Tab 1

Quick Reference for the Workshop: How to Get the Most Out of AI for Beginners

Kevin Crook

AI Basics

- **What is Artificial Intelligence (AI)?**
  - **Formal:** John McCarthy (1956 Dartmouth) - machines that simulate human intelligence (learning, reasoning, problem-solving).
  - **Funny:** A machine that can clean my house, cook, mow the lawn, do laundry… without being told.
  - **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.
  - **Non-brain-based:** reflexes, muscle memory, physical skills.
  - Most AI progress targets brain-based intelligence.
- **Animal Intelligence**
  - Animals have unique capabilities humans lack.
- **Artificial General Intelligence (AGI)**
  - AI that matches or exceeds human intelligence across a wide range of intellectual tasks.
- **Intelligence Explosion**
  - Humans create AI smarter than themselves.
  - That AI creates even smarter AI.
  - Each new generation is smarter and arrives faster.
- **Neat AI vs Scruffy AI**
  - **Neat AI:** Clean logical rules (easy for computers).
  - **Scruffy AI:** Messy, biology/evolution-inspired.
  - **Moravec's Paradox:** Computers crush hard things (chess, math) but struggle with easy human things (common sense, walking).
- **AI Functionalism**
  - We care about what the AI can do, not whether it truly thinks or feels.
  - Do machines think?
  - Do submarines swim?
- **Philosophical Zombies (P-Zombies)**

    - AIs that act and talk exactly like humans but have no inner feelings or consciousness.
- **Asimov's Three Laws of Robotics**
    - 1) A robot may not injure a human being or, through inaction, allow a human being to come to harm.
    - 2) A robot must obey orders given by human beings except where such orders would conflict with the First Law.
    - 3) A robot must protect its own existence as long as such protection does not conflict with the First or Second Law.
- **AI Dystopia**
    - **Rossum's Universal Robots (R.U.R.):** Coined the term robot and features a robot uprising against humans.
    - **Paperclip Maximizer:** Harmless goal of making paperclips accidentally destroys the world.
    - **Tale of the Omega Team:** AI that takes over the world's computing resources with no physical robots involved.
- **AI Utopia**
    - AI solves all major human problems (disease, poverty, climate change, etc.)
    - Creates a world of unprecedented abundance, creativity, and leisure.

LLM Basics

- **Large Language Model (LLM)**
    - ChatGPT is an LLM
- **Tokens**
    - Words or parts of words that LLMs read and process.
    - Example token counts:
        - King James Bible:
            - 1.2 million tokens / 785,000 words
        - Shakespeare's Complete Works:
            - 1.2 million tokens / 885,000 words
        - Tolstoy's War and Peace:
            - 750,000 tokens / 587,000 words
- **Language as Lists of Numbers**
    - Words, sentences, paragraphs are turned into lists of numbers (high-dimensional vectors).
- **Why do Graphics Cards speed up LLMs?**
    - Tensor: A multi-dimensional array of numbers.
    - High-dimensional vectors are a form of tensor.
    - Einstein used and popularized tensors. He compared them to riding a horse instead of walking.
    - Computer graphics rely heavily on tensors, so GPUs (graphics cards) are excellent at the math LLMs need for high-dimensional vectors.

    - TPUs (Tensor Processing Units) are specialized chips built purely for LLM tensor calculations.
- **Attention**
    - Mechanism that lets the LLM calculate how important each token is compared to all other tokens in the text.
    - Transformer architecture (with multi-headed attention) made this possible.
    - Allow the model to focus on relevant information anywhere in the text.
- **Training Cutoff Date**
    - LLMs are trained on text up until a certain cutoff date and have no built-in knowledge of later events.
    - Solution: Retrieval-Augmented Generation (RAG) - LLMs pull in fresh web or document text.
- **Context Window**
    - The maximum amount of text an LLM can process at once.
    - If a chat is very long, only the most recent portion that fits in the context window is used and older parts get dropped.
- **Prompt Engineering**
    - The skill of writing clear, effective prompts that get much better results from LLMs.
- **LLM Engineering**
    - Advanced techniques beyond prompt engineering that deliver much better results from LLMs, especially when prompting alone reaches its limits.
- **Challenges**
    - **Dick and Jane Effect**
        - Just like children learning to read from Dick and Jane books pick up extra story details, LLMs 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**
        - LLM confidently makes up facts that sound real but aren't.
        - Always fact check important information.
    - **Lost in the Middle**
        - In very long chats or documents, the LLM can ignore or forget information buried in the middle.
    - **Context Rot**
        - Over a long chat, earlier prompts, responses, files, etc. get diluted or forgotten.
    - **Task Rot**
        - In a long complicated prompt, the actual task gets diluted or forgotten.
        - A common, unwanted side effect of formal prompt engineering frameworks.
- **Ethics**
    - **You are the Product**
        - If you are not buying something, then you are the product.

