aruntastic

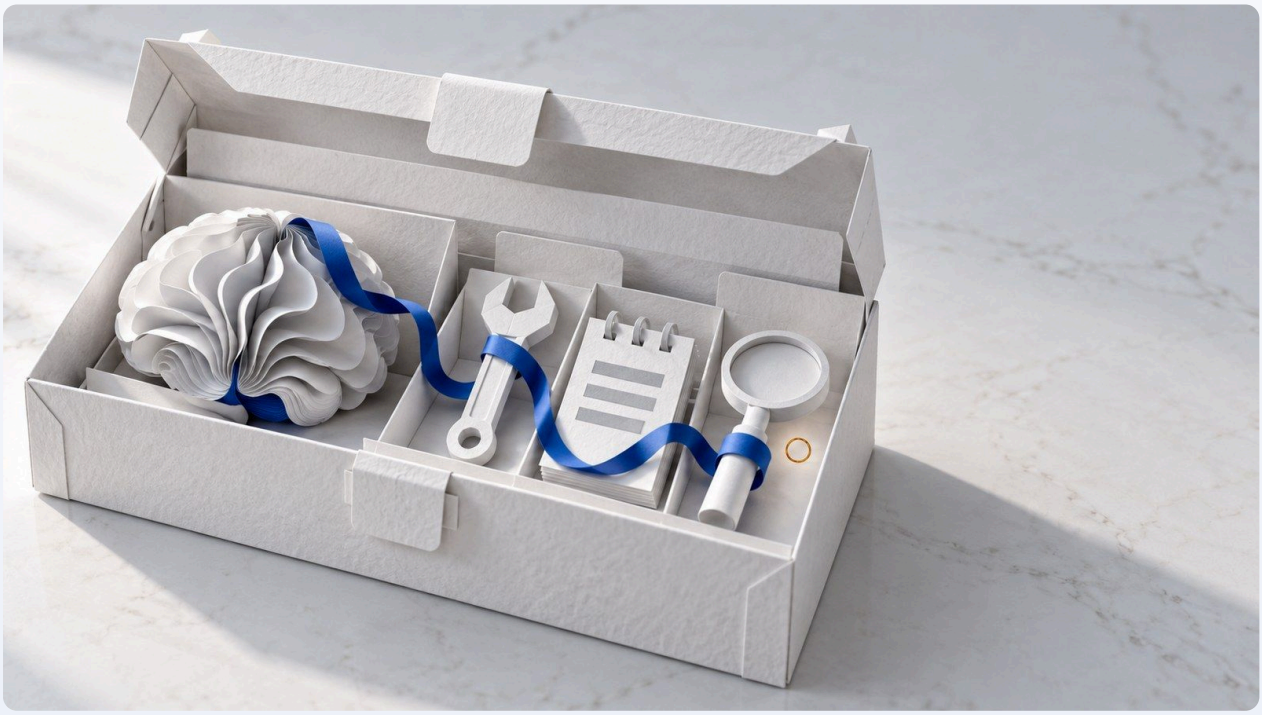
Study guide · Section 2

INTRODUCTION TO AGENTIC AI AND AI AGENTS

Understand How AI Agents Work (Models, Tools and Memory)

What each part of an agent does, and how to compare models on your own job

Keep this guide beside the Section 2 videos. Daniel and his invoices are teaching data, and the model comparison is optional practice.





Arun Nagarathanam

14 lectures · prompts, ideas and checks

BEFORE YOU START

How to use this guide

Keep it open on your phone or a second screen while the lecture plays on your main screen.

ONE

Keep it beside the video.

Watch the lecture, then try the steps yourself.

TWO

Use the prompts.

Every dark window holds a prompt from the lecture. Select the text and copy it.

THREE

Check yourself.

Each lecture ends with one question. Answer it before you move on.

ALL LINKS AND COPY BUTTONS

Every link for this section, and every prompt with a working Copy button, is on one page of our website.

Open the Section 2 page

IN THIS SECTION

14 lectures

2.0	What Is Inside an AI Agent?	Page 4
2.1	Why Can AI Explain a Job Without Doing It?	Page 6
2.2	How Does an AI Agent Take Action?	Page 9
2.2A	How to Get the Most From This Course	Page 12
2.3	What Is an API?	Page 15
2.4	How Does an AI Agent Work?	Page 18
2.5	Can You Spot an Agent's 4 Moves in Your Work?	Page 21
2.6	What Does an AI Agent Remember?	Page 24
2.7	What Is an LLM in an AI Agent?	Page 27
2.8	What Can AI Read and Create?	Page 29
2.9	How Do You Choose an AI Model?	Page 32
2.10	How Can You Compare AI Models?	Page 35
2.11	What Changes When You Switch AI Models?	Page 42
2.12	What Can You Now Explain About an Agent Run?	Page 45

2.0

What Is Inside an AI Agent?

Ask what lets an agent move from understanding a request to changing the files and tools a team uses.



IDEAS TO KEEP

01

In the webinar job, the agent read the planning-meeting transcript, prepared corrected invitation slides for Omar to review and saved a short update for the manager.

