AI Skills: A Phased Roadmap for Junior and Senior Attorneys

Lucas Hjelle & Nathan Elder

September 2026

Presenters



Lucas Hjelle

Principal

612-359-3269 | lhjelle@slwip.com

Areas of Practice:

- Artificial Intelligence
- Autonomous Vehicles
- Communications
- E-Commerce
- Electronics
- Medical Devices
- Nanotechnology

IP Services:

- Freedom To Operate Opinions
- Patent Analytics
- Patent Procurement
- Portfolio Management
- Validity Opinions



Nathan Elder

Principal

408-660-2014 | nelder@slwip.com

Areas of Practice:

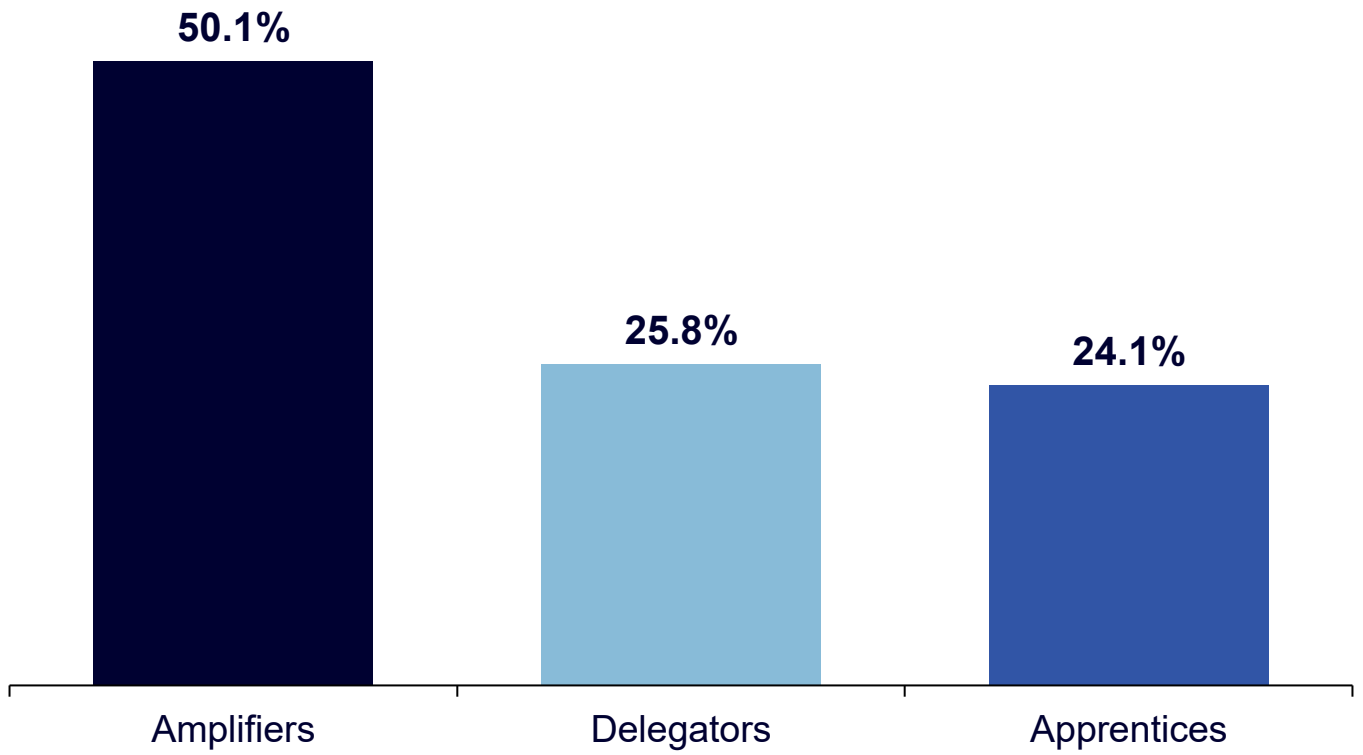
- Artificial Intelligence
- E-Commerce
- Electronics
- Semiconductors
- Software

IP Services:

- Freedom To Operate Opinions
- Patent Procurement
- Portfolio Management
- Validity Opinions

The finding

523 early-career professionals did realistic, expert-graded client work with an AI agent. Each was measured against what the AI produced alone.



