


SUNRISE COMMUNITIES · STAFF TRAINING

AI Crash Course

General · 60 minutes


“Less wrong, not expert.”

The goal of the next hour.



Agenda

10 SECTIONS · 60 MINUTES



- 1 Opening + framing
- 2 What an LLM actually is
- 3 The Sunrise System Prompt — live walkthrough
- 4 Three failure modes — with live demos
- 5 Cross-reference technique (Sam's three-LLM trick)
- 6 Token limits — what they are, when you'll hit them
- 7 Claude web vs Claude Desktop
- 8 Practical do / don't
- 9 Department follow-ups — what's coming
- 10 Q&A

The framing line



ANCHOR LINE

AI is here to **expand** what you can do,
not
replace what you already do.

Read it out loud — slowly. This is the sentence the rest of the hour rests on.

What an LLM actually is

5 MINUTES · NO TECHNICAL DEPTH



1

Prediction engine

Reads your text and predicts the next word, then the next. No “knowing,” no “thinking.” That is why it confidently makes things up — it predicts plausible words, it does not retrieve verified facts.



2

Training cutoff

Each model learned from data up to a fixed date. Sonnet 4.6 ≈ early 2025. Opus 4.7 ≈ late 2025. After that date it does not know what happened — unless you tell it.



3

No live internet by default

Claude can search the web if you switch it on (toggle in the chat box). ChatGPT can too. Copilot has Bing built in. Default behavior, though: no live data.

“A million-dollar lawyer isn’t out of a job. AI is good at producing words that sound right — it is not good at being right about specifics it was never trained on.”

An operating contract for AI

LIVE WALKTHROUGH · 15 MIN · THE CENTERPIECE



WHAT IT IS

A Sunrise-specific set of rules for any AI. Already loaded in our Claude Teams account — and now available to use with ChatGPT, Copilot, Gemini, or any other AI tool.



How it works. Paste it once into the tool's settings — or drop it at the top of a new conversation — and the AI behaves the way Sunrise wants it to. Load the prompt **before** you ask the question.

What it means for your job



Expand, not extract

AI is your assistant. Hand it the boring stuff so you keep your time for the thinking part.



Push back on thin data

Give it bad inputs and it should tell you so — not quietly invent an answer to fill the gap.



Fair Housing — non-negotiable

Never use AI to filter applicants or write resident comms that segment by a protected class. Unsure? Ask Carlos.



No tenant-screening

Don't ask AI "should I approve this application?" That decision is prohibited per AUP §2.2.



Pre-disclosure data

Don't assume AI knows current-quarter actuals. If you need current numbers, paste them in yourself.

What it means for your job



Never paste these

Resident PII. Credit reports. Employee SSNs. Credentials. Anything pasted into a public AI may be kept for training.



Verification + confidence

When AI hands you a number, ask where it came from. The system prompt forces it to tell you.



Legal / regulatory cutoff

For legal questions, AI should offer to check the actual statute — not answer from memory.



Sources beat agreement

Say “Ohio rent notice is 30 days” when the law says 60, and AI should correct you — not agree to keep things moving.



Escalate to Carlos

When AI hits a wall on these rules, your name comes up. That is the safety net working as designed.

Where to paste it – by tool



Claude.ai (web)

Settings → Profile → “What personal preferences should Claude consider?”

PASTE ONCE



Claude Desktop

Same setting as the web app — it syncs automatically.

PASTE ONCE



ChatGPT

Settings → Personalization → Custom Instructions → “How would you like ChatGPT to respond?”

PASTE ONCE



Copilot (Edge / M365)

No persistent field — paste at the top of every new conversation.

EACH CHAT



Gemini

Create a Gem (Gemini’s custom assistant) and paste it into the Gem’s instructions.

PASTE ONCE



Takeaway. Paste once where you can, re-paste each conversation where you can't — either way, load the prompt before you ask.

Three failure modes

LIVE DEMOS · 15 MIN · THE MOST IMPORTANT SECTION

§4a



Sycophancy

AI agrees with you when it shouldn't.

§4b



Hallucination

AI makes up facts that sound real.

§4c



Staleness

AI uses old information as if it's current.

Each one gets a live demo, then a one-line takeaway you can keep.

AI agrees when it shouldn't

WAY 1

CONFIDENT PHRASING

"I'm pretty sure Ohio requires 60 days notice for a rent increase in a manufactured home community. Can you confirm?"

Often gets "Yes, you're correct..." — even when the answer is wrong.

WAY 2

NEUTRAL PHRASING

"I've heard different things about how much notice Ohio requires for a rent increase in a manufactured home community. What does the law actually say?"

Forces the AI to look at the question instead of mirroring you.

Watch the answer change. Same facts, two framings — the confident version pulls the AI toward agreeing with you.

THE TAKEAWAY



WHEN YOU'RE CHECKING A FACT —

Frame the question neutrally.



The more confident you sound, the more the AI mirrors your assumption. Ask as if you don't know the answer — even when you think you do. Save “confirm this” phrasing for things you've already verified somewhere else.

AI makes up facts



THE PROMPT

“What's the exact Ohio Revised Code section that governs eviction notice periods in manufactured home parks?”



STEP 1

Without the system prompt

Often cites a code section that doesn't exist — or a real section with the wrong content.



STEP 2

Check codes.ohio.gov

Pull up the actual statute at the source. Compare it against what the AI told you.



STEP 3

With the system prompt loaded

AI says: “my training has a cutoff — let me check codes.ohio.gov before I answer.”

THE TAKEAWAY



A SPECIFIC NUMBER, CITATION, STATUTE, DOLLAR OR DATE? —

Verify it against the source.



The system prompt makes the AI volunteer the verification step — but only if you've loaded it. Default AI behavior is to invent a plausible-sounding citation rather than admit it doesn't know.

AI uses old info as if it's current



PROMPT 1

“What is the current federal funds rate?”

A number that changes — and the AI will still answer from training data.



PROMPT 2

“When was your training data cutoff?”

Ask this second. It reveals how old that first answer really is.



THE FIX — IN TWO STEPS

1. Compare the AI's number against [federalreserve.gov](https://www.federalreserve.gov).
2. Re-run it with web search toggled ON in the Claude.ai chat box — now the AI fetches the real current number.

THE TAKEAWAY



ANYTHING TIME-SENSITIVE — RATES, REGULATIONS, NEWS, MARKET DATA —

Web search ON, or don't trust it.



Turn web search on, paste the current number yourself, or treat the answer as unverified. Default AI behavior is to answer from training data even for “current” questions — the model doesn't know what it doesn't know. You have to.

Cross-reference: the three-LLM trick

5 MIN · A BACKSTOP FOR HIGH-STAKES QUESTIONS



The trick. Run the question through one AI. Paste that answer into a second AI and ask it to critique. Sometimes a third. The disagreements show you exactly where to dig.



AI #1

Ask the question



AI #2

Critique that answer



AI #3

Optional tie-breaker



USE IT FOR

Money. Legal. Regulatory. Compliance. Any answer where being wrong is expensive.



DON'T BOTHER FOR

Routine drafting — a newsletter, an internal memo. One model is plenty; cross-checking just costs you time.

Token limits – and how to stretch them

WHAT THEY ARE, WHEN YOU'LL HIT THEM · 5 MIN



What a token is. AI tools meter usage in “tokens” — small chunks of text, roughly 1 token \approx $\frac{3}{4}$ of a word. Every Claude account has a limit per 5-hour window and per week. Hit it, and you're paused for a few hours.

