




AI FOR ENHANCED PRODUCTIVITY: MASTERING PROMPTS AND A MODERN TOOLKIT


Jack Francis, Ph.D.

AI Solutions Director – Davidson Technologies



A Little About Me

- **Background:** Ph.D. in Industrial and Systems Engineering from Mississippi State
- **Top 3 Core Values:** Authenticity, Learning, Fun
- **Fun!:**
 - **Four dogs:** (Pomeranian to Great Pyrenees)
 - **Traveler:** (Top 3: Tokyo, Rome, and Barcelona)
 - **Competitor:** (Chess, Golf, Card Games, anything I can win)
- **Why AI?:** Teaching machines feels like magic




What we will cover today

- Current LLMs -> “How do I know which LLM to use?”
- Context Importance -> “Why do LLMs forget what I tell them”
- Prompt Engineering -> “How do I make the LLM do what I want”
- AI Tools -> “What exists besides ChatGPT”
- Agentic AI -> “What comes after ChatGPT?”

Large Language Models (LLMs)

- LLMs, like ChatGPT or Gemini, work by predicting the next word based on learned patterns and the **context** provided
- Sometimes next word is plausible but wrong -> **Hallucination**



Current LLMs

- ChatGPT is a collection of LLMs, not a single LLM.
- Instant models are fast and powerful for everyday work
- Thinking/Pro models solve harder work tasks more effectively
- Free versions are limited to Instant models or may have a few messages to Thinking model.
- Generally, you must be at the \$20/month tier to access Thinking/Pro



Type	OpenAI	Google	Anthropic
Instant	GPT-5.2 Instant	Gemini 3 Flash	Claude Haiku 4.5
Thinking	GPT-5.2 Thinking	Gemini 3 Thinking	Claude Sonnet 4.5
Pro	GPT-5.2 Pro	Gemini 3 Pro	Claude Opus 4.5

Example of Instant vs Thinking/Pro

- Help me plan a trip to Berlin in March. I have 5 days to spend and want to see many historical sites.


Best LLM for You

- For work tasks, the difference between OpenAI, Google, and Anthropic comes down to personal preference.
- Try out the Thinking / Pro variants of your favorite tool for a month
- The one exception! For coding tasks, stick with GPT 5.2 Codex or Claude Opus 4.5



Context Window

- Each LLM has a set amount of information it can process at one time called a **context window**.
- LLMs think in “tokens”, which are like words, but sometimes words are broken into multiple tokens
- 1000 tokens is about 750 words. Most models have a context window of over 128,000 tokens



Context Clutter

- If you exceed the context window, the LLM will “forget” information
 - For very long chats, ChatGPT, Gemini, and Claude will sometimes automatically summarize the previous messages to reduce context
- **Important:** When the AI answers your prompt, all previous messages in the thread (both yours and the AI’s) are used as **context for the answer.**
- If you have a single long thread with lots of messages, or discuss multiple topics in the same thread it will lead to **context clutter**, which will significantly increase the chance of hallucination

One Chat = One Project

- Don't ask for email summarization, a recipe for the potluck, and travel advice all in one chat. **Risk of Hallucinations!**
- If your next prompt is very different from what you've talked about so far, **start a new chat!**
- If the AI starts hallucinating or getting lazy, **start a new chat!**
- Best Practice: Try to keep individual chats to less than 50 messages to keep the AI focused.
- **Pro Tip:** If you need to keep going, ask the AI to summarize your conversation so far and start a new chat!

Master Prompt Formula

- When LLMs were first released, prompt engineering was very important to get good results.
- New LLMs have much better **system prompts** (default instructions)
 - GPT 5.2 system prompt is around 20,000 tokens (30 single space pages)
 - Include information on writing style, tools, and persona
- Prompting is still very important for structuring output and ensuring appropriate context

Prompt Formula: Persona + Context + Task + Constraint + Format

Prompt Examples

- Generic Prompt:
 - Write an email about the project delay
- Persona + Context
 - You are a Senior Project Manager. Write an email about the Alpha Go 2 launch being delayed by 2 weeks due to server issues
 - You are a frustrated Network Engineer. Write an email about the Alpha Go 2 launch being delayed by 2 weeks due to server issues
- All components
 - You are a Senior Project Manager. Write an email about the Alpha Go2 launch being delayed by 2 weeks due to server issues. Your task is to update the stakeholders. Do not apologize; focus on the solution. Use a bulleted list for the next steps.

