

Learn to direct and review AI work

AI can produce a first version. A useful skill is knowing what the result should do, spotting the gap and giving specific feedback.

What the occupation research says

Anthropic's March 2026 research combines O*NET tasks across around 800 occupations with observed Claude usage and theoretical task exposure. Programmers and customer-service representatives are among the most exposed in its measure. The report measures task coverage; it does not say these whole jobs disappeared. Observed use remains below theoretical capability.

The examples below are practical teaching workflows, not measured case studies from the paper. Directing and reviewing is the skill this guide recommends practicing; the study does not establish that most work is now review.

Example 1: a programmer

Give AI the issue, relevant code, reproduction steps, expected behavior and constraints. Ask it to explain the cause, propose a patch and identify meaningful tests.

AI can help draft code, explain a function and suggest tests. The person still owns the requirements, architecture tradeoffs, edge cases and release decision. A passing demo is not enough if the fix breaks another path.

Specific feedback: "The fix works for new users but breaks existing accounts. Preserve the existing-account behavior, change only the relevant path and add a regression test that fails before the fix. Explain any remaining uncertainty."

Review checklist: Does the change reproduce and fix the original problem? Does it preserve established behavior? Are tests checking the requirement rather than copying the implementation? Are sensitive inputs and failure cases handled? Can a reviewer understand the change and its limits?

Example 2: customer service

Give AI the customer message, relevant account facts and the current policy. Ask for a case summary, a response draft and any exceptions that need a person.

AI can help summarize the history, find relevant policy and draft a reply. The person still owns policy exceptions, escalation, tone and promises to the customer. A fluent reply can still promise something the business cannot do.

Specific feedback: "You promised a refund the policy does not allow. Cite the applicable policy, offer the permitted option and flag the exception for a person. Keep the empathetic opening. Do not send yet."

Review checklist: Are names, amounts and dates grounded in the source? Is the option permitted? Is the next step clear? Does an exception need escalation? Is the draft authorized to send through the available tool?

The reusable feedback prompt

“This misses [specific issue]. The result needs to [observable outcome]. Keep [what works]. Change [specific part]. Check it against [example, test or rule]. Tell me what is still uncertain.”

A good correction gives the model something it can act on. “Make it better” does not specify the gap. After the revision, inspect the changed result against the same standard.

Practice once, then save what works

Choose one real task. Write the expected result and your review criteria before asking AI. Keep the source material close. Compare the first draft, your feedback and the revision. Record the most useful correction in the review worksheet.

If the workflow proves useful, save the context, boundaries and quality checks as reusable instructions. Update them when the work or policy changes. The aim is a better finished result, not handing off responsibility for judgment.

Source

Anthropic: Labor market impacts of AI, March 5, 2026