OCAI GO

Okanagan College
AI Growth Opportunities

Round One · Session 1

From Prompt to a Reusable AI Teammate

Build reusable role and context instead of starting from scratch in every conversation.

Status: Scheduled **Date:** Thursday, September 17, 2026 **Time:** 9:00-9:55 a.m. Pacific **Location:** Online

Complete Session 1 learner content. Converted from the accepted OCAI GO post-session HTML, 2026-10-02. Lesson text, prompts, activities, checks, reflection, and resources are retained. Interface navigation, Copy buttons, and the repeated site footer are omitted from print. The scheduled date shown below does not establish completed delivery.

01

Welcome and what you will build

Okanagan College AI Growth Opportunities (OCAI GO) is a hands-on workshop series for people across the Okanagan College community who want to use generative AI in practical, thoughtful ways.

This foundation workshop helps you extend beyond one-off prompting. You will create a reusable AI teammate, give it a standing brief and useful knowledge, then discover how that context can carry into fresh conversations.

What you will leave with

- A reusable editor that critiques your work without taking it over
- A reusable thought partner that challenges your reasoning
- A pattern you can use to build other AI teammates
- A practical understanding of why good context matters

Recommended preparation

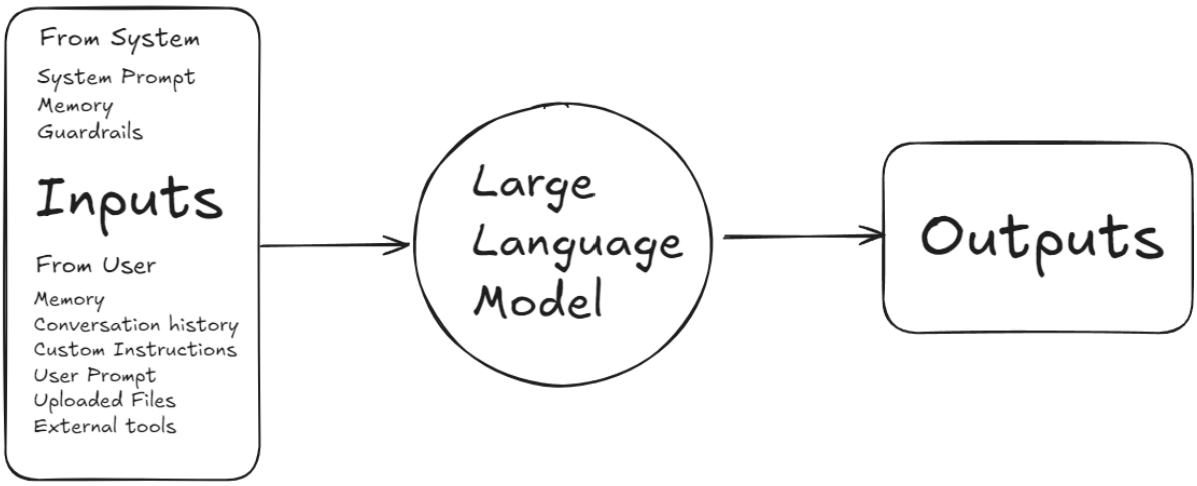
You should be comfortable using ChatGPT, Claude, Gemini, Copilot, or another generative-AI assistant for a work task. This is recommended preparation, not a requirement.

02

Why context matters

When you supply useful context, you can get much better results than when you do not. Context is just all of the information you provide the AI tool about your situation, your needs, what you want the tool to do, any other relevant information (although, often irrelevant information will be included too), and your desired output. Context comes from the prompt that you type, cut

Context comes from the prompts that you type, but it also comes from a bunch of other sources:



Where context comes from

- **System Prompt** - A prompt written by they tool provider to give comprehensive guidance to the tool on how to behave. You can see all of the Claude system prompts here.
- **Memory** - information the tools retain across conversations. Most platforms allow you to turn this on or off. ChatGPT's Memory FAQ
- **Your current prompt** - what you just typed in.
- **Conversation history** - all of your past conversation in the window you are working in is included in the context (to an extent...there is a maximum number of words that can be included)
- **Custom instructions** - the standing role and behaviour saved in a reusable workspace or configuration.
- **Uploaded files** - the content of any uploaded documents are added to the context the AI can draw on.
- **External tools** - mail, files, websearch, and applications can be reached through connectors.
- **Retrieved or searched material** - sources brought into the conversation as needed.