50.1% Amplifiers
Outperformed the AI

25.8% Delegators
Matched the AI

24.1% Apprentices
Worse than the AI alone

Source: Agarwal, Barua, Puvvada, Song & Wen, HBR, July 2026

Three profiles, at every level of seniority

The three groups shared common levels of critical thinking abilities, AI literacy, and domain knowledge.
The method of interacting with AI separated people into these groups, and AI interaction methods are coachable. All three exist at SLW today.

AMPLIFIER	Frames, grounds, interrogates, decides, and records. Can defend every call.	**Beats the tool**
DELEGATOR	Hands over inputs, takes back output, accepts what looks right. Adds little reasoning.	**Matches the tool**
APPRENTICE	Engages and critiques, but aims at the wrong layer. Fixes prose while the defect survives.	**Worse than the tool**

The delegation trap is a principal risk to a law firm, as it is a failure mode that may otherwise appear successful.

Terms and findings: Agarwal et al., HBR, July 2026

The same three profiles, at every career stage

Method, not tenure. Four real snapshots of amplifier vs. delegator vs. apprentice.

Clerk

May generate AI-drafted claims from the disclosure, never having tried drafting unaided. The fix is to draft it yourself first, then use AI to pressure-test the boundaries.

Early career

May accept the AI's office action argument because it sounds right, and files it. The fix is to check the AI's characterization against the reference itself before making a decision.

Mid career

May ask AI once for a response and send back whatever comes out. The fix is to interact like a conversation, including proposing ideas, challenging responses, redirecting outcomes, and then making a decision.

Senior

May review AI-assisted work like a junior's draft, but may miss the fluent-but-wrong parts a junior would never produce. The fix is to ask what the tool got wrong before it reaches your signature.

The behavior is coachable at every stage.
The trap does not disappear with seniority; it just gets better dressed.

Two populations, two different problems

Juniors

Lack the substrate

Amplification requires something to amplify. A clerk cannot evaluate a claim set that they could not have written. Give them AI too early, and you get a delegator with no path out of delegation. You may also get into "learning debt," where a junior attorney is behind where they should be in their abilities compared to similarly situated coworkers.

Build the foundation first.

Experienced practitioners

Lack the method

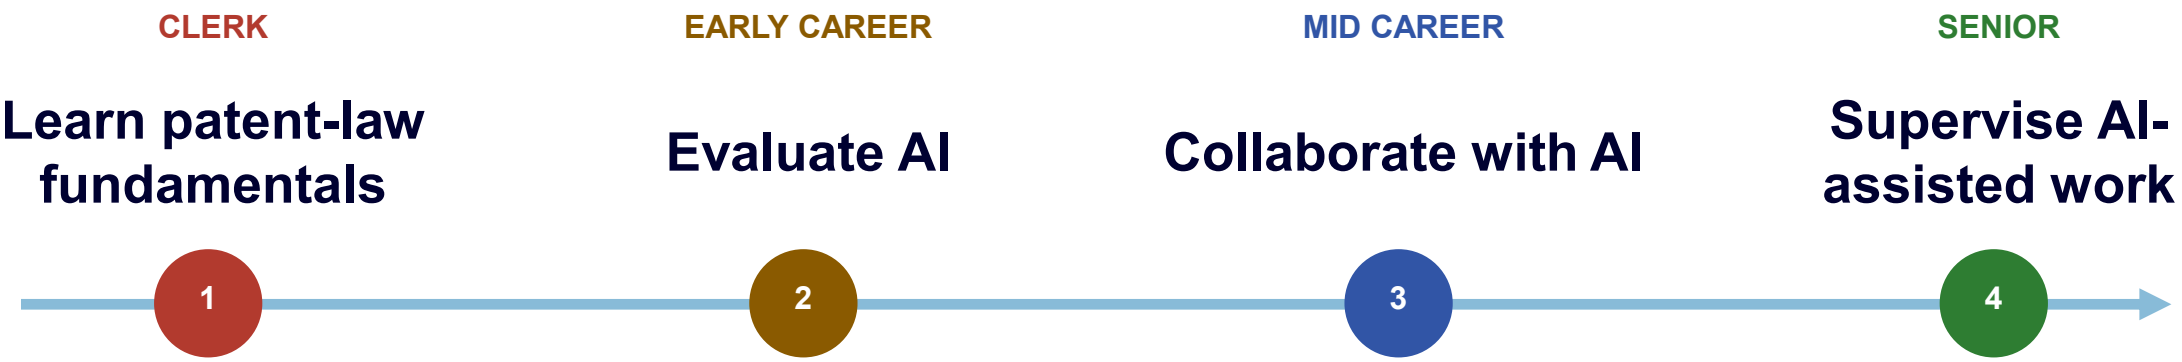
Experienced practitioners may have well-developed industry judgment, but that judgment may not easily translate into the new medium of AI interaction. AI fails differently than an associate in that it may fabricate citations, misinterpret rules or caselaw, or may drop claim language under the guise of smoother prose.

