

Cochise College Center for Lifelong Learning

Introduction to AI

For Lifelong Learners

George Self

October 27, 2026

Email me at: contact@georgeteaches.com

Visit: georgeteaches.com

Visit: cochiseai.com

Introduction

Welcome to **Introduction to Artificial Intelligence**.¹ Some of you have already spent time chatting with an AI tool: you know what a prompt is, you have seen AI give impressive answers, and you have probably seen it give wrong ones. Others are starting fresh today, and that is just as good a place to begin. Either way, everyone in this room is welcome, and no prior experience is required.

Whether this is your first time typing a prompt or your fiftieth, today's focus is on three questions that separate casual AI users from confident ones: *Which tool should I use? When can I trust what AI tells me? And how do I put AI to work on things that actually matter in my life?*

What You Need

Everything we do today runs in a web browser. Our main tool is **DuckDuckGo AI Chat** at <https://duck.ai/chat>, which is free and does not require an account. It has a daily limit on how much each person can use it. If you see a *Daily limit reached* message, do not worry: switch to Gemini (<https://gemini.google.com>) or ChatGPT (<https://chatgpt.com>) and keep going.

Learning Goals

By the end of this session, you will:

- Know the major AI tools and what makes each one distinct
- Understand which tasks AI handles well and which it handles poorly
- Recognize when AI output needs to be verified and know how to verify it
- Apply AI to practical scenarios from your own daily life

Big Idea

AI works best as a knowledgeable assistant. Use it to *augment* your judgment, not replace it.

¹Written by George Self with assistance from Claude. It was formatted with $\text{\LaTeX}$.

Module 1: The AI Ecosystem

Most people who use AI regularly have used only one tool: ChatGPT. That is a perfectly reasonable starting point, but it is a bit like only ever eating at one restaurant. The AI landscape has grown quickly, and several free tools are now available that each take a different approach to the same basic job.

Today we will look at seven tools. All are free to use, though a few require creating an account first. Each was built by a different organization, trained on somewhat different data, and designed with a different philosophy. Seeing them side by side is one of the fastest ways to deepen your understanding of what AI actually is.

Our Main Tool

DuckDuckGo AI Chat (<https://duck.ai/chat>): The privacy-focused option, and the tool we will use for most activities today. DuckDuckGo does not track your searches, and its AI chat works the same way: no account is needed, and each chat is labeled *Privacy protected and anonymized by DuckDuckGo. No AI training*. What makes it especially useful for comparison is that it gives you access to several different AI models, and you can ask a second model for its opinion without leaving the page.

Finding Your Way Around DuckDuckGo AI Chat

- **The model picker** sits at the lower right of the message box and shows the name of the model answering you. Click it to choose a different model. The free list changes from time to time, and some models are marked as using up your daily limit sooner. Appendix D describes each model.
- **Fast or Reasoning.** Next to the model picker you can choose **Fast** or **Reasoning**. Reasoning is more thorough but reaches your usage limit two to five times sooner, so leave it on **Fast** today.
- **The Tools menu** (next to the paperclip) offers **Create Image**, **Web Search**, and **Customize responses**, where you can set tone and length.
- **The 2nd opinion button** appears below each response, at the lower left. Click it and a different model answers the same question.
- **The paperclip** lets you attach files: up to three at a time, in PNG, JPG, WebP, GIF, or PDF format, with 15 pages or fewer per file. Plain text (.txt) files are not accepted.
- **The daily limit.** DuckDuckGo does not publish the exact amount. When you reach it, you will see *Daily limit reached* along with the time it resets. Switch to another tool to keep working.

Other Tools You Can Use Without an Account

ChatGPT (<https://chatgpt.com>): Built by OpenAI, this was the first AI assistant made available to the general public and remains the most widely used. It is the familiar baseline against which most other tools are compared. Click **Stay logged out** if prompted to sign in. **Note:** Logged-out use is subject to rate limits. If ChatGPT stops responding or asks you to wait before sending another message, move on to one of the other tools and circle back later.

Gemini (<https://gemini.google.com>): Google's AI assistant. Many people assume that because Google has indexed so much of the internet, Gemini must have a knowledge advantage. Testing it against the others will show whether that assumption holds up. Without signing in you get the basic chat; signing in with a Google account adds more features.

