

A PRACTICAL GUIDE

Collaborating with AI

For business teams that want to get real work
out of AI, not just experiments.

CRAFT the input. **OWN** the output.

01 Who this is for

This is a guide for business teams that want to get real work out of AI.

This is not a tool review. Not a list of prompt tricks. It is a practice for the work that has your name on it. Customer communication. Proposals. Status updates. Research. The work where "good enough" is not when you finish typing. It is when someone you respect reads it and agrees.

Two things you need to get that work done well. A way to communicate with the AI so it produces something useful. And a way to judge what it produced before your name goes on it.

CRAFT the input. OWN the output.

A VENDING MACHINE, NOT A COLLEAGUE

The vending machine problem

Most people use AI the way they use a vending machine. Put in a question. Get out an answer. Move on.

The answer is too long, or too generic, or aimed at nobody in particular. The tone is off. It says things you already know. It says things you cannot verify. So you conclude that AI is not quite ready for your work and go back to doing it yourself.

The tool is not the problem. The input is.

A vague prompt produces a vague response. A clear, structured prompt produces something you can actually use. But "write better prompts" is not real advice. Teams need a framework. The same way they need patterns for running meetings or writing documents. Something repeatable. Something that connects to the discipline they already practice.

And input is only half the problem.

AI fails differently than a colleague. At worst, a colleague admits they do not know. AI will invent a statistic instead, or give you a confident answer to a question it never actually looked into. The output will read just as polished either way. So the other half is knowing how to judge what you got back before you act on it.

This guide covers both: input and judgment. Each is a habit. Together they are a practice.

02

CRAFT: Structured Input

CRAFT is a five part framework for getting better results from AI on high value tasks. Context. Role. Ask. Fit. Tune.

You already do most of this when you brief a teammate. You tell them the situation. You name the perspective that matters. You make the request specific. You say who the output is for. And you stay in the loop while they work. The only thing new here is treating the AI the same way.

C **Context** / Situation, constraints, audience

The single biggest lever. Without it, the model guesses and gets it wrong.

R **Role** / The expertise or perspective for the AI

Shapes the lens for everything that follows.

A **Ask** / One clear, specific request

Sharpens what you actually want from the model.

F **Fit** / Format, length, tone, audience

Turns the output into something that belongs in your workflow.

T **Tune** / Iterate. Refine. Redirect.

Where the real lift comes from. Ask the AI to ask you clarifying questions before it starts.

The rest of this section walks through each letter using one realistic task: writing the opening of a proposal to a new prospect. The same prompt gets better each time we add a letter. The point is not any one output. The point is feeling how much more useful AI becomes when you give it what it needs to help you.

AN EXAMPLE

The scenario

You work at a custom fabrication company. Twenty years of experience in precision components for aerospace. A referral just surfaced a promising opportunity: a medical device manufacturer is standing up a new product line and is looking for a fabrication supplier. You have the precision experience. You do not have direct medical device history. You have three days to send a proposal.

You sit down to draft the opening. You type the obvious prompt:

```
Write a proposal intro.
```

What you get back reads like a company brochure. Three paragraphs about manufacturing excellence, commitment to quality, and decades of experience serving demanding industries. It is not wrong. It is aimed at nobody. You could send it to any prospect in any industry and it would land the same way. Which is the problem.

This is the first instinct to push past. The generic response is not the AI's limit. It is what happens when the AI has nothing to work with except the shape of the request.

This is where CRAFT comes in.

c is for Context

Context is the situation, background, and constraints. It is the single biggest lever for getting useful output. Without it, the model guesses. And guesses wrong.

Context includes what you are working on and why, what you have already tried or decided, and constraints the AI would never know. Timeline. Audience. Company specifics. The opportunity in front of you.

Add context to the prompt:

```
We are a custom fabrication company with twenty years of precision work in aerospace. A referral surfaced a new prospect: a medical device manufacturer launching a product line that needs a fabrication supplier. We have the precision chops but no direct medical device history. I need to draft the opening of a proposal to them.
```

```
Write a proposal intro.
```

The output shifts. It now talks about precision. It acknowledges a move into medical devices. But it is still overcooked. It spends too long on the company and not enough on the reader.

R is for Role

Role tells the AI what expertise or perspective to bring. It shapes the lens for everything that follows.

Add role:

```
(same context as above)
```

```
You are a senior account executive writing to a VP of Operations at the prospect.  
They have never worked with us.
```

```
Write a proposal intro.
```

Now the voice is different. It talks to someone, not about something. It leads with what a VP of Operations actually cares about: capability, reliability, and fit for their schedule.