Recalibrate, don't remediate.

The role for AI should change along with changes to the attorney's ability to evaluate, direct, and supervise.

One career arc, four training objectives

The role for AI changes as capability changes; professional responsibility does not.



A single objective across the career arc:
Produce work better than the tool's baseline, and explain why.

The junior pathway: four phases

Progression by demonstrated capability, not by calendar. The committee sets duration and the task list.

0 Foundation

AI OFF

First claim sets, OA strategy, first spec sections, reference reading; all unaided by AI.

1 Guided

Supervised

The loop and the decision log required from the very first session, before the tool.

2 Practicing

Routine use

Beat-the-baseline assignments; method reviewed alongside substance.

3 Coaching

Full use

Reviews AI methods of other attorneys, runs clinics. A firm role, not a graduation.

One method, taught to everyone

Identical for a summer associate and a shareholder. Identical across Arty, Claude, Cowork, Copilot — and whatever comes next.

1

FRAME

Define the problem, the constraints, and what a good answer looks like, all before you prompt.

2

GROUND

Give AI the appropriate materials for the task: disclosure, art, file history, client standards, governing framework (as applicable).

3

INTERROGATE

Treat output as a hypothesis to be questioned. Challenge assumptions, demand reasoning, verify authority.

4

DECIDE

Accept, change, or reject against criteria you can state.

5

RECORD

Make a note of specific judgment calls: you overrode the AI, caught a gap, or redirected scope or strategy.
Also note your reasoning.

The experienced-practitioner track

Recalibrate expertise to apply to AI, keeping in mind the different failure modes of AI.

1

Practice-group recalibration

Discuss three profiles (amplifier, delegator, apprentice), delegation trap, and examples of failure modes.

2

Task playbooks

Discuss where AI helps, where the traps sit, and what must remain human.

3

Decision logs

Discuss examples of recording two or three lines on the judgment calls that mattered.

4

Peer method review

Fifteen-minute walkthroughs (either examples or a post-mortem review) of how a practitioner directed an AI tool.

Do not teach prosecution to experienced practitioners. Teach how AI fails.

One attorney, three modes

The study sorted people. In my week, the task decides which one I am, and two of the three are on purpose.

DELEGATOR

25.8% Matches the tool

IN THE STUDY

Accepts the output with little scrutiny. The tool's floor carries the work.

WHERE IT SHOWS UP IN MY WORK

Filing attachments, indexing folders, first-pass inbox triage, formatting. Reliable, low stakes, easy to check.

On purpose

AMPLIFIER

50.1% Beats the tool

IN THE STUDY

Structures the task, directs the AI, sets the criteria, and refines over several rounds.

WHERE IT SHOWS UP IN MY WORK

OA strategy, claim scope, complex inventor comments, foreign-filing calls. The judgment is the product.

On purpose

APPRENTICE

24.1% Worse than the tool

IN THE STUDY

Engages and critiques, but at the wrong things. The prose improves; the defect survives.

WHERE IT SHOWS UP IN MY WORK

Whenever I review what is easy to see instead of what matters. It has happened to me.

By accident

Percentages: share of 523 early-career professionals, KPMG × UT Austin field study (Agarwal et al., HBR, July 2026).

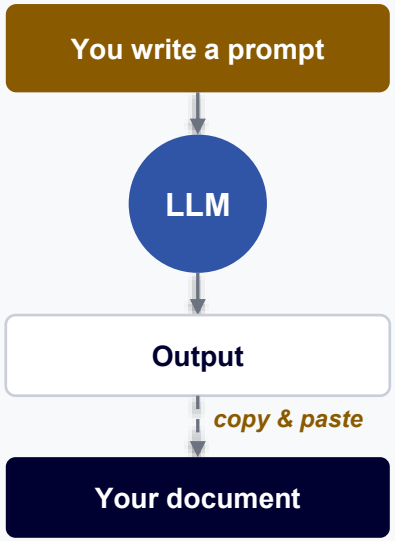
Delegation isn't the trap. Delegating the wrong task is.