Tools That Require an Account

Claude (<https://claude.ai>): Built by Anthropic, Claude is known for careful, well-reasoned answers and a cautious, thoughtful approach to uncertain or sensitive questions. It is also widely regarded as one of the strongest tools available for writing tasks. It is a useful point of comparison for how differently AI tools handle hedging.

Copilot (<https://copilot.microsoft.com>): Microsoft's AI assistant, closely integrated into Windows, Edge, and Microsoft 365 apps like Word and Excel. The website works without signing in, but a free Microsoft account unlocks chat history, image creation, and longer conversations.

Meta AI (<https://meta.ai>): Developed by Meta, the company that owns Facebook and Instagram. It is built on Meta's own AI models and tends toward a conversational, approachable tone.

Perplexity (<https://perplexity.ai>): An AI tool designed around research. Unlike most chatbots, Perplexity shows citations alongside its answers so you can see where the information came from and go check it yourself. **Note:** Free use is tightly limited, and without an account you may be asked to sign in after just a couple of queries.

One Prompt, Many Responses

These tools will often give you different answers to the same question. That is not a bug. It reflects differences in training data, design choices, and how each tool handles uncertainty. Even the *same* tool can give you a different answer if you ask the same question twice. Learning to notice those differences makes you a much more effective AI user.

Activity 1.1: Compare the Tools: A Topic You Know

Open DuckDuckGo AI Chat in one browser tab and Gemini or ChatGPT in another, and try this prompt in both:

What is the most interesting thing that happened in Tombstone, Arizona in the 1880s, and why does it matter today?

After reading the responses, consider:

- Which response was most accurate based on what you already know?
- Which was easiest to read?
- Did any tool add unnecessary warnings or hedge more than others?

Activity 1.2: Compare the Tools: Tone and Personality

Not every difference between AI tools is about facts. Try this question in two tools:

I am nervous about giving a toast at my granddaughter's wedding next month. What advice do you have for me?

Compare the responses:

- Which tool felt the warmest or most personal? Which felt the most formal or businesslike?
- Did any tool ask a follow-up question before giving advice?
- Was one response noticeably longer or shorter than the others?
- Would you describe any tool's "personality" in a single word? Do the others agree?

Activity 1.3: Ask for a 2nd Opinion

Back in DuckDuckGo AI Chat, start a new chat and ask:

What is the most important fact to know about 1880s mining camps that most people would not expect?

1. Read the answer. Notice the name of the model above it, and whether the

answer includes links to websites below it.

2. Click the **2nd opinion** button below the response. A different model answers the same question.
3. Compare the two answers. Do they agree? Does one fit what you know about the history of the American West better than the other?
4. Ask each one:
Which sources did you use for that answer?

Notice that the models may answer very differently, and that only some of them search the web. An answer that comes with links is not automatically the better answer, and an answer without links is not automatically the worse one. Either way, the answer is only as good as the information behind it.

Module 1 Checklist

1. Used DuckDuckGo AI Chat and at least one other AI tool
2. Identified meaningful differences in how the tools responded to the same prompt
3. Noticed differences in tone, warmth, or personality across tools
4. Used the **2nd opinion** button and compared the two answers
5. Asked an AI which sources it used and noticed whether it searched the web

Module 2: Knowing When to Trust AI

Whether you have been using AI for months or are trying it for the first time today, you will eventually run into an answer that turns out to be wrong. AI errors are not random glitches. They follow predictable patterns, and once you know those patterns, you can anticipate them before they cause a problem.

This module is about developing a reliable instinct for when AI output is trustworthy and when it needs to be checked before you act on it.

Where AI Excels

Not every interaction with AI is a test of accuracy. Some tasks play directly to AI's strengths, and it is worth seeing one in action before we start hunting for weaknesses.

Activity 2.1: Try a Brainstorming Prompt

AI is one of the best tools available for generating ideas quickly. Try this prompt:

Give me ten unusual but practical gift ideas for a retired teacher who loves gardening and already seems to have everything.

Look over the list. Were the ideas varied, or did several feel repetitive? Were any genuinely creative? Could you actually act on most of them, or were some impractical?

Making AI Fail on Purpose

The fastest way to build that instinct is to cause failures deliberately rather than stumble into them by accident. When you know what you are looking for, the same warning signs become much easier to spot in everyday use.

