COPILOT → COLLEAGUE

HOW TO GET BETTER RESULTS FROM AI

The PINPOINT Prompt Engineering Framework v1.3

© 2024–2026 Ahmed Muzammil · CC BY-SA 4.0

START HERE

Before you blame the model, inspect the request.

Better prompting is not about finding magic words.

It is about specifying the work.

WHY ANSWERS GO WRONG

Most bad answers come from the request, not the AI.

Too vague

It fills the gaps with guesses and gives you something generic.

Too tight

One example or rule is followed so literally that it stops thinking.

Made up

With nothing solid to work from, it can produce something that sounds right but isn't. This is often called hallucination.

The rest of this deck is how to close each gap.

EVERYDAY WORK

PIN

Start with the smallest amount of structure that works.

P: Purpose

What do you want to achieve, and for whom?

I: Introduce Context

What does AI need to know? What is the source of truth?

N: Narrow the Ask

What exactly should a good result look like?

Practical rule of thumb: PIN is enough for roughly 80% of everyday prompting.

WHEN CONSEQUENCE RISES

PINPOINT it.

Add structure when the work needs greater **control, verification or confidence.**

P
Purpose
What it's for, and who it's for

I
Introduce Context
What it needs to know

N
Narrow the Ask
What a good answer looks like

P
Persona, Ethics, Values & Role
Whose shoes it stands in

O
Outline Examples
Show it good and bad

I
Induce Verification
Make it check its work

N
Normalise & Standardise
Review the output and version the prompt, so every change can be undone.

T
Tweak & Test Again
Change one thing and try again

THE FIRST THREE

Specify the task.

Purpose

Define the outcome and audience.

Introduce Context

Give it the documents and data it needs, and say which one is the final word. Use the latest approved version, and share only what the reader is allowed to see.

Narrow the Ask

Say how long, what format, what's in and out of scope, and when it should stop and ask you.

> **SOURCE RULE** If two sources disagree, say which one wins, or ask the AI to point out the clash.

ORIGINAL

Summarise this report.

→

PIN

Summarise this quarterly report for the executive committee. Focus on the three issues requiring a decision. Use the attached report as the source of truth. Five bullets and a recommendation.

TWO EVERYDAY EXAMPLES

Same idea, two different jobs.

Priya, bank relationship manager

Weak request

Summarise this client's account.

Better request

Summarise this client's last six months for my review meeting on Friday. Use only the core banking statement attached. Five bullets: changes, risks, one question to ask.

Dana, clinic front desk

Weak request

Write a reminder message.

Better request

Write a friendly reminder for tomorrow's 9am appointments. Use the clinic's booking list as the only source. Under 40 words, no medical details, include the cancel number.

Both say what it's for, which source to trust and what good looks like.

ADD JUDGEMENT & CLARITY

Persona and Examples

Persona, Ethics, Values & Role

Set a useful professional stance, responsibilities, principles and boundaries. Persona is not cosplay.

Outline Examples

When showing is easier than explaining, show what good looks like and, where useful, what bad looks like.

Examples are optional. Use them when they remove ambiguity.

ASK IT TO CHECK ITS OWN WORK

Induce Verification

What needs to be checked before this result can be relied upon?

Evidence & sources	Assumptions & uncertainty
Facts & calculations	Constraints & Purpose

The level of verification should match the consequence of getting it wrong.

LOW CONSEQUENCE

Basic factual check

MEDIUM CONSEQUENCE

Show sources, list assumptions, say what it isn't sure of

HIGH CONSEQUENCE

Show sources, say what it isn't sure of, raise doubts, and a named person signs it off

FOUR LABELS FOR EVERY CLAIM

Make uncertainty visible.

[Verified Source]
A source you trust says so

[Inference]
It worked this out from the evidence

[Assumption]
It's taking this as true without proof

[Unverified]
It can't back this up yet

Uncertainty can be acceptable when visible. **If a doubt could change the decision, raise it with someone before acting.**

USE IN A PROMPT

“For each material claim, label it as [Verified Source], [Inference], [Assumption], or [Unverified]. Cite the source where available. If an unverified or assumed point could materially change the conclusion, flag it for review rather than guessing.”