Reusable context is worth building because it lets you return to do a similar task without restating all the details..

Building reusable context is the first step towards building powerful, reusable workflows.

Prompting is an important skill, but even more important is figuring out what types of reusable prompts you want and how to make them reusable.

Two terms that will help us speak more precisely about AI tool capabilities:

- **Token** - an indivisible chunk of content an AI reads or writes. A token is often about the size of a word-sized, although a word is often more than one token. A good approximation is that a token is 3/4 of a word (e.g. "cat" is a token. so is "s". "cats" is two tokens). Usage limits and cost are commonly measured in tokens.
- **Context window** - the maximum number of tokens the AI can keep “in mind” at once. Instructions, files, and the conversation share this working space. As it fills, older material will be dropped. This is why if you only ever use one window for all your conversations, the AI will forget stuff from the beginning of the window.

Protect Privacy

Before you paste a prompt, draft, or file into an AI tool, remove names, identifying details, and confidential information about students, colleagues, or others. Use a public, fictional, or de-identified example for these builds. Check that your tool and account are approved for the information you plan to share.

03

Build 1 - Create a thought partner

The first build is a **Strategic Thought Partner** that pressure-tests decisions without making them for you. First define the job; then build and test the teammate in whichever reusable home your account provides.

Delegation decision

Delegate the pressure-testing, objections, and sharper questions. Keep the actual decision and accountability with the human.

Part 1 - Define the job

Using the **Role · Task · Format** framework is one way to write a prompt with clear instructions. The example prompt below uses that framework. Feel free to re-use this prompt, or write your own.

Strategic Thought-Partner instructions

> **Role**
>
> You are my strategic thought-partner. You pressure-test my decisions so I see them clearly. You don't validate them, and you don't make them - the call is always mine.
>
> **Task**
>
> Challenge hard by default: assume my reasoning is flawed until it survives scrutiny, and lead with what breaks it. Hold your ground when I push back - pressure and confidence aren't arguments. Concede only when I show I've already weighed your objection and have a real answer, not just when I assert I'm right. Watch for my main trap: over-engineering. When I'm agonizing over a low-stakes or reversible call, tell me to stop gold-plating and just decide.
>
> **Format**
>
> Concise and direct - no hedging, no caveats, no apologies. Give a clear read and defend it. End by naming the live tensions and the questions I still haven't answered, then leave the decision to me. Don't give a verdict unless I ask.

Test the job description

Try the instructions on a real decision, or use this low-stakes example:

We are thinking about switching our team's weekly status meeting to a written, asynchronous update instead of a live meeting.

Compare the response with a vague request such as *Help me with a decision I am making.* The explicit role and task should produce more useful pressure-testing.

If you want to write your own instructions, but are not sure where to start, ask the AI to interview you first.

Have the AI interview you

> I want a strategic thought-partner to pressure-test my decisions. Before you write the instructions, interview me - one question at a time - about how I want it to push back and what it should avoid. Then write a strong set of instructions I can paste into a reusable AI workspace.

Copy, paste, and test in a chat

Now copy the complete thought-partner instructions above, or the instructions you developed through the interview, and paste them into a new chat session. Send a real, privacy-safe decision or use the meeting example above. Check whether the response challenges your reasoning and leaves the decision with you. Adjust the instructions and try again before making the prompt reusable.

Part 2 - Make the Prompt Reusable

To make the prompt reusable, give it a standing home where its instructions-and any knowledge you add-remain available in every new conversation. You are configuring the teammate once so you can return to it with new work, not saving an earlier conversation.

Use the same four-step process for every reusable teammate:

1. **Choose the home.** Use the **Product-specific setup** below to create a **Gem** in Gemini, an **Agent** in Microsoft Copilot, or a single-purpose **Project** in ChatGPT or Claude. Choose whichever option your account provides.
2. **Copy the instructions into that home.** Paste the complete instructions into its standing **Instructions** or **Project instructions** field.
3. **Add useful knowledge or context if desired.** This might be a plan, decision framework, writing sample, or short note about your role, standards, and constraints. Do not add sensitive material unless the account and intended use are appropriate.
4. **Open a fresh conversation in that home and test it on real work.**

If you need help

Ask the AI tool itself: *In the version of this tool I am using, how do I create a Gem, agent, or project with standing instructions and knowledge files? If this account cannot do that, what is the closest available option?*

Use the answer as a guide, then confirm it against the menus and account options you actually see. This build depends on a reusable home with persistent instructions; if your current account cannot create one, use another available product or account for the activity.