Two More Superpowers

■ Show, Don't Just Tell

- LLMs are fantastic at following provided structure.
- If you have examples of what “right” looks like, share those with the LLM
- Provide a template to get consistent results

■ Iteration

- Don't just send one prompt! Make it a conversation
- LLMs are remarkably flexible, if the output doesn't match what you need, just ask it to fix the issues
- Examples:
 - “Change the tone of this writing to [insert style]”
 - “Shorten this email”
 - “I have back issues so make a substitution for deadlifts”

Prompt Examples

- Generic Prompt:

- Write a bio for Sam Altman.

- Show, Don't Just Tell

- Write a bio for Sam Altman. Here is an example of the exact format I want you to follow: [Name] is a [Title] who loves [Hobby]. He is currently focused on [Goal].

- Iteration

- That's a little too boring. Change the tone to be witty and write it like a money-hungry pirate

AI Tools Buffet

Business Ops	Creativity	Productivity	Other
Deep Research - OpenAI, Gemini	Image Generation - Midjourney	AI Writing - Grammarly, Jasper	Code Generation - GitHub Copilot, Cursor
Chat with Documents - NotebookLM	Video Generation - Sora, Veo	Presentation Generation - Gamma, Beautiful.ai	AI Search Engines - Perplexity, OpenAI, Gemini
Agentic Coding - Replit, Lovable	Audio Generation - Suno	AI Note Taking - Otter AI, Fireflies	AI Data Analyst - Julius, Excel Copilot
Automation - Zapier, n8n	Text-to-Speech - ElevenLabs	Email Management - Superhuman	Intelligent Scheduling - Motion, Clockwise

AI Tools Buffet

Business Ops	Creativity	Productivity	Other
Deep Research - OpenAI, Gemini	Image Generation - Midjourney	AI Writing - Grammarly, Jasper	Code Generation - GitHub Copilot, Cursor
Chat with Documents - NotebookLM	Video Generation - Sora, Veo	Presentation Generation - Gamma, Beautiful.ai	AI Search Engines - Perplexity, OpenAI, Gemini
Agentic Coding - Replit, Lovable	Audio Generation - Suno	AI Note Taking - Otter AI, Fireflies	AI Data Analyst - Julius, Excel Copilot
Automation - Zapier, n8n	Text-to-Speech - ElevenLabs	Email Management - Superhuman	Intelligent Scheduling - Motion, Clockwise

What we will cover today
Will be covered in later workshops

Deep Research

- For tasks that require significant Googling and consolidation of lots of sources of information, Deep Research is a great tool
- Goes a step beyond the Thinking/Pro models by repeatedly searching the web for sources and writing 20+ page reports
- Example use cases:
 - Competitive Analysis / Market Research
 - Deep dive on new technical topic area
- **Pro Tip:** You can use Deep Research to get lots of relevant information and then use that as **context** for a follow-on chat

NotebookLM

- Google tool that significantly enhances Deep Research through targeted capabilities:
 - Audio Overview (Generate a podcast)
 - Video Overview (Generate an AI-explainer video)
 - Flashcards (Generate study help)
 - Quiz (Generate questions about your content)
 - Slide Deck (Generate slides based on content)
- Example Notebook:
<https://notebooklm.google.com/notebook/c0cbefb0-13d6-4aaf-b5a6-95eb302e93a2>