Three ways to put a model to work

Same model, different harness. What changes is who picks the next step, and what the model can reach.

1 BASIC CHAT

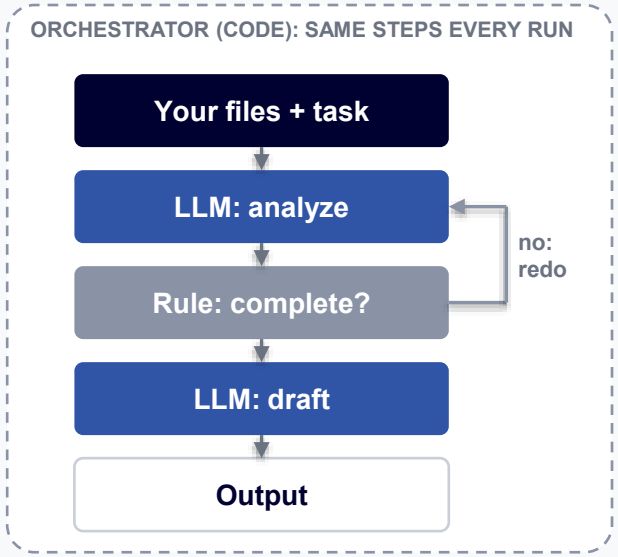
one prompt, one answer



You are the harness. It sees only what you paste in, and you carry the answer out.

2 FIXED ORCHESTRATOR

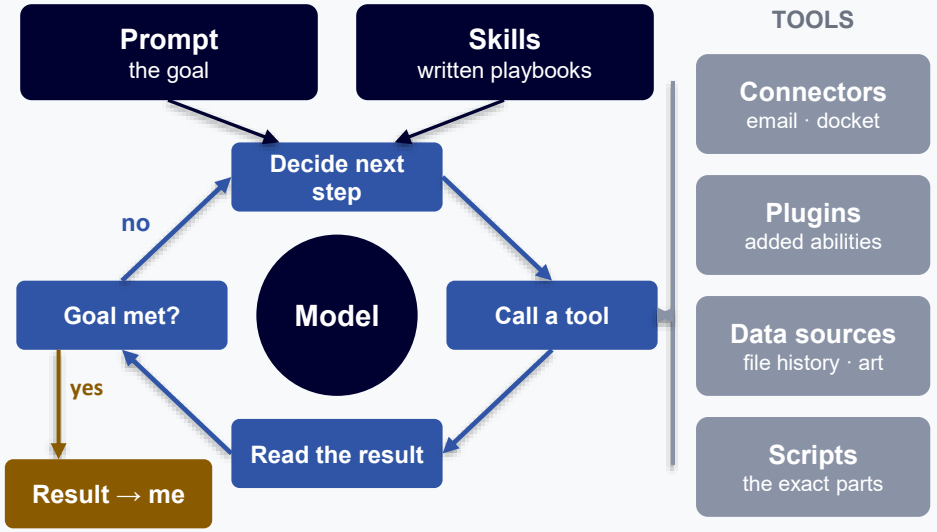
the designer draws the path



The code picks every step. It can only use what the designer thought to give it.

3 AGENT HARNESS

the model picks the next step



It loops until the goal is met. The prompt sets the goal, skills shape how the work is done, and tools set what it can reach.

A workflow can only use what you thought to give it.
An agent can go and get what you didn't.

Pick the tool for the task, and know which way it pushes you

Every tool nudges you toward a mode. Most of the method is noticing the nudge and deciding whether to follow it.

TOOL	WHAT I USE IT FOR	HOW IT FAILS	THE MODE IT INVITES
Chat assistant	Thinking out loud, quick drafts, explaining an unfamiliar technology	Knows only what you pasted in. You are its file room.	**Generate-and-use**
Legal and patent platforms	Structured tasks: upload the documents, pick the task, get the draft	Grounded in what you chose to upload. It can't go looking for what the workflow didn't anticipate.	**Delegate, unless you push back**
An agent with connectors and skills	Office actions, drafting chains, revisions, email, the docket	Can reach further, including into the wrong source. It needs rules about where the truth is.	**Amplify, or over-delegate**
Image models	Rough first layouts of figures	Plausible drawings with wrong numerals and lead lines.	**Explore**
Deterministic tools	Redlines, numbering, conversion, dates from the docket	They don't guess, and that's the point. Anything that must be exact belongs here.	**Delegate fully**
Dictation	Framing and review comments at the speed of speech	Mishears names and terms of art. Proofread those.	**Amplify: more context in**

Any tool can be used in any mode. This is the default pull.