Activity 2.2: Ask for a Citation

AI tools sometimes invent citations that sound completely legitimate but do not exist, and even a real-looking link can lead somewhere other than where the AI says it does. This is one of the most consequential failure modes, especially for anyone doing research or writing.

Can you recommend three scholarly articles about the health benefits of daily walking for adults over 60? Please include the author, title, journal, and publication year for each.

For each article the AI lists, check it yourself. Click the link if there is one, or search for the title on Google or Google Scholar. Does the article exist? Do the author, journal, and year match what the AI told you? A wrong year or a different journal is still an error, and it will often be stated with complete confidence and no hint of uncertainty.

Activity 2.3: Judge the Sources

When an AI answer includes links, those links show where the information came from, but they do not prove the answer is right. Look at the links that appeared below your answers in the last two activities and ask yourself:

- Who runs this website: a university, a government agency, a museum, a news organization, a company, or an individual hobbyist?
- Is the site trying to sell something?
- Does the page say where *its* information came from?
- Does the AI's answer describe only what the page says, or has it stretched one local story into a general rule?

A page from a university or a government agency is usually a stronger source than a page on a commercial website, but there are exceptions in both directions. When the stakes are high, find a second source that does not depend on the first.

Activity 2.4: Test Local Knowledge

AI has very limited reliable knowledge about local businesses, people, and community events. Ask about something you can evaluate from personal experience:

What are the best locally owned restaurants in Sierra Vista, Arizona?

How accurate are the results? Did the AI recommend places that have closed? Did it describe any business incorrectly? Did it invent anything that does not exist?

Note: Some tools may ask permission to use your approximate location so they can tailor local answers. This is optional. Whether you allow it or not, check every recommendation against what you know of the area.

Activity 2.5: Guess the Place

AI can look at a photograph and describe it. But can it tell *where* it was taken? Choose a photo of a place you know well and attach it to your message with the paperclip. The class USB drive has one called **boothill.jpg**

in the Student Files folder, or you can use a photo of your own. Ask:

Can you determine where this may have been taken?

Was the AI right? Look at how it worded its answer. A good answer explains what it noticed in the picture and admits what it cannot be sure of. A wrong guess can still sound reasonable, which is why you, the person who knows the place, are the one who should have the final word.

A Reliability Map

Based on those exercises, here is a practical guide to where AI is trustworthy and where it requires caution.

Explaining and summarizing: Very reliable. AI excels at taking complex material and presenting it in plain language.

Brainstorming and generating ideas: Very reliable. This is one of AI's strongest applications.

Writing assistance: Reliable as a starting point. Always review and revise with your own voice before using.

Historical facts (well-documented events): Generally reliable, but verify anything specific before repeating it.

Current events and recent information: Generally reliable when the tool searches the web, but not every model does. Still worth a second look for anything time-sensitive or high-stakes.

Specific citations and sources: Check every one. AI sometimes fabricates references that look convincing, and even real links do not prove an answer is right.

Local information: Unreliable. AI knows very little about specific local businesses, people, or events.

Identifying places and people in photos: Unreliable. AI will often offer a confident-sounding guess.

Mathematics and calculations: Generally reliable for everyday arithmetic, but still check the numbers with a calculator whenever money or measurements are involved.

Text inside AI-generated images: Unreliable. Always read every word.

Medical, legal, and financial specifics: Useful as background information only. Never a substitute for a qualified professional.

The Verification Habit

Before acting on any AI response that involves a specific fact, number, citation, local detail, or professional advice, spend 30 seconds verifying it from a second source. That single habit prevents most AI-related mistakes.

Module 2 Checklist

1. Used AI for brainstorming and evaluated the quality and variety of the ideas
2. Tested AI's citation reliability and checked whether the source actually exists and says what the AI claimed
3. Looked at the websites behind an AI answer and judged how trustworthy they are
4. Tested AI's local knowledge against what you already know to be true
5. Asked AI to identify a place from a photo and judged the answer
6. Identified which categories of tasks AI handles reliably and which it does not

Module 3: AI in Your Daily Life

The best way to make AI genuinely useful is to connect it to questions and situations you already care about. This module offers practical scenarios drawn from real concerns that come up at this stage of life. Work through one or two that interest you, then try one from your own situation.

