


FIELD BRIEFING

Where AI earns its place

Choose the task before the tool.



Which piece of work should we try first?

A business rarely needs an AI strategy for every task at once. It needs a useful first experiment. Start with work that has clear inputs, a recognisable output and someone who can judge whether that output is right. A capable model cannot compensate for an undefined decision.

Explanatory schematics and worked examples are illustrative. No live model, observed client results or measured market performance are presented.

Read the mechanism.

Separate three things: the task, the information available, and the consequence of an error. Drafting a reply from an approved returns policy is different from sending the reply. Summarising a meeting is different from deciding what was agreed. Extracting invoice fields is different from authorising a payment. The same model can be useful at one stage and unsuitable at another.

Consider an invented small repair business. The owner spends time finding the right service instructions, drafting responses to common enquiries and reconciling disputed bills. A search-assisted draft from current instructions could be a useful pilot. A disputed bill needs a responsible person to assess the facts. The distinction comes from the workflow, not a universal list of tasks AI can or cannot do.

Use a four-question screen. Is the task frequent enough to observe? Can you obtain authorised, representative inputs? Can a person identify a good result without doing all the work again? Can a failed attempt be stopped or corrected cheaply? Unknown answers are research tasks before they are automation opportunities.

Do not confuse visible speed with end-to-end benefit. A thirty-second draft that takes fifteen minutes to repair may not help. Record the existing process, then count preparation, review and rework in the trial. Pay attention to the people receiving the output: if they need to verify every detail from scratch, the work may simply have moved.

There are often three reasonable outcomes: use ordinary software or a template, use AI to assist a person, or explore a more automated workflow with explicit permissions. Keeping a task manual can also be a sound decision. A pilot is useful when it helps choose between these outcomes.

Technical references: Microsoft · Responsible AI; Anthropic · Building Evals. Links and source scope appear in the final section.

Try the idea.

INVENTED TEACHING EXAMPLE

A repair business chooses a first task

The team chooses an enquiry draft based on current service instructions. The reviewer checks the policy, promised dates and recipient before sending. A disputed bill stays with the responsible decision-maker. The exercise compares preparation, review and correction, not just drafting speed.

Go deeper.

Separate the stages

Map input collection, transformation, review and execution separately. A task can be suitable for drafting while its execution remains unsuitable for automation. Evaluate the workflow you will actually run.

Avoid a universal opportunity score

Frequency, input rights, reviewability and failure consequence are not interchangeable quantities. This guide does not combine them into a pseudo-scientific score. A failed prerequisite can outweigh an apparently attractive speed gain.

What to record

Save source examples, the approved task definition, observed preparation/review time and failure categories. Use the same acceptance rules when comparing the existing and proposed processes.

Keep the evidence close.

Reviewer

The person who can judge whether an output meets the task's requirements.

Baseline

The current process measured before a proposed change.

Recovery

How a failed or mistaken step is stopped, corrected or reversed.

Accepted task

An output that meets the stated criteria after required review.

Sources · accessed 4 October 2026

[Microsoft · Responsible AI](#)

[Anthropic · Building Evals](#)

Technical and security references support the mechanisms, not our fictional results. KAF's worksheets and examples are original editorial recommendations. No current provider pricing, comparative ranking or legal compliance is asserted. Verify the exact release, configuration and policy for your use.

Practical AI Insights · First editorial edition · 4 October 2026

Where AI earns its place · [Web edition](#)