Role matters most when the perspective matters. For a quick task, you can skip it. For anything where the reader's point of view shapes the message, name it.

A is for Ask

Make one clear, specific request.

"Write a proposal intro" is vague. What does an intro mean? One paragraph? A full first page? Does it include pricing framing? A statement of interest?

Sharpen the ask:

```
(same context and role)
```

```
Write the opening three paragraphs of the proposal. Enough to establish credibility  
and make them want to read the rest. Do not try to cover pricing, scope, or timeline.  
Those come later in the document.
```

A focused ask gets a focused answer. You also save yourself a round of deletion, because the AI is not trying to write the whole proposal in one shot.

F is for Fit

Fit shapes the output to match your actual need. Tell the AI who the output is for, what format you want, and how much detail matters.

Without Fit, you get the AI's default. Usually long, faintly formal, and aimed at nobody. Fit is where you turn the output into something that belongs in your workflow.

Add fit:

```
(same context, role, and ask)
```

```
Professional but not stiff. Acknowledge that we are new to medical devices without  
dwelling on it. Lead with what is relevant about our aerospace precision work, not  
with what is missing. One page, three paragraphs, that is it.
```

The output tightens. The tone becomes warmer. The acknowledgement of the gap is short and confident instead of apologetic.

You could stop here. Many prompts will. But the last letter is where the real lift comes from.

T is for Tune

Tune is iteration. It works in two directions.

Before the AI writes anything, you can front-load the thinking:

```
(same full prompt)
```

```
Before you write this, act as a thought partner. Ask me one question at a time.
```

The AI asks the first question. Something close to:

```
What specifically made them call us? Was the referrer someone they trust, or was it  
something we produced or did?
```

You answer. The referrer is a consultant the VP worked with at her last company, and she asked specifically for "a shop that takes precision seriously."

The AI uses that answer and asks the next thing:

So the opening probably needs to earn the referrer's weight, not just name it. What does this VP care about in her current role? Cost, quality, schedule, or qualifying new suppliers at all?

Now you are thinking. Not just answering. The second question only made sense once the AI had the first answer. And notice what happened. You came in to get a proposal written. You are now in a conversation about what actually matters to the reader of the proposal. The prompt surfaced a gap you would not have thought to fill.

That is the hinge of this whole framework. You would not have thought to put any of that into a brief to a junior colleague either. But each answer is exactly what you need to position yourself well, and the AI pulls them out of you one question at a time.

The AI started being something closer to a thought partner.

Tune also works after the first response. "Too formal." "You missed the part about their launch timeline." "Try that again, shorter, and lead with the referrer." Most people accept the first output. The best work comes from two or three rounds of feedback.

The full CRAFTed prompt

- C** We are a custom fabrication company with twenty years of precision work in aerospace. A referral surfaced a new prospect: a medical device manufacturer launching a product line that needs a fabrication supplier. We have the precision chops but no direct medical device history.
- R** You are a senior account executive writing to a VP of Operations at the prospect. They have never worked with us.
- A** Write the opening three paragraphs of the proposal. Enough to establish credibility and make them want to read the rest. Do not try to cover pricing, scope, or timeline.
- F** Professional but not stiff. Lead with what is relevant about our aerospace precision work, not what is missing. One page, three paragraphs, that is it.
- T** Before you write, act as a thought partner. Ask me one question at a time.

03 OWN: Evaluated Output

You have a stronger draft now. A real one, shaped for a real prospect. But it is not ready to send.

Getting better output is only half the problem. The other half is what happens after the AI responds. Most people read the output, decide it looks reasonable, and use it. That works until it does not. Until a confident-sounding claim turns out to be wrong. Until a number that felt right has no source behind it. Until something you shared has a line in it you cannot explain.

Three failure modes

Most problems with AI output trace to one of three patterns.

Over-trusting. Accepting output without checking, especially anything quantitative or current. The fastest way to embarrass yourself.

Under-using. So skeptical of every response that you redo the work from scratch. That is expensive double entry with none of the benefit.

Not iterating. Treating the first output as the final output, then concluding the tool does not work. That is the person who tried AI once, got something generic, and decided it was not for them.

OWN is how you stay out of all three. Three steps, each one a habit worth building.

O Observe / Did the AI research, or recall?

Watch how the output was produced before you evaluate it.

W Weigh / How much should you verify?