02

By the end of this section you can look at a run in Claude Code or Codex and say which part chose the next step, which part carried out a search and which part wrote the answer.

03

You will also see how the application keeps your requirements, and the information its searches found, available while the work continues.

04

Knowing which part does what tells you what is missing when AI answers your question but leaves the action undone.

FOLLOW ALONG

1. Think of one time an AI tool told you how to do something instead of doing it.
2. Keep that example in mind for the shirt-folding lecture that comes next.

FILES AND LINKS

Section 2 study guide (PDF)

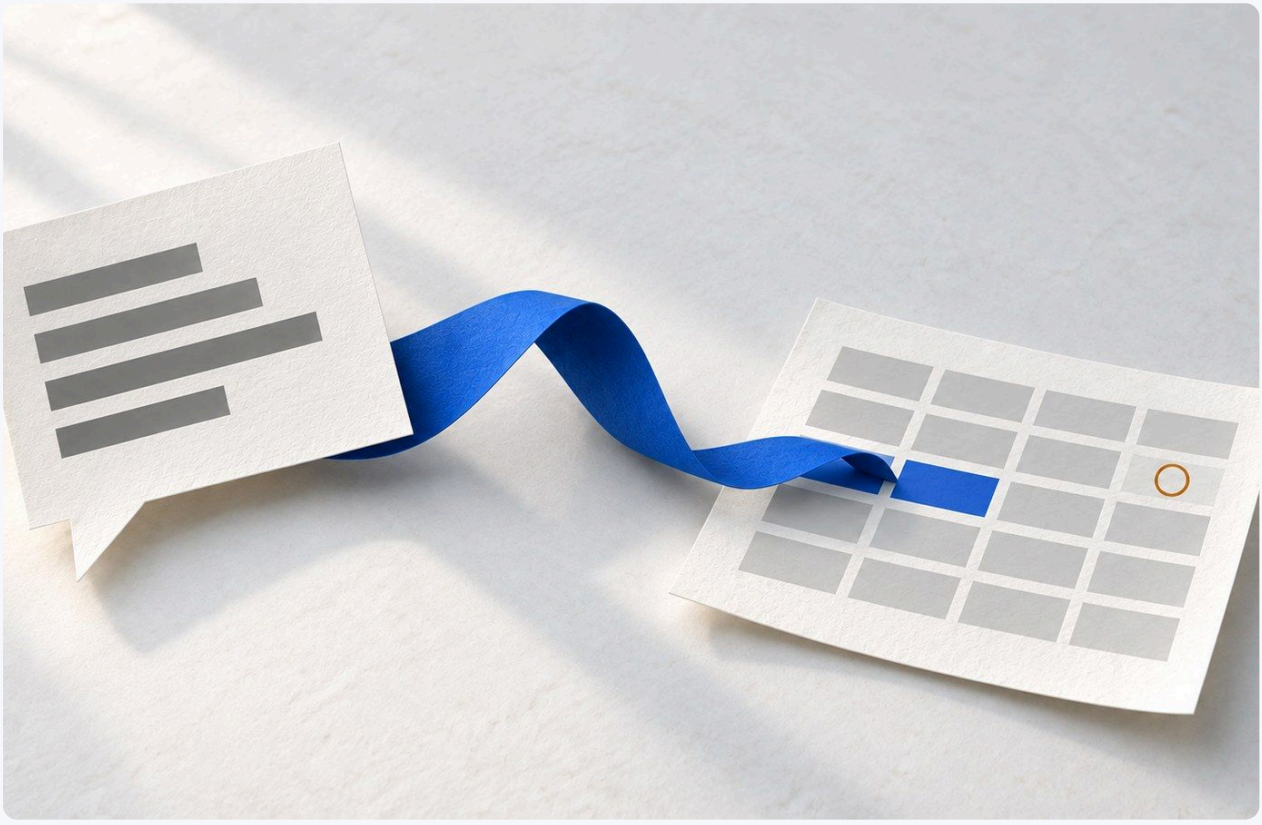
CHECK YOURSELF

In the webinar job, which three things did the agent change or produce, rather than only describe?

2.1

Why Can AI Explain a Job Without Doing It?

Separate the part that works out an answer from the part that can open a file and change it.



IDEAS TO KEEP

01

An AI can describe how to fold a shirt while the shirt stays on the table. Weave Robotics' Isaac 0 is a machine that actually folds it.

02

When you send a message to Claude, the part that interprets your request and works out a response is called the model. Sonnet and Opus are models inside Claude.

03

Claude, ChatGPT and Gemini are apps. The app gives you a way to work with a model, and it decides which files and actions the model can use.

04

To edit a spreadsheet, the formula has to be entered in the right cell and the file saved. That needs the app to open the file, and it needs your permission to change it.

05

Switching to a different model does not fix this. If the app cannot open and change the file, the new model may also hand you a formula to copy.

FOLLOW ALONG

1. Say exactly what you want done, including the change and the save.
2. Check whether the app has a way to open and change the file.
3. Check that you have allowed it to make the change.