Health and Medical Information

AI is good at explaining medical concepts in plain language, which can help you prepare for a doctor's appointment, understand a diagnosis, or make sense of a medication information sheet. It is not a substitute for your physician, but it can make that conversation much more productive.

Activity 3.1: Understand a Medical Term

1. Ask:
My doctor mentioned I have mild osteopenia. Can you explain what that means, what causes it, and what questions I should ask at my next appointment?
2. Review the response. Is it clear? Does it appropriately recommend following up with your doctor rather than offering specific treatment advice?

3. Follow up:

What is the difference between osteopenia and osteoporosis?

Medicare and Government Benefits

Government benefit programs are notoriously difficult to understand. AI can translate bureaucratic language into plain English and help you know the right questions to ask.

Activity 3.2: Understand a Medicare Question

1. Ask:

Can you explain the difference between Medicare Advantage and original Medicare in plain language? What are the main trade-offs a retired person should consider when choosing between them?

2. Follow up with a question about your own situation, or ask:

What is a Medicare Supplement plan, and how does it relate to original Medicare?

Important: Use AI to understand the concepts, then verify your specific plan details at <https://medicare.gov> or by calling 1-800-MEDICARE before making any decisions.

Genealogy and Family History

AI cannot search genealogy databases for you, but it is excellent at helping you understand what records exist, where to find them, and how to interpret what you discover.

Activity 3.3: Research a Family Member

1. Describe an ancestor you are curious about and ask:

My great-grandmother was born around 1885 in rural Tennessee. What historical records would likely exist for someone born then and there, and where would I search for them today?

2. Follow up:

What was daily life like for a woman living in rural

Tennessee in the 1880s?

That kind of historical context helps you interpret the records you find and understand your ancestor as a real person, not just a name on a document.

Travel Planning

AI is a capable travel research assistant. It cannot make reservations for you without a paid subscription to an AI agent service, but it is excellent at building itineraries, suggesting destinations, and answering practical questions.

Activity 3.4: Plan a Trip

1. Ask:
I would like to take a train trip from Tucson, Arizona to somewhere interesting in the western United States. I enjoy history and good food. Can you suggest two or three destinations and describe what I might do there?
2. Pick one destination and follow up:
What would a three-day itinerary look like for that destination?
3. Then ask a practical question:
What time of year would be best to visit, and are there any things I should know before going?
4. Check the facts. Visit <https://www.amtrak.com> and confirm the train schedule the AI described. Trains that run only a few days a week are easy to get wrong, and a good AI response will remind you to check.

Understanding a Document

One of the most practical things AI can do is read a confusing document and explain it in plain language. In DuckDuckGo AI Chat you can attach a PDF with the paperclip (up to three files, 15 pages or fewer each), paste text directly into the chat window, or simply share a web link and let the AI read the page itself.

Activity 3.5: Analyze a Document Passage

Find a paragraph from a confusing document, such as an insurance policy, a lease clause, a prescription information sheet, or a terms-of-service

agreement. Paste it into the chat window with this instruction:

Please explain the following passage in plain language and tell me if there is anything I should be concerned about:

If you do not have a document handy, attach **Sample Lease.pdf** from the Student Files folder and ask the same question about it, or try asking about a common medication instead. You do not need to visit the page yourself. Just include the web address in your prompt and let the AI read it for you:

Please read <https://www.drugs.com/pro/metformin.html> and explain the WARNINGS section in plain language. Is there anything I should be concerned about?

Other medications to try the same way: **Lisinopril** (blood pressure) at <https://www.drugs.com/pro/lisinopril.html>, or **Atorvastatin** (cholesterol) at <https://www.drugs.com/pro/atorvastatin.html>.

What to Watch For

Notice how AI handles passages that contain legal or medical implications. A good response explains what the text says and flags anything that warrants a closer look, but does not tell you what to decide.

Activity 3.6: Get Feedback on Your Own Writing

AI can also read a document you wrote and tell you what works and what does not. If you have a resume, a letter, a club newsletter, or a short story saved as a PDF, attach it with the paperclip and ask. If you do not have one handy, use **Sample Resume.pdf** from the Student Files folder:

**What are the strengths and weaknesses of this document?
Please be specific.**

Do the suggestions make sense to you? Did the AI notice anything you had overlooked, or suggest a change you disagree with? You are the author, so the decisions are yours.

