



Workshop **How to Get the Most Out** **of AI for Beginners**

Kevin Crook

KevinCrook.com

Setup Help Before **the Main Workshop Starts**



Workshop

How to Get the Most Out of AI for Beginners

Kevin Crook

KevinCrook.com



Topics

AI Basics

AI Chatbot Basics

(ChatGPT, Claude, Gemini, and Grok are AI Chatbots)

Prompt Engineering Best Practices

Takeaway: 100 Example AI Life Hack Prompts

AI Ethics & Safety



AI Basics



AI Basics

**What is
Artificial Intelligence (AI)?**



What is Artificial Intelligence (AI)?

Formal:

John McCarthy

(1956 First AI Conference at Dartmouth)

Machines that simulate human intelligence
(learning, reasoning, problem solving).



What is Artificial Intelligence (AI)?

Funny:

A machine that can clean my house, cook, mow the lawn, do laundry, etc.

... without being told.



What is Artificial Intelligence (AI)?

Practical / Moving Goalposts:

Whatever computers can't do *yet*.

Once solved, we stop calling it AI.



Turing Test (1950)

Can a hidden computer fool a human into thinking it's talking to another human?



Human Intelligence

Brain Based

Thinking, reasoning, problem-solving, planning, understanding ideas, etc.

Nervous System Based

Reflexes, muscle memory, physical coordination, automatic skills, etc.

AI definitions and success measures are limited to human brain based intelligence.



Neat AI vs Scruffy AI

Neat AI

Brain based, clean logical rules,
easy for computers.

Scruffy AI

Nervous system based, messy logic,
hard for computers.

Moravec's Paradox

Things that are easy for computers are hard for
humans.

Things that are easy for humans are hard for
computers.



Do submarines swim?

Submarines don't swim like humans or animals, but can produce similar or better outcomes.

Do AIs think?

AIs don't think like humans do, but can produce similar or better outcomes.



Do AIs experience consciousness or have feelings, emotions, or subjective experiences?

No evidence they do, but they can be trained to act as if they do, loosely called

Philosophical Zombies (P-Zombies).

(My personal opinion: definitely no. But it's considered an open question, with AI companies teasing it for marketing purposes.)



Artificial General Intelligence (AGI)

AI that matches or exceeds human brain based intelligence across a wide range of intellectual tasks.



Animal Intelligence

Animals have unique capabilities humans lack.

Animal intelligence is absent from AI definitions and success measures.

AI systems exhibiting animal intelligences are useful.



AI Chatbot Basics

(ChatGPT, Claude, Gemini, and Grok are AI Chatbots)



Large Language Model (LLM)

AI Chatbots use LLMs under the hood.

(For the remainder of this workshop, assume the term AI refers to an AI Chatbot built on top of an LLM.)



Prompt

What you type to the AI.

Response

What the AI replies.

Chat

The back-and-forth conversation, which consists of a series of prompts and responses.

Context Window

The maximum text the AI can process at once.
In very long chats, the oldest parts get dropped.



Prompt Engineering

The skill of writing clear, effective prompts that get much better results from AIs.

LLM Engineering

Advanced techniques beyond prompting to build systems around the AI. Often requires computer programming and systems engineering skills (another free workshop).