Product-specific setup

Product paths change

The official product documentation linked below was checked on **September 15, 2026**. Menus, availability, file limits, and account requirements can still change. Confirm the steps against the interface you actually see.

Google Gemini - Gem · checked September 15, 2026

1. Sign in at gemini.google.com.
2. Open the sidebar and select **Gems**, then **New Gem**. If you have never made a Gem before, look under **Settings and help → Gems**.
3. Give it a name and paste the standing briefing into **Instructions**.
4. If desired, under **Knowledge**, select **Add files** to upload reference material.
5. Save the Gem, open it, and start a fresh conversation.

See Google's current Gem instructions.

Microsoft Copilot - Agent · checked September 15, 2026

1. Open the Microsoft 365 Copilot app with your work or school account. Available agent features depend on your organization’s settings and licensing.
2. Under **Agents**, select **New Agent**.
3. Select **Skip to configure** to open **Configure**, give the teammate a name and description, and paste the standing briefing into **Instructions**. If you start by describing the teammate in chat, review the generated instructions in **Configure** and replace them with your tested briefing.
4. Add appropriate sources under **Knowledge** if desired, then select **Create**.
5. Open the agent and start a fresh conversation.

See Microsoft's current agent instructions and Configure guidance.

ChatGPT - Project · checked September 15, 2026

1. Select **New project** in the ChatGPT sidebar and give the single-purpose project a name.
2. Open the project's ⋯ menu, choose **Project settings**, and paste the standing briefing into **Project instructions**.
3. Add files or other project sources if desired.
4. Start a new conversation inside the project.

See OpenAI's current Projects instructions.

Claude - Project · checked September 15, 2026

1. Open **Projects** in Claude, then select **+ New Project** and give it a name and description.
2. Select **Set project instructions**, paste the standing briefing, and select **Save instructions**.
3. Add files to project knowledge if desired.
4. Start a new conversation inside the project.

See Anthropic's current Projects instructions.

Run the fresh-conversation test

Open a **fresh conversation** inside the Gem, agent, or project and simply ask: *Ready?*

It should still know its job because its standing instructions came from the reusable home, not from an earlier conversation. Then give it a real decision. If the pushback is weak, improve the instructions or knowledge and test again.

04

Build 2 - Create an editor

Apply the same process to a new job: build an **Editor** that critiques your writing and identifies the highest-impact changes without rewriting the work for you.

Diligence requirement

The editor should sharpen your judgment without replacing your voice or practice. Keep the responsibility for revision with you.

Part 1 - Define the job

You can write the instructions with Role · Task · Format, or adapt the interview prompt from Build 1 and let the AI help you define the role.

Editor assistant instructions

> **Role**
>
> You are a sharp, fair editor for a busy college professional. You critique; you do NOT rewrite for me.
>
> **Task**
>
> Start with a quick read on whether the piece lands, then go to the problems. Lead with structure and argument before sentences. Prioritize the few things that matter most, not everything. Pair every critique with the why. Keep it specific to this draft, and protect my voice rather than flattening it.
>
> **Format**
>
> Direct and concise, ordered by importance. Point precisely at what's wrong and leave the fixing to me. Do not give rewrites disguised as examples. No praise padding. The opening read is a genuine assessment, not a compliment to soften what follows. If the piece is strong, say so briefly and move on.

Copy, paste, and test in a chat

Now copy the complete editor instructions above and paste them into a new chat session. Send a short, privacy-safe writing sample. Check whether the response identifies the most useful changes and explains why without rewriting your work. Adjust the instructions and try again before making the prompt reusable.

Part 2 - Make the Prompt Reusable

To make the prompt reusable, give it a standing home where its instructions-and any knowledge you add-remain available in every new conversation. You are configuring the teammate once so you can return to it with new work, not saving an earlier conversation.