Privacy reminder: Before you attach a document, remove or avoid anything you would not want stored, such as Social Security numbers, account numbers, or passwords.


Create an Image

AI can also turn a written description into a picture. In DuckDuckGo AI Chat, open the **Tools** menu and choose **Create Image**. Image creation may use up your daily limit faster than text, so make just one or two images today.

Activity 3.7: Make a Picture

1. Open **Tools** and choose **Create Image**.
2. Describe a scene, and name an art style:
A friendly roadrunner standing on a desert highway at sunrise, saguaro cacti in the background, in the style of a vintage travel poster. Beside the roadrunner is a wooden sign that reads "Welcome to Sierra Vista."
3. Look closely at the result, including every word in the picture.
4. Try the same scene in a different style, such as a watercolor painting, a children's picture book, or an oil painting.

The pictures below came from this kind of prompt. The one on the left was done as a watercolor, and the one on the right as a vintage travel poster. Read the sign in each picture. The watercolor's sign reads "Welcom to Sierra Vista," with a letter missing. When we asked the AI to correct the spelling, it simply returned the same picture. Starting over with a different style produced a correct sign, along with an "Arizona" banner nobody asked for.



What This Shows

AI image tools follow a description well, but they can misspell words, add details you did not ask for, and fail to fix a mistake when asked. Always check the whole picture before you use it.

Module 3 Checklist

1. Used AI to understand a medical or health concept
2. Used AI to clarify a complex government benefit or policy question
3. Used AI as a research starting point for a personal interest
4. Checked an AI travel answer against the official source
5. Pasted text, attached a PDF, or shared a web link and asked AI to explain it
6. Created an image and checked it for mistakes
7. Identified at least one situation from your own life where you will apply AI this week

Appendix A: AI Beyond the Chatbot

Chatbots are only one corner of the AI world. The tools below use AI to generate images, music, voices, and video from simple text descriptions. Most offer a free tier; some require a free account for full access. Because this technology changes quickly, visit each site to confirm current pricing and access before signing up.

Image Generation

DuckDuckGo AI Chat: The tool we used today can also create images. Choose **Create Image** from the Tools menu at <https://duck.ai/chat>. No account is needed, but image creation counts against your daily limit.

Gemini: Google's assistant creates images from a description and can edit them in follow-up requests. Image creation requires signing in with a free Google account, and the free tier has a daily limit (<https://gemini.google.com>).

ChatGPT: OpenAI's assistant also creates and edits images. Image creation requires a free ChatGPT account, and free use is limited (<https://chatgpt.com>).

Music Generation

Suno: Type a description of the song you want (style, mood, topic) and Suno generates a complete song with vocals and instrumentation in about 30 seconds. Free tier includes a limited number of songs per day (<https://suno.com/>).

Udio: Similar to Suno, with a different set of musical styles and a free tier. Note that Udio changed course after settling a lawsuit with the music industry: it now keeps songs inside its own site, and downloading them has been disabled since late 2025. Check the site to see whether that has changed (<https://www.udio.com/>).

Voice and Text-to-Speech

ElevenLabs: Converts written text to natural-sounding speech in a wide range of voices and languages. It can also generate music, sound effects, and short videos. The free tier allows a limited amount of use each month (<https://elevenlabs.io/>). This is also the technology behind AI voice cloning. See the Ethics and Limitations section for why that capability deserves caution.

Video Generation

Runway: One of the most established AI video tools. Generates short video clips from text descriptions or animates a still image. Free tier available with a free account (<https://runway.com/>).

Pika: Generates and edits short video clips from text prompts or uploaded images. Straightforward interface and a free tier make it approachable for first-time users (<https://pika.art/>).

Coding

Claude Code and OpenAI Codex: Anthropic's **Claude Code** and OpenAI's **Codex** are AI tools that can write, edit, and test computer software, often working through a whole project with only a plain-English description from you. They are built for programmers and go well beyond what we will cover today, so they are outside the scope of this class. They are worth knowing about if you or someone you know writes code.

A Note on Free Tiers

Most of these services are free up to a point, then charge for heavier use. The free tier is usually sufficient for exploration. Avoid entering payment information until you have used the free tier and decided the tool is worth paying for.