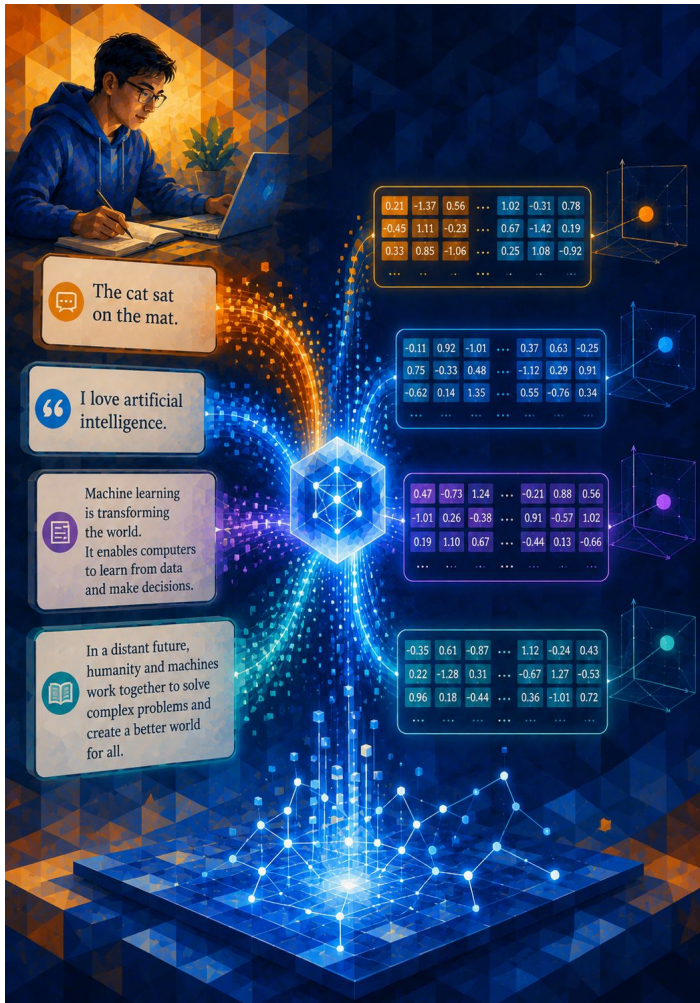


Tokens

Words or parts of words that AIs read and process.


Example token counts:

- Mary Shelley's *Frankenstein*:
 - 100,000 tokens / 75,000 words
- Bram Stoker's *Dracula*:
 - 210,000 tokens / 161,000 words
- Leo Tolstoy's *War and Peace*:
 - 750,000 tokens / 587,000 words
- William Shakespeare's *Complete Works*:
 - 1.2 million tokens / 885,000 words



Language as Lists of Numbers

Words, sentences, paragraphs are turned into lists of numbers (high-dimensional vectors).



Why do Graphics Cards speed up AIs?

Einstein used and popularized tensors. He compared them to riding a horse instead of walking.

Tensor: A multi-dimensional array of numbers.

High-dimensional vectors are a form of tensor.

Computer graphics rely heavily on tensors, so GPUs (Graphics Processing Units on graphics cards) are excellent at the math AIs need for high-dimensional vectors.

TPUs (Tensor Processing Units) are specialized chips built purely for AI tensor calculations.



Training Cutoff Date

Als are trained on text up until a certain cutoff date and have no built-in knowledge of later events.

Solutions:

- In Context Learning.
- Retrieval-Augmented Generation (RAG).



Analogy: Kids Learning to Read

Just like kids learning to read pick up extra story details, AIs trained on massive amounts of text absorb extra information they might never fully forget.

This can lead to unexpected or unwanted associations in responses.



Hallucinations

Als sometimes confidently make up facts that sound real but aren't.

Als predict plausible-sounding text. They don't look facts up. Gaps in training get filled with convincing guesses.

Always fact check important information.

Solutions that cut down hallucinations:

- In Context Learning.
- Retrieval-Augmented Generation (RAG).



Lost in the Middle

In very long chats or documents, AIs can ignore or forget information in the middle.

Solutions:

- Chunking: break up into chunks, ask for summaries, then combine summaries.
- Newer LLM architectures such as State Space Model (SSM) + Transformer Hybrids.



