produces more alerts than an early warning system has defeated its own premise.

The predictable failure mode

Clinicians will learn that specific triggers generate notifications and vague triggers do not. The rational response is to write vaguer triggers.

This must be measured (§9.1) and addressed through education and governance rather than through enforcement. If trigger specificity degrades over the pilot, the capability is producing worse documentation than no capability at all, and that finding should stop deployment.

8.5 Regulatory positioning

The monitoring capability sits close to the boundary of medical device regulation, and how it is designed and described determines which side it falls on.

FDA's clinical decision support guidance excludes software from the device definition only where all four statutory criteria are met, including that the software does not acquire, process, or analyze images, IVD signals, or patterns from signal acquisition systems, and that it enables the clinician to independently review the basis for its recommendations. Time-critical alerting, sepsis and stroke alerts are the standard examples, has generally been treated as device functionality, on the reasoning that time sensitivity precludes independent review. Software interpreting continuous physiologic monitoring streams is likewise device functionality.

The design constraints in §8.4 are what keep this capability on the non-device side:

Criterion	Why the design supports non-device status
No image, IVD signal, or signal-pattern analysis	Evaluates discrete documented data only, lab results, recorded vitals, medication administration, order status
Displays or analyzes medical information	The monitored condition and the underlying data are both ordinary clinical information
Supports rather than directs	The recommendation originates with the clinician; the system returns the clinician's own stated plan
Enables independent review of the basis	The basis is a sentence the clinician wrote and can read at the moment of notification

The fourth criterion is the strongest position available to this design and should be protected accordingly. It is also where time-criticality is now evaluated, which makes the plan-adherence framing load-bearing rather than rhetorical.

Language has regulatory consequence. Describing this capability as AI that predicts patient deterioration invites device classification. Describing it accurately, a system that tracks whether clinician-authored plans are executed, does not. Organizations and vendors should obtain regulatory counsel review before publishing product descriptions of this capability.

9. Measurement

A framework that cannot be evaluated cannot be adopted. The measure set below is deliberately small. Organizations consistently over-specify pilot measurement and then lack statistical power on any of it.

9.1 Primary measures: declare thresholds before baseline

#	Measure	Definition	Direction
1	Documentation time	Minutes per note, pilot cohort vs. baseline	No material increase
2	Documentation completeness	Structured audit score on presence and clarity of etiology, acuity, anticipated course, monitoring, escalation criteria	Improvement vs. baseline
3	Paired-statement rate	Proportion of Predict statements with corresponding Act content	High and stable
4	Trigger specificity	Proportion of documented escalation triggers stated as conditions with thresholds	High and not declining over time
5	Trigger closure	Proportion of met triggers closed by documented action or documented reassessment, and time to closure	Improvement; short intervals
6	Retrospective query volume	Post-discharge clarification queries per 100 discharges	Reduction
7	Unsupported diagnosis rate	Audited rate of diagnoses without adequate clinical support	No increase
8	Clinician-reported usefulness	Structured survey, including perceived alert burden	Majority report clinical value

Measures 1, 4, and 7 are **balancing measures**. The framework fails if it improves completeness by increasing burden, if it drives clinicians toward vaguer triggers, or if it produces diagnoses the record does not support. Declare thresholds for these before baseline collection, and treat a decline in measure 4 as a stop signal rather than a tuning problem.

9.2 Secondary measures: descriptive only

Reported without inferential claims: query response time, concurrent review coverage, final DRG agreement, coding rework, notification volume per patient-day, escalation tier distribution, rapid response and unplanned ICU transfer rates, discharge barriers identified, length of stay, medical-necessity and clinical-validation denials, avoidable write-offs.

Financial outcomes belong here, not in the framework. Claim them as observed effects with appropriate caveats for claim lag, case-mix variation, and concurrent initiatives.

Clinical outcome measures are hypothesis-generating only. A single-unit 90-day pilot cannot demonstrate an effect on failure-to-rescue rates or mortality. Report them; do not claim them.