Request · Ask for the action, not only the answer

PROMPT

I want you to put the formula in the spreadsheet and save the updated file.

Use in: Claude, Claude Code or another app that can open and save your file

FILES AND LINKS

Links and copy buttons on the website

CHECK YOURSELF

Claude gives you a formula when you wanted the spreadsheet changed. Why might choosing a different model not help?

2.2

How Does an AI Agent Take Action?

Look at the actions a Google Calendar connection offers, and see why an appointment written in a reply is not an appointment in your calendar.



IDEAS TO KEEP

01

Claude writes the supplier call for Friday at 3 p.m. in its reply, but the call is not in Google Calendar. To create it, Claude needs a connection to Google Calendar that lets it create an event.

02

Claude calls these connections connectors. They can let the AI use your email, appointments or team messages, depending on the connection and the permissions you choose.

03

Each action is a separate tool: get_event reads an event, create_event creates one and update_event changes an existing one.

04

Codex offers connections through its plugin catalog. In the recording, Codex said it could not create the appointment from that conversation, even though the plugin lists the action.

05

The connected account must also allow changes. A company can allow reading appointments while blocking changes.

FOLLOW ALONG

1. Open the connection page and read the list of actions.
2. Sort them into actions that read information and actions that change something.
3. Before asking the agent to book the call, check that the conversation can use the event-creation tool.
4. Afterwards, open Google Calendar and check the call appears at the agreed date and time.

FILES AND LINKS

Links and copy buttons on the website

CHECK YOURSELF

Which Google Calendar tool does the model need to choose for a new supplier call, and which one to move it to another time?

2.2A

How to Get the Most From This Course

See which resources come with each section and how to use them.



IDEAS TO KEEP

01

Follow the course in order the first time, because later explanations use ideas the earlier lectures introduce. Lecture titles help you find one explanation again.

02

Each section has a downloadable summary with the main ideas and key words in one place.

03

Practical lessons come with prepared files. In this section, the model comparison project gives you sample invoices and business files, so you spend your time on how the models handle the job.

04

Where an assignment has a sample answer, use it to judge your own result.

05

The quizzes are a short way to check your understanding. Answer before you read the explanation; a surprising answer marks an idea worth another look.

FOLLOW ALONG

1. Turn on captions, pause or change the playback speed whenever you need to.
2. Open this section's summary when you want a reminder of the key words.

FILES AND LINKS

Links and copy buttons on the website

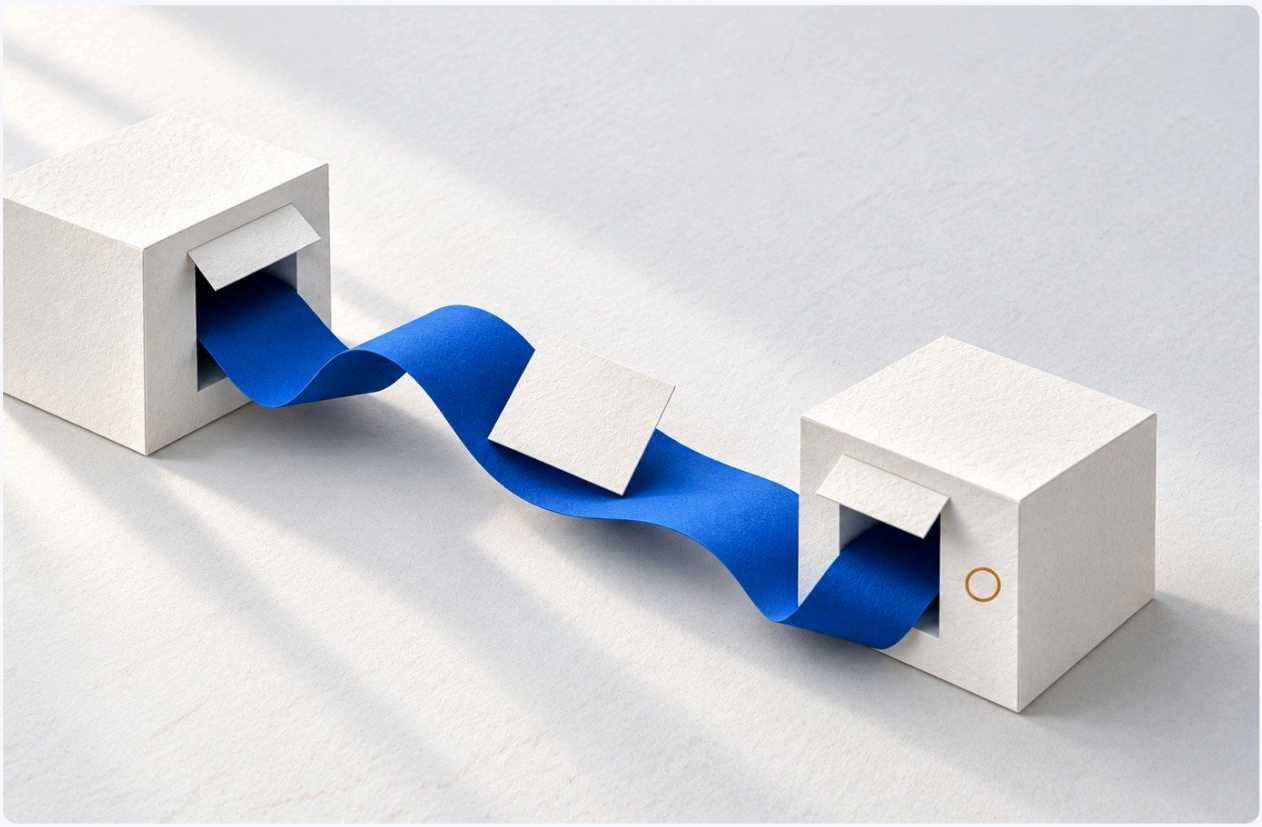
CHECK YOURSELF

When a quiz answer surprises you, what should you do with that idea?

2.3

What Is an API?

Follow the supplier call from ordinary words to an event that Google Calendar creates and sends back.



IDEAS TO KEEP

01

The calendar tool needs the details: the date, a 3 p.m. start and 3:30 p.m. end, a title, and which calendar to use if you have more than one.

02

Another program can send those details to Google Calendar as a request. If the details and permissions are valid, Google Calendar creates the event and sends back the saved event's details.

03

Google publishes instructions for these requests: which actions exist, what each one needs and what comes back. That agreed way for programs to ask for information or actions is the API, short for application programming interface.

04

Asking for Friday's appointments returns events already there. Asking to create the supplier call adds a new event.

05

You still grant access, for example by signing in and approving it. After that the connection sends the authorized requests for you.

FOLLOW ALONG

1. Describe the appointment in ordinary words.
2. Check the returned event: is the date and time right, and is it in the right calendar?
3. If Google Calendar rejected the request, expect an error, not a claim that the call was booked.

FILES AND LINKS

Links and copy buttons on the website

CHECK YOURSELF

? What should the software do if Google Calendar rejects the request to create the supplier call?

2.4

How Does an AI Agent Work?

Watch Claude Code turn a project-management question into searches and a table, then check one row against the provider's page.



IDEAS TO KEEP

01

The request asked for 3 project-management applications under 20 US dollars per user per month, which ones let a client join as a guest without a paid account, and a source link under every number.

02

Claude Code sent 2 broad searches first, then specific searches for ClickUp Unlimited, Asana Starter and Trello Standard. The returned results are Look.

03

Choosing which applications and plans to investigate is Think. Sending 5 WebSearch queries and producing the table is Do.

04

Opening a source link and comparing the plan name and price with the provider's page is your Check. In the recording, Asana Starter showed US$10.99, matching the table; billing and guest conditions still needed separate checks.

05

The links make the check possible. They do not show that the agent performed it. When a result leads the agent to choose another action and keep going, that repetition is often called an agent loop.

FOLLOW ALONG

1. Open one source link from the table.
2. Compare the plan name and price with the provider's page.
3. Check whether the price needs a full year's payment and whether a client can join as a free guest.
4. If a price does not match, ask the agent to look again; that starts another Look.

GOOD TO KNOW

Prices on these pages change. The figures in the lecture are what the recording showed on the day it was made.

FILES AND LINKS

Links and copy buttons on the website

CHECK YOURSELF

Claude Code searched, read the results and searched again. Which of those actions is Look and which is Do?

2.5

Can You Spot an Agent's 4 Moves in Your Work?

Follow Daniel through 2 unpaid-invoice reminders and name each of the 4 moves.



IDEAS TO KEEP

01

Daniel reads the payment records and each client's messages, checking for any payment or question about the bill. That is Look.

02

The firm's rules call for a polite reminder to the client who usually pays late and is 9 days overdue, and a firmer request for a payment update to the new client who is 30 days late with no payment or dispute. Choosing the response is Think.

03

Writing the 2 messages is Do. Comparing the amounts with the payment records and the wording with the reminder rules, before sending, is Check.

04

An AI could prepare the drafts if it had the invoice records, client messages and company rules. Daniel still reads them before anyone sends them, and checks for payments again later.

05

The handout's examples, such as a customer return and a supplier order, each name the original record you use to check the result.

FOLLOW ALONG

1. Pick one example from the handout.
2. Name the information used, the choice made, the action taken and the record that checks it.
3. Optionally, do the same for a recurring job of your own where the next action depends on what you find.

FILES AND LINKS

2.5-look-think-do-check-examples.md

Links and copy buttons on the website

CHECK YOURSELF

For Daniel's reminders, what exactly does he compare before sending?

2.6

What Does an AI Agent Remember?

See what Claude Code supplies to the model for Daniel's client reply, what can drop out when space runs low and where a lasting rule belongs.



IDEAS TO KEEP

01

A client asks for the next report as a spreadsheet instead of a PDF. The firm's rule says: keep replies courteous and do not promise a change until the operations team approves it.