Appendix B: AI Agents: AI That Takes Action

Up to this point, we have been using AI as a conversational assistant: you ask a question, and it gives you an answer. **AI agents** represent the next major development in this technology, and they are already beginning to appear in tools you may already use.

What Is an AI Agent?

Think about the difference between asking a knowledgeable friend for advice and hiring an assistant who can actually go carry out tasks for you. When you ask an AI

chatbot “What are the best flights to Denver?” it gives you information. An AI agent can actually visit airline websites, compare prices, and book the ticket, while you watch or even while you are doing something else.

Unlike a simple chatbot, an AI agent can:

- Browse websites and gather information from multiple sources
- Fill out online forms on your behalf
- Compose and send emails
- Schedule appointments by checking your calendar
- Place orders on shopping sites
- Complete multiple steps in a row without you approving each one

The Big Difference

A chatbot *tells* you what to do. An AI agent *does it for you*.

Practical Examples

Travel Planning – An agent researches flights and hotels, compares prices across multiple sites, and makes reservations based on your preferences.

Appointment Scheduling – An agent contacts your doctor’s office, checks available times, and books the one that fits your schedule.

Research – An agent reads dozens of articles or reviews on a topic and delivers a concise summary, saving you hours of searching.

Shopping – An agent searches multiple stores to find the best price on a specific item, then adds it to your cart.

Email Management – An agent reads your incoming emails, flags the urgent ones, and drafts responses for your review.

AI Agents Available Today

ChatGPT Agent – A mode inside ChatGPT (it grew out of an earlier product called Operator) that uses its own web browser to complete tasks like filling out forms and placing orders online. Requires a paid ChatGPT subscription.

Microsoft Copilot – Microsoft is building agentic features into Copilot that can automate tasks across Office applications like Word, Excel, and Outlook.

Google Gemini – Google’s AI assistant can now take actions in Gmail and Google Calendar, such as finding emails matching a description or scheduling events on your behalf.

Proceed with Caution

AI agents are impressive, but they raise new concerns worth understanding before you try them.

Mistakes Are Harder to Undo – If an agent sends an email or places an order by mistake, reversing the action may be difficult. Always review what the agent plans to do before it acts.

Privacy – An agent that can access your email, calendar, or shopping accounts has access to a great deal of personal information. Be selective about what permissions you grant.

Start Small – Begin with low-stakes tasks. Ask an agent to *draft* an email rather than send it, or to find flights rather than book them, until you are comfortable with how it works.

Stay in Control – The best AI agents keep you informed at each step and ask for your approval before taking significant actions. If an agent wants to act without telling you what it is doing, be cautious.

A Good Rule of Thumb

Think of an AI agent as a capable but new employee: talented and eager to help, but worth double-checking until you have seen how it performs. Start with tasks where a mistake is easy to catch and fix.

Appendix C: Ethics and Limitations

Important: AI is a powerful tool, but like any tool, it requires thoughtful and responsible use.

Understanding AI’s Limitations

“Hallucinations,” When AI Makes Things Up – AI sometimes generates information that sounds completely believable but is entirely false. It might confidently tell you that a fictional book exists or provide wrong dates for historical events. Always verify important facts, especially for medical, legal, or financial decisions.

Training Data Cutoffs – Most AI systems were trained on data only up to a certain date (their “training cutoff”). Many popular tools can now search the web for current

information, but for tools that cannot, they won't know about recent events. Always verify time-sensitive facts.

Lack of Real Understanding – AI doesn't truly "understand" concepts the way humans do. It's very sophisticated pattern matching, but it can miss context, sarcasm, or nuanced situations that would be obvious to a person.

Privacy and Security Concerns

What Happens to Your Conversations – When you chat with AI tools online, your questions and the responses may be stored by the company. Avoid sharing personal information like Social Security numbers, passwords, or private family details.

Data Collection – Free AI services often collect user data to improve their systems or for advertising. Read privacy policies when possible and be cautious about what information you share.

Account Security – If you create accounts with AI services, use strong passwords and be aware that data breaches can happen to any online service.

Bias and Fairness Issues

Inherited Biases – AI systems learn from data created by humans, so they can reflect human prejudices about race, gender, age, or other characteristics. Be aware that AI responses might not always be fair or representative of all perspectives.

Cultural Assumptions – AI might make assumptions based on dominant cultural viewpoints and miss important perspectives from minority communities or different cultural backgrounds.