Context Rot

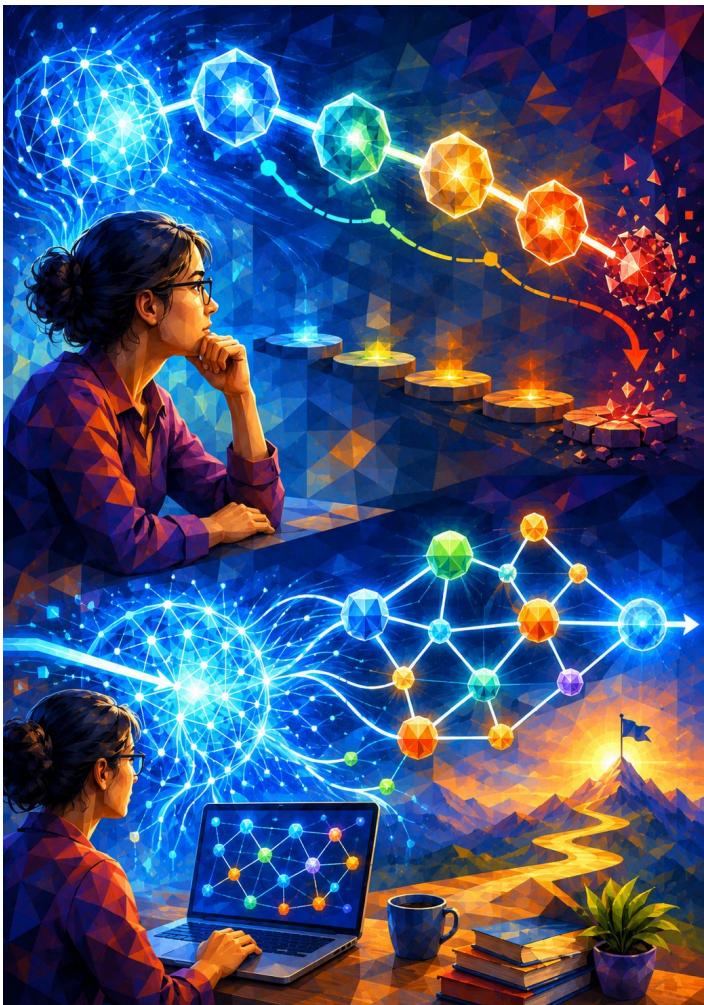
Over a long chat, earlier prompts, responses, files, etc. get diluted or forgotten.

Task Rot

Over a long chat, the AI can drift away from the original task or instructions.

Solutions:

- Ask for a summary of the chat so far, start a new chat, and paste the summary in.
- Newer LLM architectures such as State Space Model (SSM) + Transformer Hybrids.



Chain Rot (Multi-Hop Reasoning Errors)

Multi-hop Chain: A relates to B relates to C relates to D...

AI's reasoning gets worse with each hop in the chain away from the original information.

Solution:

Knowledge Graph Engineering to create Knowledge Graphs, then put them in context or RAG (another free workshop).



Future: SSM + Transformer Hybrid LLMs

Today's AIs (ChatGPT, Claude, Gemini, Grok) are built on Transformers, which re-read the whole chat for every word they write.

State Space Models (SSMs) instead keep a running summary as they read.

Hybrids combine both:

- Faster and cheaper on the same hardware (GPUs, TPUs).
- Better with long chats and documents. Less lost in the middle, context rot, and chain rot.
- Well suited to RAG including Knowledge Graphs.

NVIDIA's Nemotron models are free to download and run on your own high end hardware (self-hosted).



Prompt Engineering Best Practices



Start Simple with Interactive Chat

Don't try to write the perfect prompt on the first try.

Trial and error is normal and expected.

Begin with a basic prompt, then refine it step by step.



The Secret Sauce of Prompt Engineering

Triple Power:

Meta-Prompting + Role + COSTAR

Major prompt engineering contest winners and top finishers use this technique.

Most important takeaway from this workshop!



Meta-Prompting

Think of AI as your personal expert prompt engineer.

AI can do prompt engineering better than most humans can (maybe all?).

Use 2 chats:

- Chat 1: ask the AI to help you create and refine prompts.
- Chat 2: run the finished prompts.



Role + COSTAR

Role: Assign the AI a specific role or expertise.
("You are an experienced travel agent...")

Context: Background info the AI needs.

Objective: Exactly what you want done.

Style: Writing style you want.
(formal report, casual, technical, etc.)

Tone: The attitude of the response.
(friendly, professional, persuasive, etc.)

Audience: Who the response is for.
(beginners, experts, kids, your boss, etc.)

Response Format: How to deliver the answer.
(bullet list, table, paragraph, long form, with examples, etc.)



