

Putting AI to Work in Operations

Where AI genuinely earns its place in a workflow — and where it's just expensive noise. A grounded starting point from someone who builds the tools, not just talks about them.

Most AI-in-operations initiatives stall not because the technology fails, but because it was pointed at the wrong problem. The skill isn't prompting — it's judgement about where AI fits. This guide is about that judgement.

Is this a good fit for AI?

Before automating anything, hold the task against these two lists. If it lands mostly on the left, it's a strong candidate. Mostly on the right, and you'll spend more time correcting the AI than the task ever cost.

Good fit

- High volume, repetitive, draining people's time
- Language-heavy: reading, summarising, drafting, classifying
- A "good enough" draft a human then checks is genuinely useful
- Clear inputs and a checkable output
- Mistakes are visible and cheap to catch

Poor fit

- Rare, high-stakes decisions with no room for error
- Needs a guaranteed-correct answer every time
- The real problem is a broken process, not a slow one
- No one will own checking the output
- Mistakes are silent and expensive

Five patterns that reliably pay off

A Triage & routing. Read an incoming ticket, email, or request and sort it — category, priority, who should handle it. Removes the dull first-pass that clogs every support and ops queue.

B Extraction. Pull structured data out of messy documents — invoices, forms, PDFs, logs — into a clean table. Quietly one of the highest-return uses in operations.

C Drafting. First-draft the report, the reply, the summary, the update. The human edits rather than starts from blank — often a 10x speed-up on routine writing.

D Knowledge retrieval. Answer "where's the policy on X?" from your own documents, so people stop interrupting each other for things already written down.

E Sense-checking. A second pair of eyes on a plan, a dataset, or a decision — surfacing what a tired human might miss, without replacing the human's call.

How to run your first pilot

Pick one task from the good-fit list that's narrow, painful, and low-risk. Set a plain success measure before you start ("cuts triage time by half"). Keep a human in the loop checking outputs. Run it for two weeks against the old way, side by side. Then decide with evidence, not enthusiasm.

WHERE IT REALLY STALLS

Most AI projects don't fail at the pilot — they fail at adoption. The tool works, but it never becomes how the team actually operates. Design for the day-to-day habit from the start, not the demo. That transition is the whole game.

Three pitfalls to avoid

- **Automating a broken process.** Fix the process first. AI on top of a mess just makes a faster mess.
- **No owner for the output.** If no one is accountable for checking what the AI produces, quality quietly erodes until trust collapses.
- **Boiling the ocean.** One narrow win that people feel beats a grand platform nobody uses. Earn the next step.