9.3 Required analytic practices

- Define denominators before baseline collection

- Stratify by clinician group, service line, and relevant patient population
- Report clinical and financial measures separately and in that order
- Report unintended consequences alongside benefits
- Avoid attributing causality to DIPA where concurrent initiatives affect the cohort
- Include qualitative clinician feedback alongside quantitative results
- Maintain the §6.4 separation between medical record content and aggregate analysis throughout evaluation

9.4 Equity

Review documentation completeness, trigger setting, notification distribution, response time, and override patterns across relevant patient populations. Where AI supports lever two, review its performance across the same dimensions. Investigate and remediate unexplained variation. Differential response time to met triggers is a patient safety finding, not a workflow metric.

10. Governance

DIPA should be owned as a clinical framework, not as a revenue cycle, technology, or AI initiative. Ownership determines perception, and perception determines whether clinicians adopt it.

Executive sponsor: Chief Medical Officer, Chief Quality Officer, or Chief Clinical Informatics Officer. A revenue cycle executive should participate but should not be the sole sponsor.

Participants: physician champions; nursing and APP leadership; CDI leadership; HIM and coding leadership; quality and patient safety; physician advisors and utilization management; compliance, legal, and risk management; privacy, security, and IT; revenue integrity; and vendor representatives where applicable and subject to organizational governance.

Where §8.4 is implemented, patient safety leadership and the organization's PSO relationship owner are required participants, not optional ones.

Core policy statement:

DIPA does not direct patient care, establish diagnoses, or replace a licensed clinician's independent professional judgment. Predictions and escalation triggers are authored by clinicians. Technology supporting DIPA must be evidence-based, transparent, auditable, and governed through human oversight. Financial impact may be measured, but it cannot be the sole or primary success criterion.

11. Licensing and use

DIPA is published as an open framework. Adoption is the objective.

- The framework, its elements, and the documentation structure may be used, adapted, and implemented by healthcare organizations without license or fee
- Attribution is requested for derivative or published work
- The name and mark **DIPA: Describe. Interpret. Predict. Act.** are retained by the author
- Organizations publishing implementation results are encouraged to share methodology and findings, including negative findings

Commercial products claiming DIPA alignment should be evaluated against the required design principles in §8.3, the prohibited uses in §8.2, and the design elements in §8.4. Alignment is a property of behavior, not of marketing.

12. Invitation

DIPA is an emerging framework intended for collaborative refinement and responsible testing. Contributions are invited on:

- Framework refinement and specialty-specific application
- Documentation and workflow design
- Trigger monitoring design, routing models, and alert burden management
- Medico-legal structure, including patient safety work product architecture
- EHR integration and user experience
- Clinical safety, regulatory, and compliance validation
- Pilot design, evaluation, and publication, including studies that find no effect

The organizing question:

How can we make clinical reasoning visible, and ensure it is acted on, early enough to improve care and protect patients before the record is all that remains?

13. Closing

A complete note does not merely state what the patient has. It explains what is happening, why it is happening, what may happen next, and what the care team is doing about it.

Predict without Act is an unanswered warning. Act without Predict is a plan without a reason. A plan documented and not executed is the failure that quality improvement has spent decades trying to name.

DIPA makes the clinical story explicit, and closes the loop between what the clinician anticipated and what the team actually did, while it can still change the outcome.

Version 3.0 | September 2026 | Working draft for review and comment

Use it. Adapt it. Tell us where it breaks.

DIPA is published openly because a framework is only worth something if people use it. There is no license, no fee, and no product you have to buy to implement it.

Four documents are available: this framework, an ambulatory companion, a clinical pilot charter, and a vendor capability and RFI guide. Contributions are invited on specialty-specific application, trigger design, clinical safety validation, and published evaluation, including studies that find no effect.

reveiqai.com/dipa.html



Matthew Ayinde, MBA, BSN, RN

Emergency Nurse · Clinical Documentation Integrity · AI Workflow and Prompt Architecture

Comments and corrections: dipa@reveiqai.com

Working draft for professional education. Not legal, regulatory, or clinical advice.