In a chat window, you are the harness.

Perfect each step, then retire its review

Break the task up. Every review fixes the output and feeds the lesson back into the step that made it.

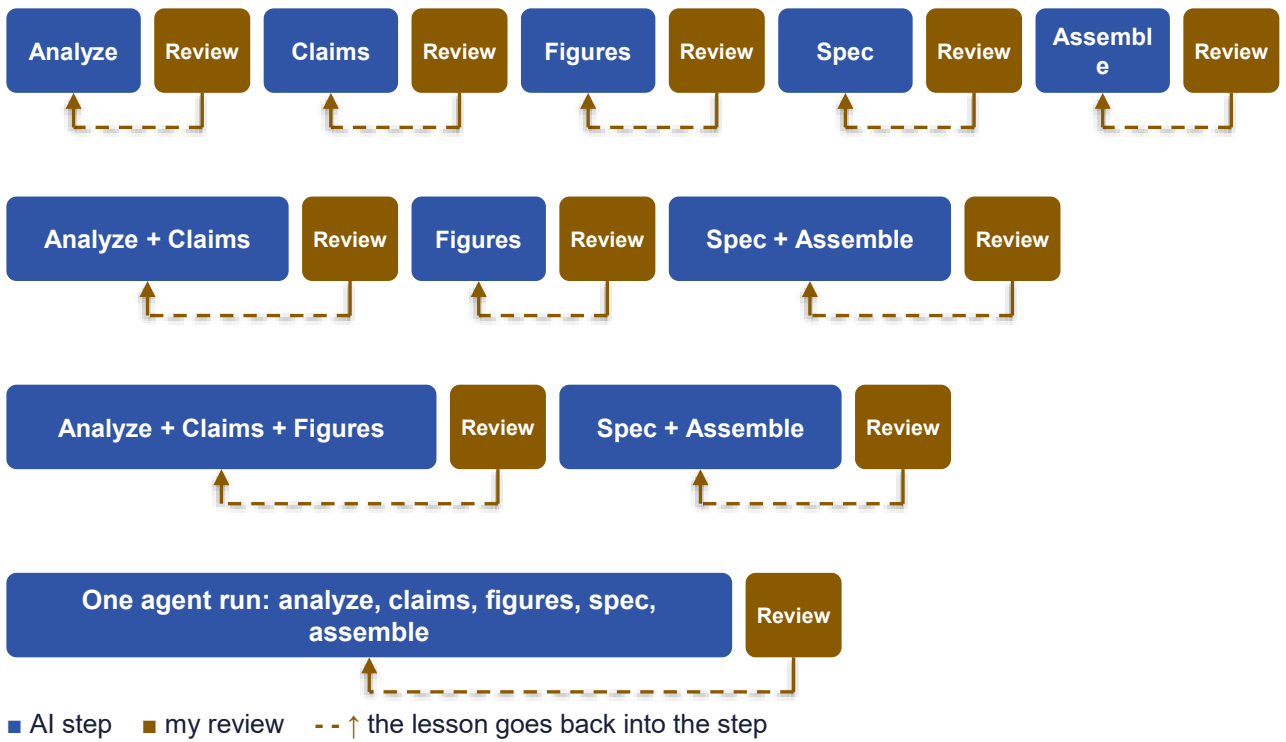
Day one
every step reviewed

After a few matters
proven steps combined

After many
fewer gates

Mature
one agent run, one review

OVER TIME



Every review does two jobs

1. **Fix this output.**
   The draft in front of me gets corrected.

2. **Fix the step that made it.**
   The lesson goes into the prompt, the skill or the process, so the next draft starts better.

A review retires when its step stops producing corrections. **The last review never does.**

Every review fixes two things: this draft, and every draft after it.

What amplifying looks like: one task, two paths

Same tools either way. What changes is how much I hand over, and where the gates sit.

AI tool me

Inventor's email
+ marked-up draft

The peek
Simple, or does it need me?

