

The Clinical AI Autonomy Ladder

AI as Clinical Labor · Part 2 · Version 1.1

Jivesh Sharma, M.D. · HealthIT.com
Editorial date September 30, 2026

A stronger clinical team for a healthier tomorrow.

Same purpose. Greater possibilities.

Scope

An original delegation framework, not a validated maturity, clinical-risk or regulatory scale. Examples are illustrative; classify the deployed task, not the product name. Higher autonomy is not the goal.

This text edition describes the full reference, including examples and human duties. The six rungs map delegated authority, rather than clinical risk or demonstrated benefit.

1. Observe

Flags a signal for human assessment.

AI. Detects signals and surfaces risk.

Human. Interprets the signal and decides the response.

Illustrative examples. Sepsis risk scores, imaging triage flags, wearable alerts.

Fits this rung when the system flags a signal and a person decides the clinical response. High clinical risk can occur at any rung.

2. Suggest

Offers options for human judgment.

AI. Presents options, evidence and uncertainty.

Human. Selects, modifies or rejects.

Illustrative examples. Treatment options, differential diagnosis, candidate trial matching.

Recommendations support human assessment. Candidate trial matching does not establish eligibility or enrollment; generated rationale is not proof of faithful reasoning.

3. Draft

Prepares a work product for review.

AI. Drafts notes, messages or other text.

Human. Reviews and finalizes before use or send.

Illustrative examples. Ambient note drafts, draft inbox replies, prior-auth drafts.

Draft text remains under human review. A system that prepares an executable submission awaiting approval belongs on rung 4 for that step.

4. Act with sign-off

Stages an action for human approval.

AI. Prepares an action; releases it after approval.

Human. Approves or rejects before release.

Illustrative examples. Pended refill orders, referrals awaiting approval, pended order sets.

The executable action requires individual authorization before release. Automatic protocol-based refill or referral actions belong on rung 5 under this framework.

5. Act and report

Executes specified, preauthorized actions.

AI. Follows an approved protocol; logs actions.

Human. Responds to exceptions and reviews outcomes.

Illustrative examples. Protocol scheduling, approved result notices, care-gap reminders.

Use explicit rules and timely escalation. Result notices mean approved nonurgent communications, not unsupervised interpretation, treatment changes or critical-result management.

6. Bounded autonomy

Makes a defined, validated decision.

AI. Decides within a specified task boundary.

Human. Sets limits, monitors performance, handles escalation.

Illustrative examples. Autonomous DR interpretation within labeled use (LumineticsCore).

The autonomous decision here is interpretation of retinal images within the device's intended use. Staff acquire images; clinicians handle results, referrals and treatment. The AI does not perform the whole eye-care pathway.

The diabetic-retinopathy example

LumineticsCore, formerly IDx-DR, autonomously detects more than mild diabetic retinopathy in adults aged 22 or older with diagnosed diabetes and no previous diabetic-retinopathy diagnosis, using the supported Topcon NW400 camera. This is bounded diagnostic interpretation, not a complete eye examination or autonomous treatment. Trained staff capture images. The responsible clinical team reviews outputs, manages inadequate-image results and arranges referral or follow-up. Refer to current labeling for the complete indication, exclusions and warnings.

Use the two frameworks together

The Healthcare AI Labor Stack maps overlapping work functions. This ladder maps delegated authority. Map each workflow step to its function or functions and its rung. One product can occupy several cells. Neither axis proves clinical benefit or permitted use.

Choose appropriate controls

Define the task. Name the function or functions, inputs and decision boundary.

Name the human duties. Record who approves, responds to exceptions and can stop the task.

Match the controls. Use clinical risk, autonomy and reversibility to choose safeguards.

Measure the whole burden. Count time saved after review, correction and exception handling.

Clinical and organizational accountability remain with people.

Sources

[FDA IDx-DR De Novo decision summary, DEN180001](#)

[Current LumineticsCore indication, workflow and warnings](#)

[Epic Art: reviewed draft notes and messages](#)

[NIH TrialGPT FAQ: potential trial identification](#)

[CMS CLIA guidance, section 493.1291\(g\): critical-result alerting](#)