TRUST NEEDS TWO CONTROLS

Checked is not the same as approved

What the AI checks

Induce Verification

Ask AI to check evidence, assumptions, uncertainty, consistency and constraints.

What you decide

Normalise & Standardise

You review the output, fix it and bring it into your house format before anyone uses it. Keep a version history of the prompt, so every change can be traced and undone.

AI verification does not remove human accountability.

THE PERSON WHO CHECKS

A reviewer only helps if they can say no.

1983: he broke the loop

Stanislav Petrov's early warning system reported incoming US missiles. Protocol said escalate. He questioned it instead, and it was a false alarm caused by reflected sunlight. He was willing to say no.

A rubber stamp isn't a check

The real danger is people who stop thinking because the answer sounds confident. Signing off is a responsibility, not a formality.

People get tired

We make worse calls when we're tired, rushed, or late in a long run of decisions. One more approval step doesn't fix that.

About 17% → 5%

Share of planted dangerous commands people caught, falling after about 50 approval requests.

13.6% vs 89%

Dangerous commands caught by the people approving, compared with an automated check.

These numbers come from AI coding sessions, and we're assuming they carry over to other kinds of review. Source: claude.com/blog/auto-mode-default-in-claude-code · 7 August 2026 · 1,053 testers.

HUMAN REVIEW

Is it fit for use?

1. Did it achieve the Purpose?	2. Is it grounded in the source of truth?
3. Did it follow constraints and permissions?	4. Any unsupported claims or material uncertainty?
5. Is it suitable for the intended decision or workflow?	**Then decide:** use it as it is, fix it, or send it back.

> A tired reviewer approves almost anything. If you're checking many answers a day, check the high-stakes ones properly, spot-check the rest, and ask the AI to point out what it's least sure about.

ITERATION

Tweak & Test Again

Test → Learn → Change → Retest

One good answer isn't enough.

Review what worked. Identify what missed the mark. Then change one meaningful variable at a time (the context, the constraint, the example, or the verification) and test again.

For repeatable work, test across normal cases, edge cases, and failure cases. Look for consistent performance, and learn where it breaks.

The goal isn't a prompt that worked once. It's a prompt you know when to trust.

Purpose → Context → Ask → Examples → Verify → Review → Tweak 

THE FULL JOURNEY



PIN the task.
PINPOINT the result.
BLUEPRINT the colleague.

CHEAT SHEET

The eight steps on one page

The first three, highlighted, are your everyday steps. Add the rest when the work matters more.

Purpose

What's it for, and who's it for?

Introduce Context

What does it need to know, and which source wins?

Narrow the Ask

What does good look like, and when should it stop?

Persona, Ethics, Values & Role

Whose shoes should it stand in?

Outline Examples

Can you show it good and bad?

Induce Verification

What must it check before you rely on it?

Normalise & Standardise

What will you check before use, and how will you version the prompt?

Tweak & Test Again

What one thing will you change next time?

TAKEAWAY

Don’t make every prompt longer.

Make delegation clearer.

Start with PIN. Add the PINPOINT elements the work actually needs.

WHAT COMES NEXT

When the AI needs to act, not just answer

PINPOINT helps you ask an AI well, one task at a time. You read the result and you decide what happens next.

Once an AI starts doing things in your systems, like looking things up, updating records or sending messages, you have to decide in advance what it's for, what it may touch and who checks its work.

BLUEPRINT is the companion framework for that. It walks through those decisions one letter at a time, in plain language, with two example agents: one in a bank and one in a clinic.

Read it at ahmedmuzammil.com/blueprint or download the BLUEPRINT PDF (v2.4).

PIN the task. PINPOINT the result. BLUEPRINT the colleague.