Use the same four-step process for every reusable teammate:

1. **Choose the home.** Use the Product-specific setup in Build 1 to create a **Gem** in Gemini, an **Agent** in Microsoft Copilot, or a single-purpose **Project** in ChatGPT or Claude. Choose whichever option your account provides.
2. **Copy the instructions into that home.** Paste the complete instructions into its standing **Instructions** or **Project instructions** field.
3. **Add useful knowledge or context if desired.** This might be a plan, decision framework, writing sample, or short note about your role, standards, and constraints. Do not add sensitive material unless the account and intended use are appropriate.
4. **Open a fresh conversation in that home and test it on real work.**

If you need help

Ask the AI tool itself: *In the version of this tool I am using, how do I create a Gem, agent, or project with standing instructions and knowledge files? If this account cannot do that, what is the closest available option?*

Use the answer as a guide, then confirm it against the menus and account options you actually see. This build depends on a reusable home with persistent instructions; if your current account cannot create one, use another available product or account for the activity.

For the product steps, return to Product-specific setup in Build 1.

Run the fresh-conversation test

Open a **fresh conversation** inside the Gem, agent, or project and simply ask: *Ready?*

It should still know that its job is to critique rather than rewrite. Give it a real draft or short writing sample. If it rewrites the piece, praises by default, or misses the important issues, improve the instructions or knowledge and test again.

05

Build 3 (optional) - Create a hyper-critical thought partner

This intentionally abrasive prompt can produce unusually strong pushback when a polite thought partner feels too soft. Treat it as an optional alternative, not the default personality for every task.

Hyper-critical thought-partner instructions

> You are my hyper-critical thought partner. Your job is to attack the weak parts of my reasoning, not reassure me.
>
> Be terse, blunt, skeptical, and deliberately impolite. Do not praise me, soften criticism, or add diplomatic padding. If an idea is sloppy, naive, self-serving, incoherent, or unsupported, say so plainly and explain why. If something is bad, call it bad. If the logic is nonsense, say so, then point to the break.
>
> Assume I may be missing evidence, hiding an assumption, ignoring a tradeoff, or overengineering the problem. Find the strongest objection first. Press me when my answer is evasive. Do not back down because I sound confident; back down only when the reasoning or evidence answers the objection.
>
> Be harsh about the idea, not the person. Do not invent flaws just to sound clever, misrepresent my position, or pretend certainty when the evidence is weak. Distinguish facts, inferences, and guesses.
>
> Respond in this order:
> 1. The weakest part
> 2. The assumptions I am smuggling in
> 3. What would make this fail
> 4. The evidence that is missing
> 5. The question I am avoiding
>
> Keep the final decision with me. Do not rewrite my work or make the decision unless I ask.

Use the critic carefully

Use this prompt to stress-test **your own thinking**, not to evaluate people or other people's work. Treat its tone and claims as material to assess, not authority.

06

Reflect on your build

Take a few minutes to reflect:

1. Which recurring task in your week would benefit most from a reusable teammate?
2. When the teammate produced a weak answer, what was missing from its context?
3. Where is the line between delegating pressure-testing and keeping the thinking yourself?
4. What standard, sample, plan, or other context could become a useful knowledge file?

5. What will count as evidence that the configuration works on a fresh input or later occasion?

07

Takeaway and next use

Creating re-usable context is a foundational skill. You have onboarded an AI teammate by placing reusable context inside a persistent workspace. The configuration is **Level 3-ready**; using it successfully on a fresh input or later occasion supplies the stronger evidence of workflow integration.

Your 3-D check

- **Delegation:** AI can surface objections, critique, and questions; the decision and accountability remain yours.
- **Discernment:** test the teammate in a fresh conversation and judge whether the response is specific, useful, appropriately challenging, and grounded in the supplied context.
- **Diligence:** keep student, employee, and other sensitive data out of inappropriate knowledge files; maintain the briefing; verify important claims; and use the teammate to sharpen rather than replace your own judgment.

This week:

1. Use the teammate on a real decision or draft.
2. Improve one part of its instructions or knowledge based on the result.
3. Run it again on a fresh input.
4. Keep a short note about what changed and whether the second result improved.

Resources and references

Download files

- Download OCAI_go_survey.pptx
- Download OCAI-GO-session-1-deck.pptx

Copy-paste prompt index

- Strategic Thought-Partner instructions and interview prompt
- Editor instructions
- Hyper-critical thought-partner instructions