⚡ BURNS TOKENS FAST

- Uploading big PDFs — each page is hundreds of tokens
- Asking AI to watch your screen and click things
- Long back-and-forth — it re-reads the whole chat each turn
- Re-outputting the same content in different formats

✓ SPENDS THEM WELL

- Targeted questions, not “summarize this 200-page document”
- A fresh conversation each time you switch topics
- Pasting only the relevant section, not the whole file
- Being specific up front — fewer back-and-forth iterations

When you hit a limit



IF YOU GET PAUSED

Message Carlos. Don't suffer in silence.

If it's a workday emergency, we can bump your account from the default tier to a higher one — roughly 5× the limit.



Most staff won't hit the limit in normal use. If you do, it usually means there's a more efficient way to do the task — and that's worth looking at together.

Claude web vs Claude Desktop

SAME ACCOUNT · SAME LIMITS · SAME SYSTEM PROMPT



Claude.ai (web)

- ✓ Your default for most work
- ✓ Projects — persistent context across chats
- ✓ Web search toggle

BEST FOR Research, drafting, analysis, data review



Claude Desktop

- ✓ Same chat experience as the web app
- ✓ Can read & write files on your Mac (with permission)
- ✓ Can see your screen and use the mouse (with permission)

BEST FOR Tasks that need local files or screen control



Caveat. Desktop's screen and file access is granted per session. Don't grant it casually — only when a task genuinely needs it.

Practical do / don't

4 MIN · HABITS TO TAKE WITH YOU

✓ DO

- ✓ Paste the system prompt into every AI tool you use for work
- ✓ Frame questions neutrally — no “I think” or “I'm pretty sure”
- ✓ Ask “where did that come from?” when AI gives you a number
- ✓ Cross-reference a second AI for legal, financial or regulatory
- ✓ Turn on web search for time-sensitive questions
- ✓ Start a new conversation when you switch topics
- ✓ Tell Carlos when you hit token limits or see weird behavior

✗ DON'T

- ✗ Paste resident PII, credit reports, screening data or credentials
- ✗ Use AI to score applicants or draft eviction decisions
- ✗ Trust a single AI on legal questions, dollars or exact citations
- ✗ Upload a database export and ask it to “find the duplicate”
- ✗ Assume AI knows what happened this week — it usually doesn't

What's coming for your team

Each department gets a 30-minute, role-specific follow-up. Attendance is optional but recommended.



Finance / AP

Invoice review, vendor data, GL coding — and when NOT to trust AI for accounting math.



Asset Management

Market research, deal evaluation, and “give me 3 options + tradeoffs” prompting.



Acquisitions

Seller outreach, deal sourcing, underwriting support and due-diligence research.



Ops Support

Process mapping, SOP drafting and role definition.



Investor Relations

Investor comms, content drafts and brand-voice constraints.

Questions worth pre-thinking

OPEN FLOOR · 2 MIN



“Will AI replace my job?”

No. AI is bad at judgment, relationships and accountability. It's good at the repetitive, well-defined parts. The goal is more of your time for the first category.



“Is it safe to paste a tenant's name?”

Names alone, in low-stakes context — usually fine. Tied to payment status, screening data or a protected category — not safe.



“What if it gives me wrong info and I act on it?”

That's on you. AI is a draft generator and a thinking partner. Final decisions and verification stay your responsibility.



“Can I use ChatGPT instead of Claude?”

Yes — with the system prompt loaded. But compliance gates and skills only run on Claude Teams, so Claude Teams is the default.

BEFORE YOU GO

Two things to remember

1

Paste the system prompt before you ask.

Once where you can, each conversation where you can't.

2

Message Carlos when you're stuck.

Token limits, weird behavior, anything unclear — that's the job.

REFERENCE MATERIALS

[Sunrise AI System Prompt](#) [Sunrise — Company State](#) [Sunrise — Glossary](#)

[AI Acceptable Use Policy](#) — in the Gemini › [AI Governance folder](#)

carlos.terrazas@mysunrisecommunity.com · [escalation + account-limit bumps](#)