



Before You Automate

Prioritize AI ideas before you spend money on the wrong project.

SEPTEMBER 2026

WRITTEN BY

Allison Marx

technical-humanist.com



The AI use case prioritization workbook

How to use this workbook

You probably have more AI ideas than time. Every one of them sounded promising in the meeting where it came up. This worksheet exists to slow that down just enough to make a good decision.

Pick one proposed use case. Work through the steps with the people closest to the work. You'll land on one of five honest answers: pilot it now, design it further first, fix a foundation problem first, keep it human-led, or park it.

Use this when...

- You have several AI or automation ideas competing for attention.
- A vendor demo has created pressure to choose quickly.
- Your team is experimenting unevenly, with no shared way to compare ideas.
- You need one defensible recommendation for what should happen next.

This will not...

- It will not guarantee ROI. No worksheet can.
- It will not replace legal, security, or technical review for consequential use cases.
- It will not turn a poorly understood workflow into a safe automation just by filling in boxes.

Why this framework exists

The majority of generative AI pilots fail to produce measurable business impact. That's because most teams start in the wrong place. They pick a tool first, then go looking for a problem to justify it. That's backwards, and it's exactly how AI pilots become AI graveyards: interesting experiments that never connect to an outcome anyone actually cares about.

Start here

This worksheet flips the order. Instead of asking "what can this tool do," it asks whether a specific idea holds up once you look closely at **value, workflow, data, risk, feasibility, and the people who have to live with it.**

USE CASE: What are we trying to make easier? What problem is this solving?

CURRENT WORKFLOW: How is the work being done now?

INPUT & OUTPUT: What information does this depend on? What's the expected result?

TIMELINE: When does this need to be solved by, and why?

OWNERSHIP: Who is responsible after this ships? What does a person still need to decide?

HUMAN QUESTION: Who benefits from this? Does anyone's role change if this is automated?

The six lenses (part 1)



1 Business Value

Does this matter enough to spend time on?

Top questions to ask:

- Which customer, team, revenue, cost, risk, or decision outcome actually improves?
- How often does this problem occur, and what does it cost in delay, rework, or error each time?
- Is the improvement meaningful enough to justify real attention, not just mild convenience?

Evidence:

Answer: ☐ Clear ☐ Partly Clear ☐ Unclear

2 Workflow Clarity

Do you actually understand how this work happens today?

Top questions to ask:

- Can the team explain the current process from trigger to outcome, without hand-waving?
- Are the handoffs and exceptions known, or does everyone do it slightly differently?
- Is this workflow stable enough to improve, or does it change week to week?

Evidence:

Answer: ☐ Clear ☐ Partly Clear ☐ Unclear

3 Data & Context Readiness

Does the information this idea depends on actually exist, and can you trust it?

Top questions to ask:

- What information would the AI need to do this well?
- Are the sources current, accessible, documented, and properly authorised?
- Are there known gaps, conflicting definitions, or missing context that would undermine the output?

Evidence:

Answer: ☐ Clear ☐ Partly Clear ☐ Unclear

The six lenses (part 2)



4 Risk & Human Impact

What happens when this is wrong, not just when it's right?

Top questions to ask:

- What happens if the output is incorrect, biased, late, or misregulated?
- Does this touch customer trust, privacy, money, employment, health, eligibility, or another consequential outcome?
- Who can review, override, pause, or correct it before harm reaches someone?

Evidence:

Answer: ☐ Clear ☐ Partly Clear ☐ Unclear

5 Feasibility & Operating Cost

Can you actually build, run, and afford to maintain this?

Top questions to ask:

- Can the needed systems connect through supported, secure interfaces, or does this require a new platform entirely?
- What will implementation, licensing, monitoring, and vendor management actually cost to keep running, not just to launch?
- Is there a real owner who can run this after launch?

Evidence:

Answer: ☐ Clear ☐ Partly Clear ☐ Unclear

6 Adoption & Change Effort

Will the people who have to use this actually use it?

Top questions to ask:

- Will the people doing the work use this process, or choose to route around it?
- What skills, training, permissions, or new working agreements does this require?
- Does this remove a real burden, or does it just add another approval layer people resent?

Evidence:

Answer: ☐ Clear ☐ Partly Clear ☐ Unclear

Choose your action path



Work through all six lenses first.

A single serious flag, especially in Risk or Data, should outweigh strong scores everywhere else.

<i>If this is true</i>	<i>Choose this path</i>
Value is meaningful; workflow, data, owner, and review are all clear; risk is manageable	Pilot now. The guardrails already exist. You're testing a process you understand, not inventing one.
Value is meaningful, but the workflow, exceptions, or success criteria are still fuzzy	Design first. The idea is good, but nobody has defined what "done well" looks like yet.
Sources, definitions, permissions, ownership, or data quality are unclear	Fix the foundation first. The AI isn't the actual problem. Something upstream needs attention first.
The consequence of error is high and there's no real review or control in place	Keep human-led. Judgment or consequence make this the wrong place to hand off control right now.
Value is weak, timing is poor, or no real owner exists	Park it. Not the next move, not a bad idea forever.

Record your next move

Decision made:

Evidence supporting the decision:

Primary unknown to resolve:

Smallest next action:

Named owner:

Review date: