

/ activate intelligence /

Claude at Excellera

Session 1

orientation, model selection, and where claude fits

What this course is.

four sessions, sixty minutes each. demo-led. no homework.

01

TODAY

orientation, model selection,
where claude fits.

02

NEXT WEEK

prompting, artifacts, memory,
and styles.

03

WEEK 3

projects: build, share, and
iterate.

04

WEEK 4

skills, deep research, chrome,
MCP, office agents.

Think of Claude as a colleague.

a specific kind of colleague.

- 01** **sharp and well-read.**
reads quickly, writes well, knows a lot.
- 02** **each conversation starts fresh.**
unless you give it a project, or turn on memory.
- 03** **confident, including when wrong.**
treat the first answer as a draft, not a verdict.

You may already use other tools.

chatgpt, gemini, copilot. they are broadly similar, and the skills transfer.

excellera has standardised on claude, so that is what we are here for.

And there is other AI work happening here too.

excellera is building specialised tools for specialised work.

knowledge base

the firm's institutional memory, searchable.

reputation measurement

reading client coverage, and what is shaping it.

policy monitor

real-time legislative tracking, evolving politor.

intelligence hub

the backbone tying them together.

claude is the general-purpose colleague that sits alongside these.

Today.

four moves, in this order.

- | | | |
|----|------------------------------|--|
| 01 | pick a model | opus, sonnet, haiku. when to use each. |
| 02 | limits and connectors | what is enabled, what to avoid. |
| 03 | where claude fits | your work, your firm, your week. |
| 04 | using it confidently | a few notes for real client work. |

01

Pick a model.

opus, sonnet, haiku. the model picker, and how to choose.

Three things to know.

before we pick a model, three quick definitions.

A MODEL

the AI itself – the thing that reads what you write and produces what you read. Anthropic offers a few, each with different strengths.

A TOKEN

roughly a word, sometimes a syllable. Everything you send to Claude and everything it sends back is counted in tokens. Your usage allowance is measured in them. A short message: dozens. A long document: thousands.

WHY PICKING MATTERS

different models, different sizes, different costs. The most capable models burn through your allowance several times faster than the everyday ones. Picking deliberately means getting the work done without running out.

There is more than one Claude.

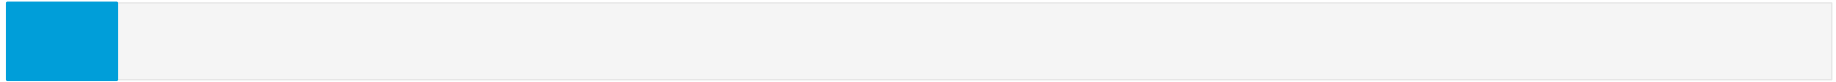
the picker sits at the top of any chat.

CHOOSE A MODEL	
Claude Opus 4.7	most capable. hardest tasks, deepest reasoning.
Claude Sonnet 4.6	everyday default. fast and capable.
Claude Opus 4.6	previous generation opus.
Claude Haiku 4.5	fastest. for short, simple tasks.

What each one costs to run.

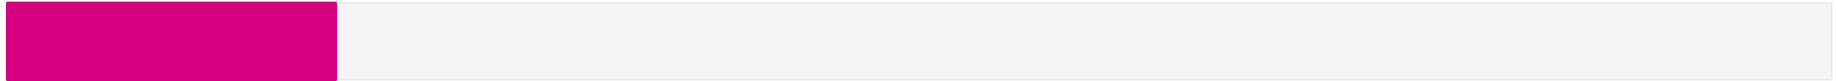
per task, in relative terms. pick opus when the quality is worth the spend.

haiku



roughly a third of sonnet.

sonnet



your everyday default.

opus



around five times sonnet.

pick opus too often and you will hit your usage limit sooner. sonnet is the default for a reason.

How to pick.

a simple rule.