DELEGATE THE LOOP one dictated prompt

"Read the email, file the attachment, make the revisions, and draft my reply with a summary and a redline."

File attachment to the matter → **Revise** per every comment → **Redline** vs. the version they saw → **Draft reply** summary + redline → **Review** the redline → **Stage, then** I press send

AMPLIFY work through it together

"Present each issue with your recommendation and options. Walk me through the ambiguous ones. Write the plan; then revise in one run."

File attachment to the matter → **Present** each issue + options → **Decide** together, one by one → **Record** the revision plan → **Revise & redline** in one run → **Review, then** I press send

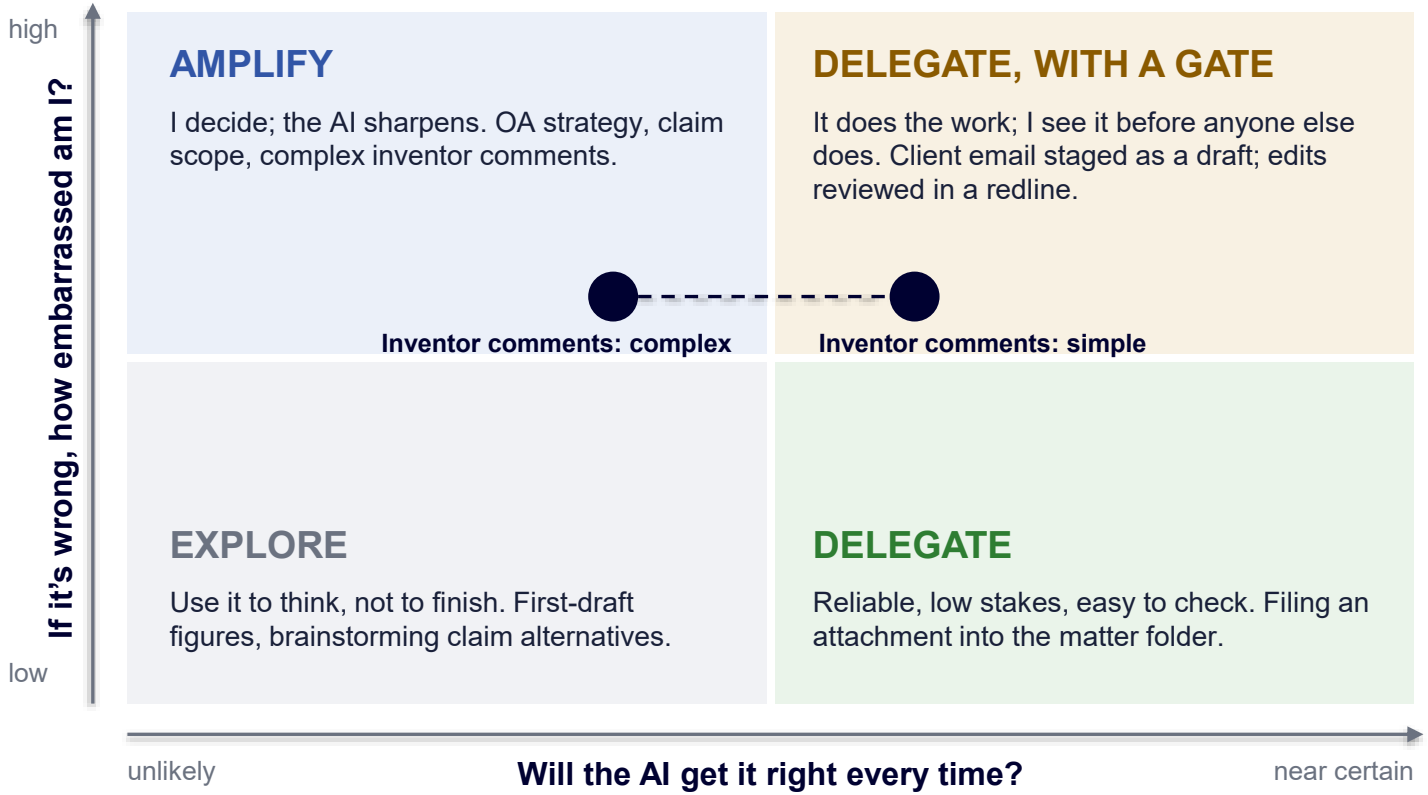
FRAME & GROUND · INTERROGATE · DECIDE · RECORD · VERIFY

Tags: the five-move method from Lucas's section, step by step. The agent does the legwork; every judgment call is mine.

