AI Fluency

Working with AI like an engineer

Assoc. Prof. Alon Grinberg Dana · alon@technion.ac.il · Department of Chemical Engineering, Technion · 11 October 2026

The goal: work with AI **effectively** (better results), **efficiently** (less time and effort), **ethically** (responsibly) and **safely** (you stay accountable). The AI executes. **You design.** Choose **augmentation**: you and the AI collaborate.

Under the hood: what to remember

- An LLM is a very large neural network trained to **predict the next token**. At each step it picks a likely token, **not the best answer**.
- Each neuron computes a weighted sum, then a **nonlinear function**. Without it, 100 layers collapse into one.
- It works now because of three things: the **Transformer** (2017), a lot of digital text, and GPU/TPU compute.
- More data and compute brought **new abilities** that nobody programmed.
- The model reads **tokens, not letters**. Hebrew costs more tokens than English. Numbers and SMILES strings split in odd places.
- The **context window** is its working memory: your prompts, its replies, files, search results, hidden system instructions. It overflows: old messages drop out. Start a new chat for a new task and restate what matters.
- **Watch out for:** hallucination (plausible but wrong), knowledge cutoff, different answers to the same prompt, and sycophancy (it agrees too easily).

The 4Ds, with what to do

D Delegation — who does what?

Before you involve AI

Know your **goal**. **Split** the work. Decide who does what, and the **mode**: automation, augmentation or agency. No clear goal, no good result.

Ask yourself

What does success look like? What is simple but slow? Where would a thinking partner help? What needs **my judgment**? Be an expert in your field first, an AI delegator second.

Platform awareness

Pick two or three systems and learn them well. Each has its own strengths. Check which **model** you use (fast, thinking, pro) and what it prioritises: speed, creativity, depth or accuracy. Check privacy settings.

Example: delegate or not?

Plotting script for a lab report: **delegate**, and check the plot. Formula sheet for an exam: **do it yourself**, writing it is the studying. Presentation on your thesis: **augment**, the story and claims stay yours.

D Description — say it well

Product, process, performance

Say **what** you want (context, format, audience, limits), **how** (steps, an example) and the **behaviour** (tone, style). The AI cannot read your thoughts. Explain it as to a new colleague: eager, smart, new to your project.

1 Give context

Say who you are and what for.

```
"I'm a 2nd-year med student, studying for pharmacology. Explain
how ACE inhibitors lower blood pressure, step by step."
```

2 Show examples

Two good examples beat a paragraph of rules.

```
"Five more, in this style:
Q: k = 0.1 s⁻¹, 1st order. t½?
A: ln 2 / k = 6.9 s.
Q: … A: …"
```

3 Constrain the output

Name the format.

```
"A table: density, strength, fatigue limit, cost per kg. Al 6061-
T6 vs Ti-6Al-4V. Then one paragraph: which, why."
```

4 Break it down

Work in steps.

```
"First the data model. Wait for my OK. Then the API endpoints.
Then the interface."
```

5 Give it space to think

Ask for reasoning before the answer, or pick a thinking model.

```
"Before answering, list the climate constraints and three
alternatives. Compare, then recommend."
```

6 Give it a role

The next word depends on who is speaking.

```
"You are a sceptical reviewer for a molecular biology journal.
Find the weakest point of my CRISPR knockout design."
```

Let the AI write your prompt

Paste **prompt_engineering.md** and your raw request. It asks 3 to 6 questions, you answer, it writes a better prompt but **never runs it**. Run that prompt in a **fresh chat**. Keep the file.

D Discernment judge it

Three lenses

Product: accurate, valuable, solves it? **Process**: logical steps, no circular reasoning? **Performance**: follows you, efficient? If you can't check the answer, **check the steps**: units, assumptions, shortcuts.

Right answer, wrong process

Bubble point of benzene/toluene (Antoine): the answer looked right, but it put 350 K into constants that expect °C, and the P vs. T graph was missing. It likely recalled a textbook value. Read the steps, not only the result.

The car wash test

"The car wash is 50 m away. Should I walk or drive?" Check that the answer makes sense for your actual goal (you are washing the car).

D Diligence own it

You are responsible

You own everything that comes out of your AI use: **rigorous, transparent, accountable**. Like driving: not just A to B, also rules, safety and other people.

Three pillars

Creation: know the system and your data. **Transparency**: disclose AI use. **Deployment**: you own its accuracy, so verify.

Ask yourself

Data: who owns and sees it? **Privacy**: is sensitive information allowed there? **Impact**: who is affected by errors? **Disclosure**: who needs to know?

Against hallucinations

Allow "I don't know", demand sources, and ask what to double-check.

```
"If you're not sure, say so."
"Quote the source for each claim."
"List what I should double-check."
```

Live: the LLM council

1 Ask several models

Paste the **same prompt** into 3 chats (for example ChatGPT, Gemini, Claude).

```
Explain the thermodynamics of protein
folding. Specifically, how folding
affects the entropy of the protein and of
the surrounding water.
```

2 Review, then synthesise

Paste the three answers, **unlabelled**, into each model for critique. One model chairs the synthesis. Agreement across independent models is the closest thing to high confidence.

3 Adversarial development

One model writes the code. A **different** model attacks it. You decide which critiques are real. Sycophancy is the failure mode; the adversary is the cure.

The files you received (`prompt/`)

File	What it is for
`prompt_engineering.md`	Turns a raw request into a better prompt. Turn 1: it asks you 3 to 6 questions. Turn 2: it writes the prompt. It never does the task itself. Attach it, then add your raw request.
`deep_research_prompt.md`	Writes **one** Deep Research prompt for your question, to paste unchanged into several tools (ChatGPT, Gemini, Claude, Perplexity).
`deep_research_consolidate.md`	Merges the answers of those tools into one audited table: checks that each source exists and says what is claimed, and shows where the tools disagree.
`example_protein_folding_entropy/`	A complete example of the council on the protein-folding question: the three answers (ChatGPT, Claude, Gemini) and `consolidated.md`.

References

4Ds Dakan & Feller, with Anthropic, AI Fluency Framework and course, CC BY-NC-SA 4.0

Hood Raschka, Build a Large Language Model (From Scratch), 2024

Hood Vaswani et al., Attention Is All You Need, NeurIPS 30, 2017

Hood Azevedo et al., Equal numbers of neuronal and nonneuronal cells make the human brain an isometrically scaled-up primate brain, J Comp Neurol 513:532, 2009

Hood Moonshot AI, Kimi K2 technical report, arXiv:2507.20534, 2025

Tools Tokenizer used in the demo: tiktokenizer.vercel.app

Prompts More prompts against hallucinations (Futurepedia)