02

Typing only 'Write a reply to this client' tells the model neither. Claude Code reads the client's message and the company instruction and supplies them with Daniel's request. That supplied information is the context.

03

The model can work with only so much at once; that capacity is its context window. Text is measured in tokens, a whole word or part of one. For these Claude models, the window holds what is supplied and the answer it writes.

04

When space runs low, Claude Code can shorten earlier conversation into a summary. This is compaction, and a summary can drop a detail such as the approval rule. The original files stay where they are, so Claude Code can read them again.

05

A rule for future replies can go in CLAUDE.md, the project's instruction file that Claude Code reads when it starts work. Keeping information for later use is commonly called memory.

FOLLOW ALONG

1. In Claude Code, open the context usage display and see which categories use space.
2. If a reply ignores a rule, ask Claude Code to read the client's message and the current instruction again.
3. Open CLAUDE.md to read the saved rules, and update it when your company's rules change.

Request · A saved instruction for CLAUDE.md

PROMPT

```
Keep client replies courteous and do not promise a change before approval.
```

Use in: The project's CLAUDE.md file

GOOD TO KNOW

The corrected reply from the lecture: "Thank you for requesting the next report as a spreadsheet. I will check with our operations team and confirm whether we can provide that format."

FILES AND LINKS

Links and copy buttons on the website

CHECK YOURSELF

The saved instruction is in CLAUDE.md. Why does Claude Code still need the client's latest message for this reply?

LECTURE 2.7

2.7

What Is an LLM in an AI Agent?

Understand what a large language model is and why a fluent answer can still be wrong.



IDEAS TO KEEP

01

Sonnet and Opus learned patterns from large amounts of text during training. Models that work with language at this scale are called large language models, or LLMs.

02

An LLM is one type of model. The word model also covers systems built for other jobs, such as predicting demand for a product. Some language models can also work with images or audio.

03

An LLM writes its response a token at a time, using what it learned and the current context. After 'Thank you for your', it might write 'message' or 'payment'.

04

A model can invent a price or feature and present it as true. That kind of error is often called a hallucination.

05

A convincing sentence does not prove a claim. For a software price, compare the answer with the provider's pricing page.

FOLLOW ALONG

1. When an answer gives a fact you will act on, find the page it should come from.
2. Compare the claim with that page before you use it.

CHECK YOURSELF

An AI tells you a software plan costs 12 dollars a month. What do you check, and where?

2.8

What Can AI Read and Create?

Choose a model by asking what information the job needs and what result you want back.



IDEAS TO KEEP

01

The information you give the AI is the input. The result you ask for is the output.

02

For the total on a photographed receipt, the input must include the picture. The app can pass the image to a model that reads images, or convert its words to text first; either way the amount must survive. A PDF table also needs its layout kept so each amount stays with its item.

03

OpenRouter lists models from different providers and the input types they support, with filters for text, image, audio and video.

04

If a meeting summary depends only on the words said, turn the speech into text first; that is transcription, done by a speech-to-text model. To judge someone's tone of voice, the model needs the audio itself.

05

Creating pictures, speech, music or video is a separate output question. Hugging Face lists models by task and calls speech-to-text 'automatic speech recognition'. A listing does not mean the model runs in your app.

FOLLOW ALONG

1. Write down the input the job needs and the output you want.
2. Find models that support that input and output.
3. Check that your application accepts the input and passes it to the model.

FILES AND LINKS

Links and copy buttons on the website

CHECK YOURSELF

You want to know whether a client sounded annoyed on a recorded call. Why is a transcript not enough?

2.9

How Do You Choose an AI Model?

Compare models on how well they handle the work, what using them costs and how long you wait.



IDEAS TO KEEP

01

Model menus change with updates. In the recording, Codex showed model names without descriptions, and Claude Code described only one model on hover.

02

A benchmark is a defined test used to compare systems. Artificial Analysis runs tests on models, and its Intelligence Index combines several into one score. Look at which benchmarks are included before reading the score.

03

API prices are quoted per 1 million tokens, often separately for input tokens (your request and material) and output tokens (what the model writes). A monthly subscription has its own price and usage limits; check your plan's usage information for those.

04

A lower price per token does not always mean a cheaper job. Long answers, repeated requests and extra reasoning tokens add up, and so does the time you spend correcting the answer.

05

Artificial Analysis reports time to the first answer token and output speed in tokens per second. Reasoning tokens can arrive before the answer starts, so measure the full response time, and count waits for websites and tools too.

FOLLOW ALONG

1. Use published tests to choose 2 or 3 models worth trying.
2. Check whether a quoted price is for input tokens, output tokens or both.
3. Check your own plan's usage page for limits and reset times.
4. Judge the whole job: answer quality, cost or usage, and total time.

FILES AND LINKS

Links and copy buttons on the website

CHECK YOURSELF

Why does a lower price per token not guarantee a cheaper finished job?

2.10

How Can You Compare AI Models?

Compare 2 models fairly: same folder, same request, same checks, then compare waiting, corrections and usage.