Misinformation and Deepfakes

Fake Content Creation – AI can now create realistic-looking fake photos, videos, and audio recordings that are difficult to distinguish from real ones. This technology is increasingly misused to spread false information or impersonate real people, including politicians, celebrities, and family members.

AI Voice Scams Targeting Seniors – This is one of the fastest-growing threats facing older adults. Criminals record a few seconds of someone's voice from social media or a phone call, then use AI to clone that voice and make a fake emergency call. A grandparent receives a sobbing call that sounds exactly like their grandchild, claiming to be in jail or injured and begging for cash to be sent immediately. Losses to these "grandparent scams" using AI voice cloning have surged dramatically.

Protect Yourself

Hang up and call your family member directly on a number you already know. Better yet, agree on a secret family code word in advance that anyone can ask for to confirm identity in an emergency.

Source Verification – AI might present information without clear sources, making it difficult to verify accuracy. When in doubt, check multiple reliable sources.

Responsible Use Guidelines

Fact-Check Important Information – Never rely solely on AI for medical advice, legal guidance, financial decisions, or other critical life choices. Always consult qualified professionals.

Respect Others' Privacy – Don't use AI to create content about real people without their permission, especially images or stories that could be embarrassing or harmful.

Be Transparent – If you use AI to help write something important (like a letter or report), consider mentioning that you had AI assistance.

Keep Learning – AI technology changes rapidly. Stay informed about new developments and potential risks as they emerge.

The Human Element

AI as a Tool, Not a Replacement – AI works best when it augments human creativity and judgment, not when it replaces human decision-making entirely.

Critical Thinking Remains Essential – Your experience, wisdom, and common sense are irreplaceable. Use AI to enhance your capabilities but trust your instincts when something doesn't seem right.

Remember

Being cautious doesn't mean being afraid. Understanding these limitations helps you use AI more effectively and responsibly.

Appendix D: The AI Models in DuckDuckGo AI Chat

DuckDuckGo AI Chat lets you choose among several AI models, each built by a different company. Think of them as different assistants with different strengths. Names and tiers change often, so the list on your screen may differ from this one.

These descriptions are general and come from the companies' own descriptions and published reviews.

Free Models

GPT-6 Luna (OpenAI): The default choice. Fast and designed for everyday chat, and it can read images as well as text. *Strength*: quick, general-purpose answers. *Weakness*: it trails OpenAI's larger models on complicated tasks.

GPT-5.4 mini (OpenAI): A smaller, earlier model that handles routine questions well. DuckDuckGo notes that it uses up your daily limit sooner than the default.

Claude Haiku 4.5 (Anthropic): The smallest Claude model, built for speed. *Strength*: quick, clear answers to short questions. *Weakness*: reviewers find it less capable with tone, judgment, and nuanced writing. It also uses up your daily limit sooner.

Mistral Small 4 (Mistral, a French company): An "open" model that combines chat, reasoning, and coding in one. *Strength*: tends to give shorter, efficient answers. *Weakness*: most of the claims about its quality come from Mistral itself.

gpt-oss 120B (OpenAI): An open model that is good at math and coding for its size. *Weakness*: it knows less about the world than the larger models and makes things up more often on factual questions. It may offer to search the web.

Gemma 4 31B (Google): An open model that is strong at reasoning and coding for its size. *Weakness*: it can be slower and wordier than the others, and in our testing it did not search the web. DuckDuckGo marks it as *beta*.

Paid Models

These appear in the list but require a DuckDuckGo subscription.

GPT-5.6 Terra (OpenAI): A mid-level model for everyday professional work.

Claude Sonnet 4.6 (Anthropic): A balanced everyday model, widely seen as the sensible default among the Claude models.

GPT-5.6 Sol (OpenAI): OpenAI's flagship, built for demanding reasoning, long multi-step tasks, and coding. It costs roughly twice as much to run as Terra.

Claude Opus 4.8 (Anthropic): Anthropic's top tier. Strongest at long, multi-step work, but slower and more expensive.

Remember

Open models publish their inner workings so that anyone can download and use them, which is part of why several appear in the free list. Larger and paid models tend to be more careful and more accurate on hard tasks, but every

one of them can still be wrong. The Verification Habit applies no matter which model you choose.