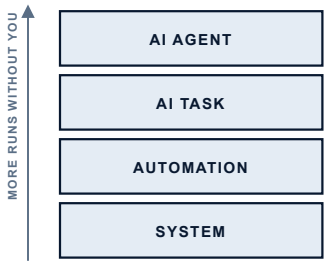
The Leverage Ladder

One practice task. Four levels. Nine questions to find which one it should be — without over-building it.

Pick one task that keeps landing on you or your front desk — confirmations, insurance verification, new-patient intake, following up on unscheduled treatment. Answer all nine questions below — one tick per row, the closest fit right now. Then add up each column. Your highest column is the level to build this task at. If two columns tie, the rule under the table decides.

The task I'm deciding about: _____



AI agent — AI chases a goal across many steps, picks its own path and tools, loops until done. Needs guardrails and a checkpoint.

AI task — one step where no rule fits: the model makes the judgment call from context you've given it (your method, your standards), and you confirm each output before it goes.

Automation — software runs it by fixed rules. Same steps, same kind of input, no judgment. It just happens.

System — a written process a person follows the same way every time. You still do the work; the doc makes it repeatable.

> **This isn't a ranking.** An agent is not a better system. Each level is the right call for a different kind of work, and a task can sit at one level for good — anything that runs on human judgment, trust, or accountability you shouldn't delegate stays a **System**, and that's a right answer. Reaching higher than the task needs is the most common — and most expensive — mistake.

One tick per row — the answer closest to true right now. Nine rows, nine ticks. Then total each column at the bottom.

	SYSTEM	AUTOMATION	AI TASK	AI AGENT
The process itself	○ Lives in my head — varies by who does it	○ Written out, identical every time	○ Written out, but one step needs a judgment call	○ Goal is clear; the steps change every time
The inputs each run	○ Every case is different	○ Consistent and structured — intake forms, fields in your practice software, known format	○ Messy — emails, voicemails, insurance documents — one at a time	○ Unpredictable — I don't know what I'll be handed
The decisions in it	○ Need my experience and context throughout	○ None — pure if-this-then-that	○ One call no rule can capture — made from context I'd give it	○ A chain of calls, each depending on the last
How often it runs	○ Rarely or irregularly	○ Constantly — many times a day or week	○ Regularly, usually in batches	○ On demand — each run is a mini-project
If it's done slightly wrong	○ A person catches it in the moment	○ Obvious and quick to fix — the rule was off	○ I catch it at the confirm step, in seconds	○ It could run several steps down a wrong path first
What I'd be handing over	○ Nothing yet — I'm still defining how it's done	○ A fixed recipe	○ One tricky step inside a fixed recipe	○ The whole "work out how to get this done"
Tools or steps to finish it	○ Just me and a checklist	○ One or two apps, always the same	○ A model call slotted into an existing flow	○ Several tools — which ones depends on the situation
Human judgment, trust, accountability	○ Central — patients expect a person here	○ Not needed — nobody minds who or what does it	○ Fine as a draft I sign off before it goes	○ OK to run semi-independently with a checkpoint
What I want from fixing this	○ Get it out of my head and consistent	○ Never touch this repetitive thing again	○ My judgment on one step, at speed, without doing it myself	○ Hand off the whole outcome, not just a task
COLUMN TOTAL — MUST ADD TO 9 →	_____	_____	_____	_____

My level (highest column): _____ The level I *want*: _____

Tie-breaker: if two columns tie, or if **System** has 3 or more ticks, treat it as a **System** first. Automation, AI task, and AI agent are all built on a documented, consistent process — without one underneath, they wobble.

BEFORE YOU BUILD IT — THE CHECK FOR YOUR LEVEL

Take the level your score landed on and run its four checks. If you can't tick all four, you're not ready to build it there yet — drop down a level, or close the gap first.

SYSTEM — before you rely on it

- ☐ The steps are written down, not just in my head
- ☐ One person owns it
- ☐ Someone else could run it from the doc and get the same result
- ☐ It's meant to stay human, or it's the next thing in line to automate

AUTOMATION — before you build it

- ☐ One specific event starts it — not a judgment call
- ☐ Inputs always arrive in a consistent shape
- ☐ Every rule is if-this-then-that — no "it depends"
- ☐ When it fails, the failure is visible, not silent

AI TASK — before you wire it in

- ☐ No rule fits this step — it takes a judgment call every time
- ☐ I can write down the context to make that call the way I would — my method, standards, edge cases
- ☐ Every output is confirmed before it goes — and it starts that way, loosening only as it proves out
- ☐ It's one bounded step, not "work out how to get this done"

AI AGENT — before you turn it on

- ☐ An automation genuinely can't do this — every run needs real judgment, not just rules
- ☐ It's one clearly-scoped job — goal, tools and limits written down, not "do anything"
- ☐ It knows the moment it should stop and come get me — a real checkpoint, not just a log
- ☐ The system, automations and written context under it are already built — the agent is the last thing you add

The four aren't ranked. Each is the right call for a different kind of work, and a task can stay at one level for good. The costly mistake is reaching past what the task needs — an agent doing a job an automation would do more reliably, or automation bolted onto a process nobody wrote down. One more check for a practice: anything that touches patient information needs a HIPAA-compliant tool and a signed BAA before it's automated or handed to AI. Bring this filled-in ladder — and your wider list of tasks — to your Revenue Waterline Audit call. We place each one, then put them in the order that frees the most of your and your team's time first.

© 2026 Digital Vision · The Leverage Ladder — Dental