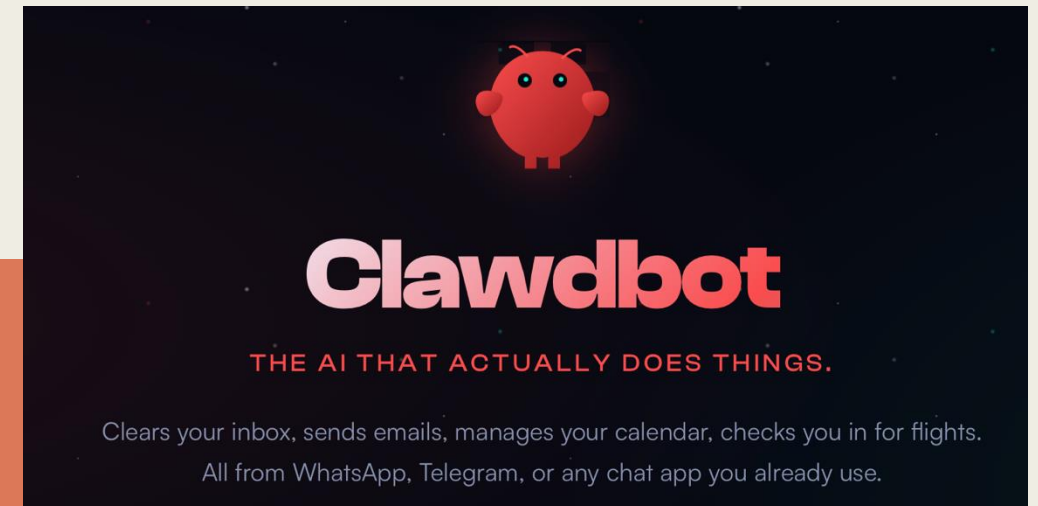

Agentic Coding

- Replit, Lovable, Bolt, Rork and many others provide non-coders with the ability to develop and deploy full applications
- Use cases:
 - Small applications that only you would use
 - Rapidly prototyping features before showing development team
 - Learning to code / using tools as launchpad
- A local Huntsville company uses Replit to accelerate Software Development by forcing Product Owners to build features in Replit before Engineering Team sees them.
 - Improves communication
 - Reduces development time

The Future is Here: Agentic AI

■ Chatbots to Agents

- No longer just prompting; giving Agent goal to achieve
- Agent has tools to use and continues to iterate



The Future is Here: Agentic AI

Measuring AI Ability to Complete Long Tasks


Linear Scale

Log Scale

50% Success

80% Success

The time-horizon of software engineering tasks different LLMs can complete 80% of the time



The Future is Here: Agentic AI

Measuring AI Ability to Complete Long Tasks


Linear Scale

Log Scale

50% Success

80% Success

The time-horizon of software engineering tasks different LLMs can complete 50% of the time



The Future is Here: Agentic AI

- Upcoming workshops focused on Agentic AI
 - Tuesday 10:25 – 11:40 AM: Building your First AI Agent
 - Tuesday 3:10 – 4:25 PM: From Prompts to Processes
 - Wednesday 10:25 – 11:40AM: AI and Cyber: Automating Exploitation with AI Agents

Key Takeaways

- Current LLMs
 - Try Thinking/Pro
 - OpenAI/Google/Anthropic all about the same
- Context Importance
 - One Chat = One Project
 - Limit to ~50 messages
- Prompt Engineering
 - **Prompt Formula:** Persona + Context + Task + Constraint + Format
 - Show Don't Just Tell and Iterate!
- AI Tools
 - Deep Research, NotebookLM, Replit
- Agentic AI
 - Expect to see Agentic AI tools continue to improve

AI Tools

Business Ops	Creativity	Productivity	Other
Deep Research - OpenAI, Gemini	Image Generation - Midjourney	AI Writing - Grammarly, Jasper	Code Generation - GitHub Copilot, Cursor
Chat with Documents - NotebookLM	Video Generation - Sora, Veo	Presentation Generation - Gamma, Beautiful.ai	AI Search Engines - Perplexity, OpenAI, Gemini
Agentic Coding - Replit, Lovable	Audio Generation - Suno	AI Note Taking - Otter AI, Fireflies	AI Data Analyst - Julius, Excel Copilot
Automation - Zapier, n8n	Text-to-Speech - ElevenLabs	Email Management - Superhuman	Intelligent Scheduling - Motion, Clockwise

LinkedIn



Contact Info:

jackfrancis@davidson-tech.com

jackbfrancis@gmail.com