    - **Privacy**
        - User data from chats is often collected, stored, and used for training, and may show up in responses to prompts from other users.
    - **Environmental Impacts**
        - Training and running LLMs uses a lot of electricity and water.
    - **Bias**
        - Garbage In / Garbage Out (GIGO).
        - Biased training data leads to biased answers.
    - **Copyright Violation**
        - LLMs are trained on copyrighted material and can sometimes produce copyrighted text.
    - **Echo Chambers / Filter Bubbles**
        - LLMs can reinforce your existing beliefs by agreeing too easily or feeding you similar perspectives.
    - **Academic Cheating**
        - Using LLMs to do school work without proper credit or disclosure.
        - Important question: Are students learning?
    - **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 LLM skills are increasingly required for many roles at all levels.
        - AI is automating tasks and making some jobs obsolete, especially entry-level. positions.
    - **Explainable AI (XAI)**
        - **Black Box AI:** the norm, no transparency, we don't see how it reaches answers.
        - **White Box AI:** rare, more transparent.
        - True XAI would clearly show its reasoning steps.

Prompt Engineering Best Practices

- **Start Simple with Interactive Chat**
    - Use an interactive chat style.
    - Trial and error is normal and expected.
    - Begin with a basic prompt, then refine it step by step.
- **Meta-Prompting**
    - Ask the LLM to help you write better prompts.
    - The most powerful technique available.
    - The most important take away from this workshop.
- **Use Role / Persona Prompting**
    - Assign the LLM a specific role or expertise.
- **Specify Output Format and Give Examples (Few-Shot Prompting)**

    - Tell the LLM exactly how you want the answer: bullet list, table, summary paragraph, long form paper, etc.
    - Give 1 to 3 examples of the kind of output you want (few-shot prompting).
    - Even one good example often dramatically improves results.
- **Use File Uploads**
    - Most LLMs allow you to upload files, documents, etc. which go into the context window.
- **Use Projects**
    - Projects keep multiple related chats organized.
    - You can set a master prompt and master files that automatically appear at the front of every context window in the project.
- **Share Chat Links**
    - Share useful chats using the LLM's share link.
- **Manage Context Rot & Task Rot**
    - Help prevent context rot by periodically asking the LLM to summarize the conversation to stay focused.
    - At the first signs of context rot or task rot, cut your losses!
    - Ask for a summary of the conversation so far, then start a fresh new chat.
- **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**
    - Add "explain your reasoning step by step" for better quality answers on complex tasks.
- **Try Different LLM Versions and Vendors**
    - LLM vendors often have multiple versions (fast, deep, beta, etc.).
    - There are multiple LLM vendors to choose from.
- **Build or Use a Prompt Library**
    - Save your best prompts in a document for easy reuse and adaptation or search for public "prompt libraries" for great starting points.
- **Use Frameworks**
    - **COSTAR+Role**
        - COSTAR is Context / Objective / Style / Tone / Audience / Response format.
        - COSTAR+Role adds Role.
        - Most popular for beginners.
        - Has won major prompt engineering contests.
    - **CRISPE**
        - Capacity / Role / Insight / Statement / Personality / Experiment
        - Popular for more complex or creative tasks.
- **Use Multi-Agent Prompting**
    - Grok only as of July 2026 (other LLMs require computer programming).
    - Default named agents have specific roles.
    - You can specify your own agents.

Life Hacks Using LLMs

- **Generate Text**
  - Create emails, essays, cover letters, social media posts, stories, poems, scripts, etc.
- **Review and Rewrite Text**
  - Edit, improve, or polish your writing (clarity, grammar, tone, professionalism).
- **Summarize Text**
  - Condense long articles, books, meetings, or videos into key points.
- **Explain Concepts**
  - Break down complex topics in simple terms (great for school, homework, or learning anything new).
- **Learning & Tutoring**
  - Get personalized explanations, practice questions, study help, or test prep.
- **Brainstorm Ideas**
  - Generate creative ideas, solutions, or options for school projects, problems, or life.
- **Plan & Organize**
  - Create travel plans, meal plans, workout schedules, project timelines, or to-do lists.
- **Decision Making**
  - Weigh pros and cons, compare options, and get recommendations for life choices.
- **Answer Questions & Research**
  - Get clear answers, comparisons, or background information on almost any topic.
- **Role-Play**
  - Act as a tutor, coach, mock interviewer, career advisor, or expert in any field.
- **Translate Languages**
  - Translate text between languages with strong accuracy.
- **Creative Work**
  - Generate ideas and content for art, images, music, videos, and creative projects.
- **Analyze Text**
  - Extract insights, tone, sentiment, or main ideas from documents or conversations.
- **Help with Coding (Computer Programming)**
  - Write, debug, or explain code (even for non-programmers).

Disclaimers

I am a lecturer at the University of California, Berkeley, in the Master's Program in Data Science and the Master's Program in Cybersecurity. These free AI workshop materials are provided on my personal website and reflect my own views and opinions, which are not necessarily those of UC Berkeley.