

The Clinical AI Autonomy Ladder

How much work can AI do on its own?

A stronger clinical team for a healthier tomorrow.

1 Observe

2 Suggest

3 Draft

4 Act with sign-off

5 Act and report

**6 Bounded
autonomy**

Choose the appropriate **delegated authority** for each task. Advancement is not the goal.

Original conceptual framework, not a validated maturity, clinical-risk or regulatory scale. Classify the deployed task, not the product name.

1

Observe

Flags a signal for human assessment

AI DOES

Detects signals and surfaces risk

PEOPLE DO

Interpret the signal and decide the response

ILLUSTRATIVE EXAMPLES

Sepsis risk scores, imaging triage flags, wearable alerts

TASK BOUNDARY

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 DOES

Presents options, evidence and uncertainty

PEOPLE DO

Select, modify or reject

ILLUSTRATIVE EXAMPLES

Treatment options, differential diagnosis, candidate trial matching

TASK BOUNDARY

Human assessment remains required. Candidate trial matching does not establish eligibility or enrollment. Generated rationale does not prove faithful reasoning.

3 **Draft**

Prepares a work product for review

AI DOES

Drafts notes, messages or other text

PEOPLE DO

Review and finalize before use or send

ILLUSTRATIVE EXAMPLES

Ambient note drafts, draft inbox replies, prior-auth drafts

TASK BOUNDARY

Drafts require review before use or send. Preparing an executable submission awaiting approval is rung 4 for that step.

4

Act with sign-off

Stages an action for human approval

AI DOES

Prepares an action;
releases it after
approval

PEOPLE DO

Approve or reject
before release

ILLUSTRATIVE EXAMPLES

Pended refill orders, referrals awaiting approval,
pended order sets

TASK BOUNDARY

Each executable action requires approval before release. Automatic protocol-based refill or referral actions are rung 5 in this framework.

5

Act and report

Executes specified, preauthorized actions

AI DOES

Follows an approved protocol; logs actions

PEOPLE DO

Respond to exceptions and review outcomes

ILLUSTRATIVE EXAMPLES

Protocol scheduling, approved result notices, care-gap reminders

TASK BOUNDARY

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



Bounded autonomy

Makes a defined, validated decision

AI DOES

Decides within a specified task boundary

PEOPLE DO

Set limits, monitor performance, handle escalation

ILLUSTRATIVE EXAMPLES

Autonomous DR interpretation within labeled use (LumineticsCore)

TASK BOUNDARY

Autonomy covers retinal-image interpretation within labeled use. Staff acquire images; clinicians handle results, referrals and treatment. It does not cover the whole eye-care pathway.