REACH FOR

Opus

when quality matters more than speed

client deliverables, nuanced editing,
complex briefings, multi-document
research. slower, but worth it.

DEFAULT

Sonnet

for everyday work

drafting, summarising, iterating. fast
enough for chat, capable enough for almost
everything you do day to day.

REACH FOR

Haiku

when speed matters more than nuance

short rewrites, quick classifications,
bulk tidying. fast and cheap.

L I V E D E M O

Same prompt, two models.

we run the same brief on sonnet, then on opus – a press release lede and an email to the board, from the same paragraph. then a quick haiku demo for speed.

02

Limits and connectors.

what counts, what to avoid, and what is wired up.

Limits, in plain language.

long conversations and big files cost more than short ones. three habits keep things smooth.

01 start a new chat when you change topic.

old context follows you through the conversation. a fresh chat is faster, and often gets you a better answer.

02 put long documents in a project, not a chat.

we cover this in session 3. a 200-page pdf re-uploaded into every chat is wasteful.

03 a few well-chosen files beats a dumping ground.

especially in projects. the smaller and tighter the knowledge base, the better claude does with it.

if you hit a limit, it is a pause, not a block. limits reset on a schedule. the admin can grant more if needed.

Connectors.

a bridge between claude and a service you already use.

/ **with your permission.**

each connector asks before it acts. nothing happens silently in the background.

/ **only when asked.**

claude reads from a connector when your prompt calls for it, not by default.

/ **what is enabled depends on admin.**

some connectors are off until excellera switches them on. skills sit behind the same toggle.

L I V E D E M O

Pulling a real document.

open a connector. ask claude to read a recognisable document and summarise the key points.

03

Where Claude fits.

your work, your firm, your week.

Where Claude helps.

the work several of you flagged in pre-work – and more.

DRAFTING

briefings, press releases, client memos, internal updates. fast first drafts you edit on top of, rather than typing from a blank page.

READING

long parliamentary texts, regulatory consultations, sector reports, multi-document research. ask questions of a document rather than reading it cover to cover.

SYNTHESIS

press review across many articles, multi-source compilation from word, slides, and spreadsheets. groups by theme, flags what is notable by absence.

TRANSLATION

between italian and english, between formal and informal, between technical and plain. handles both directions, both registers.

CRITIQUE

a sceptical read of a draft before it goes out. catches unevidenced claims, structural weaknesses, the sentence that will not survive a senior eye.

And where it struggles.

the honest limits. worth knowing before you trust it on real work.

DOESN'T KNOW THE NEWS

claude's training has a cut-off date. anything that happened recently has to be in your prompt, or it does not know.

CAN BE CONFIDENTLY WRONG

the first answer is a draft, not a verdict. check the facts that matter. confidence is not accuracy.

DOESN'T HAVE ACCESS

claude sees what you give it. nothing else. no sharepoint, no email, no client files unless you paste them in or connect them.

ISN'T A DATABASE

long conversations get long. claude can lose track. between chats it remembers nothing unless memory is on.

04

Using it confidently.

a few notes for real client work.

Three notes for real client work.

most of what you draft for clients belongs in claude. these are the practical edges to watch.

Training and residency.

conversations aren't used to train anthropic's models by default. but workspace data is stored in the us – there is no eu-only option on claude team. if a client contract requires eu data residency, flag with it; that work needs a different setup.

Embargoes and handoff.

the chat itself is fine – claude team conversations don't leak. the moment you copy the draft out, the embargo travels with the file. be deliberate about where it goes next.

Incognito chats.

available for one-off sensitive work. nothing saves to your history, nothing feeds memory. team admins retain visibility for compliance – it's incognito from your view, not from the firm's. reach for it when you'd rather a conversation didn't stick around.

Questions.

drop them in the chat.

What happens next.

/ **next week**

mollie takes you through prompting, artifacts, memory, and styles.

/ **questions between sessions**

drop them in the shared channel. we will pick them up.

Thank you.