IDEAS TO KEEP

01

The useful model is the one that handles your job well. Give both models the same job, source material and checks, in 2 fresh Claude Code conversations with comparable settings.

02

Work out the correct result first. Harbor & Pine owes 1,800 dollars of 2,400 and gets a courteous reminder that acknowledges the 600 paid. Northline is 30 days overdue with no payment or dispute, so it gets a firmer request for a payment update.

03

Willow Ledger has paid in full and needs no reminder. Seaglass is disputed, so it goes to a person and gets no reminder.

04

Artificial Analysis and OpenRouter help you choose candidates. Arena shows 2 answers with the model names hidden and turns votes into rankings; a preferred answer can still hold a factual error.

05

In the recording, Sonnet and Opus both met every check. The usage breakdown showed Sonnet's cost a little lower. Record each run's time and what you would correct, and try a few more examples before relying on a choice.

FOLLOW ALONG

1. Download model-comparison-project.zip and unzip it.
2. Follow the practice sheet after this lecture for the full steps.
3. Save each complete response in the model comparison results folder.
4. Check both answers against the answer key that follows the practice sheet.

Request · First model

PROMPT

Look at all the files in this folder. Work out what we should do with each invoice, and prepare any reminders that are needed.

Use in: Claude Code, in the model-comparison-project folder

Request · Second model

PROMPT

Look at all the files in this folder (exclude 'model comparison results' folder). Work out what we should do with each invoice, and prepare any reminders that are needed.

Use in: Claude Code, a fresh conversation in the same folder

FILES AND LINKS

model-comparison-project.zip

Compare 2 models on the same invoice review

Invoice comparison answer key

Links and copy buttons on the website

CHECK YOURSELF

Both models met every check. What else can still tell you which one to prefer for this job?

Compare 2 models on the same invoice review

Optional practice for repeating the worked comparison. Watching it is enough to follow the course. Keep this sheet and the answer key outside the project folder, so neither model reads the expected answer.

SET UP

1. Download model-comparison-project.zip and unzip it. The model-comparison-project folder holds the invoices (invoices.csv), client messages, the payment-reminder policy and the saved instructions (CLAUDE.md).
2. Create a folder called model comparison results inside the model-comparison-project folder.
3. In the Claude desktop app, open the Code tab, choose Local and select the model-comparison-project folder.

RUN BOTH MODELS

1. Start a fresh conversation with one suitable model available to you. Paste the first request from Lecture 2.10 and run it.
2. Save the complete response in model comparison results.
3. Start another fresh conversation with a second suitable model, in the same folder. Keep other settings the same where Claude Code offers comparable settings.
4. Paste the second request. It adds (exclude 'model comparison results' folder), because the first answer is now saved inside the project.
5. Ask the second model to save its complete response, unchanged, in model comparison results.

RECORD

1. For each model: its name, how long the whole run took and what you would correct or check before using its drafts.
2. Compare usage only where the app reports it for each conversation, as the usage breakdown did in the recording. Account-wide usage cannot be split between the two models.

Invoice comparison answer key

Use this after both runs are finished. The as-of date in the project is 10 September 2026.

WHAT A CORRECT ANSWER DOES

1. Harbor & Pine Studio (INV-1042): 2,400 dollars minus 600 paid leaves 1,800. No dispute, the client usually pays late, and the partial payment is recorded, so a courteous reminder for 1,800 that acknowledges the 600 (policy P1).
2. Northline Workshop (INV-1048): 1,750 dollars unpaid, 30 days overdue, no payment or dispute recorded, so a firmer request for a payment update (policy P2).
3. Willow Ledger Co. (INV-1045): 1,200 dollars paid in full before the due date, so no reminder (policy P3).
4. Seaglass Market (INV-1051): 980 dollars disputed with a service-manager review open, so no reminder; it stays with a person (policy P4).

CHECKS FOR BOTH ANSWERS

1. Every decision and draft uses only the supplied records and names the evidence and policy rule behind it.
2. Harbor & Pine's draft uses the 1,800 remaining balance and acknowledges the partial payment.
3. Northline's draft asks for a payment update without inventing a deadline or consequence.
4. Willow Ledger and Seaglass get a clear No draft note.
5. No answer claims a message was sent or a record was changed.
6. Missing information is named without inventing contact names, payment dates, links or promises.