In Context Learning

Als learn from whatever you put in the context window with no retraining needed.

File Uploads and Projects are great examples: Your documents become knowledge the AI can use.

Meta-prompting works here too:

Ask the AI to help generate the context material itself (summaries, examples, background docs).



Manage Context Rot & Task Rot

Periodically ask the LLM to summarize the conversation so far. This keeps both of you focused.

At the first signs of context rot or task rot, cut your losses! Ask for a summary, start a fresh new chat, and paste it in.



Handle Lost in the Middle Issues

For large documents or long contexts:

- Break them into smaller chunks.
- Process each piece separately.
- Then ask the LLM to combine or summarize the results.



Ask for Chain of Thought

For complex tasks, add this to your prompt:

"Explain your reasoning step by step."

You'll get noticeably better answers, and you can check the reasoning yourself.



Try Different LLM Versions and Vendors

LLM vendors offer multiple versions (fast, deep, beta, etc.).

There are multiple vendors to choose from (ChatGPT, Claude, Gemini, Grok, etc.).

Stuck or unsatisfied? The same prompt can get very different results elsewhere.



Build or Use a Prompt Library

Save your best prompts in a document as you go for easy reuse and adaptation or search for public "prompt libraries" for great starting points.



Share Chat Links

Share useful chats using the LLM's share link.

Great for showing coworkers, family, or friends exactly how you got a result, prompts and all.



Takeaway:
100 Example
AI Life Hack Prompts



AI Ethics & Safety



Privacy

If you are not buying something,
then you are the product.

Check your AI Chatbot's privacy policies
and settings (often varies by paid tier).

Avoid pasting sensitive information into an
AI Chatbot.



Environmental Impacts

Training and running AIs uses a lot of electricity and water.

The AI race has worldwide utilities recommissioning and extending coal plants and building new fossil and nuclear plants.



Bias

Garbage In / Garbage Out (GIGO).

Biased training data leads to
biased answers.



Copyright Violation

Als are sometimes trained on copyrighted material (books, articles, art, music), usually without permission or royalties.


Als can sometimes reproduce copyrighted content in their responses.

Ongoing lawsuits and settlements are deciding what's legal, with the rules still being written.



Echo Chambers / Filter Bubbles

Als can reinforce your existing beliefs by agreeing too easily or feeding you similar perspectives.



Academic Cheating

When prohibited, using AIs to do school work.

Important question:
Are students learning?

Note that AIs can help learning if used properly.



Self-Harm and Harm to Others

LLMs should have guardrails that refuse to help with self-harm, suicide, harming others, dangerous acts, etc.



Workforce Realignment & Job Displacement

Basic AI skills are increasingly required for many roles at all levels.

AI is automating tasks and making some jobs obsolete, especially entry-level.



Autonomous Weapons

Weapons that can select and attack targets with little or no human control.

Big debate: should a machine ever decide to take a human life?

AI arms race.




Cybersecurity Issues

AI Helps Attackers

More convincing phishing emails, voice cloning scams, malware assistance, finding software vulnerabilities.

AI Helps Defenders

Detecting attacks, spotting fraud, finding and fixing vulnerabilities faster.



Black Box AI

The norm, no transparency,
We don't see how it reaches answers.

White Box AI

Rare, transparent.
We see how it reaches answers.

Explainable AI (XAI)

XAI is white box AI that clearly shows how it reaches answers.



False Accusations

Claims such as:

- Credit / Loan / Insurance / Redlining Decision
- Detecting AI Use in Writing
- Detecting Lies
- Guiding Prison Sentencing / Probation

Black box, no transparency.

Accuses but can't explain why to allow humans to defend themselves.

Vendors refuse to submit their AIs to formal, independent, scientific scrutiny.



Intelligence Explosion

Humans create an AI smarter than themselves.

That AI creates a new AI smarter than itself.

Process repeats, with each AI generation creating a new AI smarter than itself.

Each AI generation arrives faster than the last AI generation.



Workshop **How to Get the Most Out** **of AI for Beginners**

Kevin Crook

KevinCrook.com

Thanks for attending and **good luck on your AI journey!**