Match effort to stakes and work type.

N Name / Would you put your name on it?

If you share it, you defend it.

The hardest part is not the prompt. It is evaluating the output for your use.

is for Observe

Before you evaluate the output, understand how it got there. Modern AI can do real work, or it can just write. Both will sound equally confident.

JUST WROTE

Training data. May be a year out of date or invented entirely.

SEARCHED

Grounded in actual information. Sources you can trace.

Back to the proposal intro. The AI asked clarifying questions, used what you told it, and composed paragraphs from general writing patterns. It did not look up the prospect or verify a single claim about the medical device market. It only has what you gave it.

That is fine for framing. It is not fine if a specific industry claim slipped in. A market size. A growth rate. A regulatory requirement. That line needs verification or cutting before the proposal ships.

A different kind of example

The Observe beat bites hardest when the task is research. Imagine a short memo on data center cooling trends. The AI produces three confident pages with percentages, market sizes, and industry quotes.

If it just wrote, those numbers may be stale or invented entirely. "I did not see you do any research" is a valid and useful thing to notice. It is a process review, not a content review.

If it searched, ask for the sources. Open the ones that matter. Confirm the numbers.

Before you judge the output, notice how it was produced.

W is for Weigh

Not every output needs the same level of scrutiny. A personal working note and a customer deliverable are different animals. Treating them the same wastes time in one direction and creates risk in the other. Two questions calibrate your verification.


What is at stake?

- Personal working note or exploration: quick sanity check.
- Shared with your team: structural review.
- Customer-facing or leadership-facing: source verification on every claim.

What type of work is it?

- Qualitative work (writing, structure, framing): the AI is generally reliable here. Review the logic and the tone.
- Quantitative work (numbers, data, percentages): the AI is unreliable here. Verify everything numeric.
- Factual work (verifiable claims, citations, current data): the AI may invent or be stale. Confirm it actually searched rather than recalling.

The intersection gives you a verification protocol:



Do not over-verify low-stakes work. That is waste. Do not under-verify high-stakes work. That is risk. Calibrate.

N is for Name

The moment you share AI-generated output, it carries your name. This is not about pride. It is about professional responsibility.

"The AI said" is not a defense. If your customer or your leadership pushes back on something you sent, "I can explain the reasoning" is the only answer that holds up.

Two practices make ownership real: distilling the output, and pressure-testing it before you send.

Volume is free. Clarity is the work.

It is easy to generate a fifteen page document. It is harder, and more valuable, to distill that down to the page that is actually digestible. The act of distilling is the act of owning. You cannot condense something you do not understand.

Use Tune to drive the cut:

Distill this to the points that actually matter. Before you cut, ask me one question at a time: who will read it, what decision they need to make, and what they already know.

The defense test

Before you share, run a simple test. Could you defend every claim in this document, out loud, in a meeting?


If not, it is not ready.

You can run the same test through the AI:

Play the role of a skeptical reviewer. Walk me through the three weakest claims one at a time. For each, ask me to defend it, then flag the ones I cannot defend without verifying.

AI makes you faster. The standard for what carries your name does not change.

04 Putting it together



Before you prompt, consider: Context, Role, Ask, Fit, Tune. Give the AI what you would need from a colleague if you were taking on this task.

Before you read the output, Observe. What did the AI actually do to produce it?

Before you verify, Weigh. Match your effort to what is at stake.


Before you share, Name. Distill it. Can you defend every claim?

The framework is simple. The habit is what takes repetition. Start with the next piece of work that has your name on it.

05 Where this goes next

Everything in this guide is the personal practice. CRAFT and OWN are habits that help one person get better work out of AI and take responsibility for what they use.

AI often exposes where real workflow opportunities live: in scattered documents, disconnected tools, personal habits, and the judgment of a few experienced people. CRAFT and OWN help one person work better. Helping a team build a shared practice around that judgment is a different conversation. Building the systems that integrate it into how the business runs is another.



RoleModel is a collaborative software development team. Since 1997, we have refined our [approach](#) to help partners turn workflow opportunity and hard-earned expertise into working systems. Our developers and non-technical teammates practice these habits daily, and have trained others to do the same. Process first, software second. AI is a new lens on a familiar craft.

Wherever you are on the path, the next step starts with a conversation.

WHAT'S NEXT

Free AI Consultation

Our goal is to listen, advise, and figure out the right next step together.

If we are the right partner for what you need, we will say so.

If we are not, we will tell you and point you to resources that can help.

Let's talk →