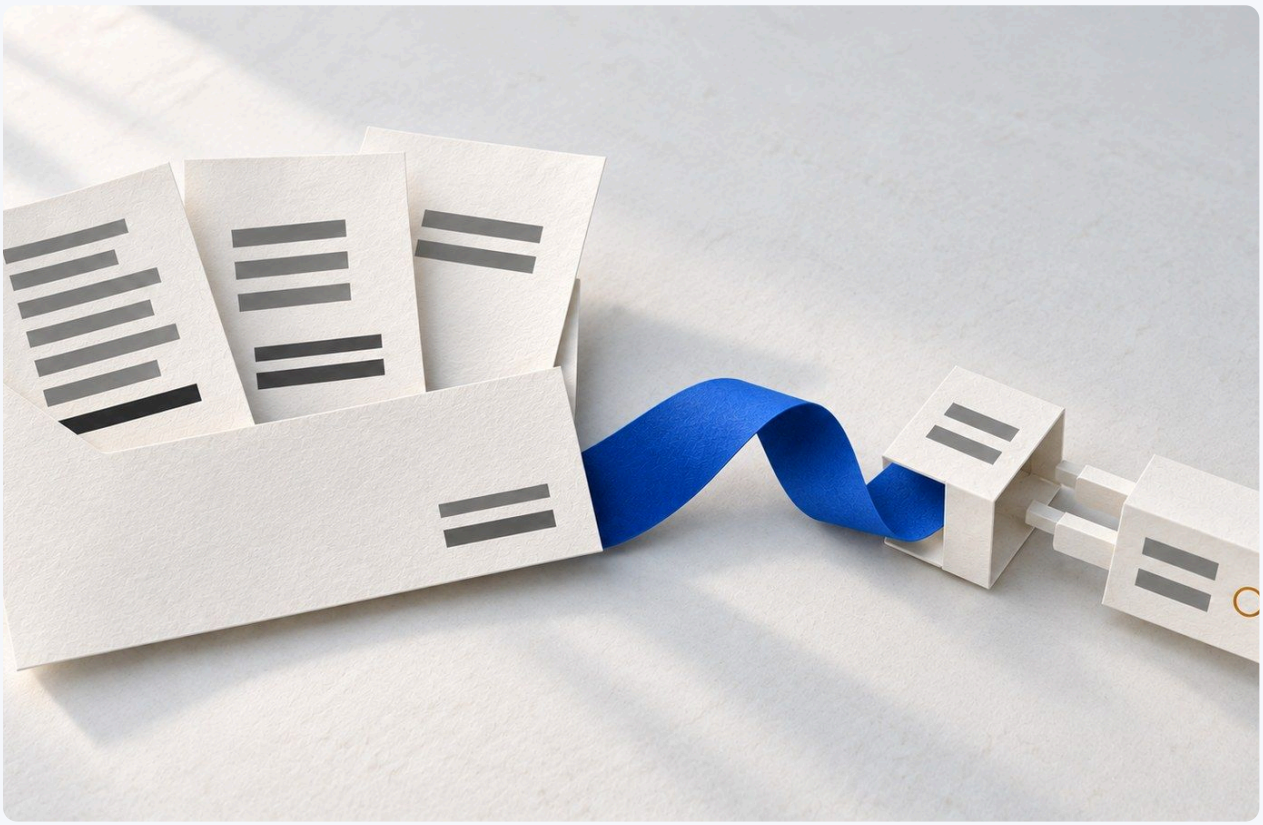
CHOOSING BETWEEN THEM

1. Neither model has to fail. If both pass every check, compare waiting, corrections and the review work each needs, and usage only where it is reported per conversation.

2.11

What Changes When You Switch AI Models?

Check what stays available when you change models in one app, and what you must set up again when you move to another app.



IDEAS TO KEEP

01

The invoice files and reminder policy stay in the project folder. Claude Code reads them and gives them to whichever model you choose.

02

Inside the same app, saved preferences and account connections such as Google Calendar may still be available, because the app manages them. The model uses what the app supplies.

03

A new model can still produce a different result from the same files and instructions, so test it on the same job.

04

Moving from Claude Code to Codex needs a separate check. You may need to connect the calendar in Codex and give it the reminder-policy file.

05

If you buy an AI product for your company, ask the seller what happens when its model is replaced: do your saved rules still apply, do connections still work, and who tests the tasks you rely on?

FOLLOW ALONG

1. After switching models, run the same job and check the balances and draft messages.
2. After moving apps, reconnect what you need and supply the instruction files again.
3. Test the action itself, such as creating an appointment, and check the saved event's date and time.

CHECK YOURSELF

?

You move your invoice work from Claude Code to Codex. Name two things you should not assume Codex already has.

2.12

What Can You Now Explain About an Agent Run?

Put the section's parts together: model, tools, API, context and memory, and checks that stay the same when the model changes.



IDEAS TO KEEP

01

In the project-management comparison, the model chose which plans to investigate, search tools carried out the requests and the returned information helped the model build the table. Opening a provider's page compared a claim with its source.

02

An API gives software a way to request an action or information. A listed tool still has to be available in the task, with the access the action needs.

03

Daniel's client reply showed why current information and saved instructions both matter. Claude Code supplies them as context within the model's context window, and an omitted rule can be supplied again from its file.

04

The invoice checks stay the same when the model changes: Harbor & Pine has 1,800 dollars outstanding, the paid invoice needs no reminder and the disputed invoice needs a person.

05

To judge an AI product, you also need to know how the product uses these parts to choose its actions. Section 3 starts with ChatGPT.

FOLLOW ALONG

1. Open the Section 2 summary and read the key words once more.
2. Explain one agent run you watched in this section using model, tool, context and check.