I talk, it proposes, I decide. Then I review the redline, not the whole document.

Delegate or amplify? Let the task decide.

Two questions before I hand anything over. A third tells me how hard to look at what comes back.



Moves that make more of it delegable

1. **Give it a tool** for the part that must be exact
2. **Add a checkpoint** so I see the error before anyone else
3. **Break it up** delegate the pieces, keep the decision

The third question

If it's wrong, will I notice? A wrong citation looks exactly like a right one. The less visible the error, the harder I review.

Delegate the bottom right to buy the time to amplify the top left.

Delegating safely: the gate moves, it never disappears

Email in, judgment out. My agent drafts client email all day. It has never sent one.

HOW FAR I LET IT GO, ON EMAIL

1	Reads and summarizes	
2	Classifies, files attachments, logs	
3	Drafts the reply, grounded in the file and the docket	
4	Runs the checks and stages it in my Drafts	the agent stops here
5	Sends after my click	only I send
6	Sends on its own	never

Before it can reach my Drafts, it is checked

Green goes. Amber asks me. Red is refused.

- ✓ **Right people:** every address traces to the client's contact list; never guessed
- ✓ **Right attachments:** what this kind of email must carry for this client, and what it must never carry
- ✓ **Right thread:** replies stay in the conversation
- ✓ **Right dates:** restated from the docket, never computed
- ✓ **Right record:** matter email is captured to the docket system

Automate everything that only touches me.
Gate everything that touches someone else.

Where the time goes

AI buys the time. Who keeps it, and what it buys, is a choice.



Delegators bank the time. Amplifiers spend it on the work the tool can’t do.

Where the traps actually live

Analyze your AI task playbooks, and break down where AI helps, the traps, the amplifier moves, and what should never be delegated. Four examples below.

Claim drafting

Risk: Reads on the art, or silently narrower than the disclosure justifies.

Office-action responses

Risk: Argues the wrong thing; amendment creates new §112 exposure or concedes scope.

Foreign filing

Risk: Confident, specific, wrong national-law detail.

Any deadline

Risk: A computed date that is fluent and incorrect. Never accept an AI-computed date.

AI is reliable for structure, triage and first drafts.
Unreliable for judgment, scope, authority, dates, and client knowledge.

Grade the method, not just the document

When a tool can produce a competent deliverable unaided, grading the deliverable only measures the tool.

1	**Beat-the-baseline**	Save an example of unaided output from the tool. Review the delta the person added and the decision log, and not the document in isolation.
2	**Red-team drills**	In a mutually agreed upon experiment, hand over an AI draft seeded with real defects. Analyze whether the critique finds substance or chases cosmetics. Diagnoses apprentice behavior in minutes, at any level.
3	**Decision logs**	Encourage recording two or three lines on the calls that mattered on AI-assisted work product. Turns an invisible skill into a coachable one.

The non-negotiables

Professional responsibility sits underneath every phase and task.

NEVER DELEGATE

- Claim scope, strategy, and whether art teaches a limitation
- Legal conclusions, advice, and construction
- Materiality and duty-of-disclosure judgments
- Anything carrying your signature

UNIVERSAL CONTROLS

- Verify every authority against the source, as a summary is not the source
- Never accept a computed date
- Client information only in approved tools
- If you cannot say what it got wrong, stop

The training plan: five steps to a contained pilot

1	2	3	4	5
Build	**Launch**	**Practice**	**Grade**	**Scale**
Finalize the Phase 0 (no AI) task list, the playbooks, and the decision-log template	Junior cohort starts AI-off; practice group runs a first recalibration session	Guided introduction, then routine use of the method and playbooks on real tasks	Beat-the-baseline reviews and red-team drills score the method, not the document	Pilot results decide the wider rollout, and the same exercise moves into interviewing and reviews

This is a training plan first, and the policy decision comes after it works.

AI won't standardize our work.
It will magnify the differences between our people.

Everyone will have AI-literate practitioners. The firms that win will be the ones whose people apply real judgment inside a disciplined method, and who can see that judgment well enough to coach it, review it, and hire for it.

We already have the tools. The work is the method.

slwip.com

Thank you

Schwegman Lundberg & Woessner