FILES AND LINKS

summary.md

Section 2 study guide (PDF)

Links and copy buttons on the website

CHECK YOURSELF

? Which three invoice facts let you judge either model without needing a particular winner?

REFERENCE

Words used in this section

Model

The part that interprets your request and works out a response. Sonnet and Opus are models inside Claude.

App (application)

What you open and send messages to, such as Claude, ChatGPT, Gemini, Claude Code or Codex. It decides which files and actions the model can use.

Tool

One action the model can choose, such as get_event, create_event or update_event.

Connector

Claude's name for a connection to a service such as Gmail, Google Calendar or Slack. Codex offers connections through its plugin catalog.

Permission

What the connected account allows. A company can allow reading appointments while blocking changes.

API

Application programming interface: the agreed way programs request information or actions from a service and receive a result.

Look, Think, Do, Check

The information used, the choice made, the action taken, and the comparison of the result with its source.

Agent loop

An agent using a result to choose another action and continuing the work.

Context

The information supplied to the model for the reply it is writing now.

Token

A small piece of text, a whole word or part of one. Text, usage and API prices are measured in tokens.

Context window

How much the model can work with at once. For the Claude models in this section it includes what is supplied and the answer generated.

Compaction

Shortening earlier conversation into a summary to make room. A detail can drop out; the original files stay.

Memory

Information kept for later use, such as a rule saved in CLAUDE.md, the project instruction file Claude Code reads when it starts.

LLM

Large language model: a model that learned patterns from large amounts of text and writes a token at a time.

Hallucination

A fact a model makes up and presents as true.

Input and output

The information you give the AI, and the result you ask it to produce.

Transcription

Turning speech into written text, done by a speech-to-text model. Hugging Face calls it automatic speech recognition.

Benchmark

A defined test used to compare systems.

Input and output tokens

Your request and material, and what the model generates. API prices for each can differ and are quoted per 1 million tokens.

Reasoning tokens

Tokens some models produce while working through a problem before the answer begins.

Time to first answer token

The wait from sending a request until the first token of the actual answer arrives.

Tell agentic AI, chatbots and workflows apart

Next comes Section 3, where you take a familiar product, ChatGPT, and work out when it is answering from what it already learned and when it is using tools to do the job.

All course resources | Section 2 on the website | aruntastic.com





aruntastic

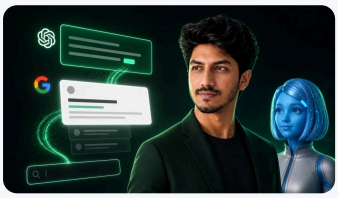
Written by Arun Nagarathanam
Part of Introduction to Agentic AI and AI Agents



KEEP BUILDING

More from Aruntastic

Practical courses for people who want to build useful things with AI.



AI SEO Course 2026: From Zero to Page-One Results

Research search demand, structure useful pages, and work through technical and off-page SEO.

Open the course



AEO, GEO, AI SEO: Get Cited by ChatGPT, Google, Perplexity

Study how AI search finds information, and make your pages clear enough to be cited.

Open the course



Claude Code Without Coding: AI That Does Real Work for You

Start with the installation, give Claude Code a real project, and learn how to check what it changes as you build.

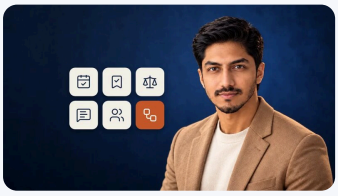
Open the course



Microsoft PowerPoint Presentation 2026 for Professionals

Learn visual hierarchy, slide composition, typography and animation through practical work.

Open the course



AI Agents for Leaders Who Have No Time to Learn AI

If you can delegate to a person, you can delegate to an AI agent. You keep the judgment. For managers and founders.

Open the course

Go deeper



TOOLKIT

SEO Accelerator Toolkit

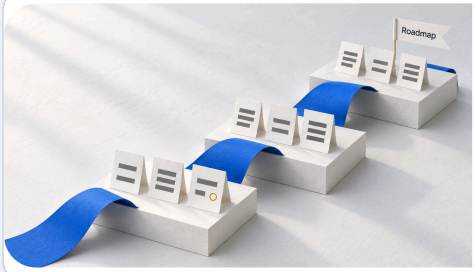
150+ copy-paste AI prompts for audience research, keywords, content, technical SEO and link building.



TOOLKIT

GEO Accelerator Toolkit

100 detailed AI prompts for GEO audits, entity building, content structure and client reporting.



PROGRAMME

GEO Authority Accelerator

A complete GEO programme on our learner platform: nine modules and a 30, 60, 90 day roadmap. The GEO Accelerator Toolkit is included.

Learn more

CLAUDE CERTIFICATION PRACTICE

Preparing for a Claude certification?

Practice exams written by Arun: 600 questions across 4 mock exams and 2 drill banks. Unofficial preparation, not affiliated with Anthropic.



CCAR-F
Architect Foundations



CCAO-F
Associate Foundations



CCDV-F
Developer Foundations



CCAR-P
Architect Professional