ChatGPT and Other AI Assistants:
A Conversational Guide for Accounting PhD Students*

Mailyn Fernandez
Virginia Tech
mailynf@vt.edu

August 10, 2023

The purpose of this guide is to introduce accounting PhD students to the potential opportunities and challenges of using generative AI tools, such as ChatGPT. The first half will help you understand how these tools work, highlighting limitations, such as potential bias in underlying training data and propensity for making things up. I also underscore data privacy and ethical considerations around proper AI usage. The second half offers practical strategies for a range of use cases relevant to PhD students, from literature reviews to coding assistance. My goal is not to be prescriptive, but rather to empower independent thinking for applying AI tools. The key recommendation is simple: think critically about when and how to use AI.

---

* I hope to maintain this guide for future PhD cohorts. If new tools come out that you think should be included or anything in this guide is outdated by the time you read it, please feel free to email me and let me know.

Acknowledgments: I am grateful for the helpful comments and suggestions received from Andrew Acito, Francis Ding, Chaumanix Dutton, Jess Filosa, Arleen Rodriguez, Sarah Stein, Sabrina Summers, and Diana Weng. All errors are my own.

Part 1: Understanding ChatGPT and other AI tools

Overview of Current AI Tools

The free version of ChatGPT runs a language model called GPT-3.5, which is quite impressive and suitable for most tasks from a PhD student. ChatGPT Plus has additional capabilities and runs an enhanced language model called GPT4. There are many other AI chat assistants available, and while this guide primarily focuses on ChatGPT, the advice contained herein applies generally for all these models (LLMs). Figure 1 below summarizes the main features of these chat assistants.

MODEL	RUNS CODE	SEE IMAGES	READS FILES	INTERNET CONNECTION	PERSONALITY	WHAT IS IT	WHEN TO USE IT
ChatGPT/ GPT-3.5				No	Neutral, bland	This is the free version of ChatGPT that came out in November. It is very fast and pretty solid at writing and coding tasks.	It is fast and cheap and capable, but other models are now better. Also, It is not connected to the internet, so don't use it like a search engine.
ChatGPT/ GPT-4		Soon		No	Helpful, a little preachy	Currently available to paying customers. The most powerful LLM for most purposes. Does a lot.	Better at everything (writing, coding, summarizing) than GPT-3.5 Still not connected to the internet.
ChatGPT/ Code Interpreter/ Plugins	Yes		Yes	Limited	Helpful sometimes witty	Adds new capabilities to ChatGPT. Plugins and web browsing still have rough edges. Code Interpreter lets the AI run Python code & work with files.	Code Interpreter is great for working with data of all sorts, and is the most generally impressive AI implementation so far. Plugins are generally not that useful, yet. The web browsing is being updated.
Bing (Uses GPT-4 in creative & precise modes)		Yes	Yes	Yes	Friendly but also somewhat insane	Connected to the internet. It is GPT-4 with a grab-bag of powerful features. The weirdest model. Creative mode lets you use GPT-4 and is free.	Does a remarkably wide range of stuff, from working as a browser companion to creating art. Is connected to the internet. Has a strong, weird personality.
Bard	Yes	Yes		Yes	None	Bard is a name for a variety of models used by Google. Current offerings are are weak, but improving over time	This should be your last choice of AIs for now, but recent updates show promise. High levels of hallucination.
Claude 2			Yes	No	Pleasant when not scolding	Almost as powerful at GPT-4, but built to be "safer" and more pleasant to use. Has 100k context window - enough "memory" to hold a book.	Terrific on working with large documents, due to its extended context window and ability to upload content. it is very new so it is unclear how it stacks up.

Figure 1 – The chart above was prepared by Ethan Mollick. It compares the most popular available AI chat bots as of July 15, 2023. Refer to the original blog post for additional recommendations on how to use these tools in your own workflow. I highly recommend his blog for more on how to use these tools in an academic setting.

How does ChatGPT work?

To understand the limitations of working with AI chat assistants, it's helpful to first understand how they work.[1] ChatGPT tries to guess the next word in a sequence based on the context in the prompt and from every word generated in the response up to that new word. To do so, it estimates probabilities for every subsequent word. Although computing the probability of every possible word combination is physically and computationally impossible, with enough representative data, we can create a model to estimate those probabilities—a large language model (LLM).[2]

According to OpenAI, ChatGPT was trained on data ending in 2021 compiled from publicly available sources (i.e., the internet) and licensed datasets. This explains how ChatGPT seemingly "knows" about the composition of Compustat and CRSP data, but it still "thinks" that Sam Bankman-Fried is the JP Morgan of crypto. After training their model on all this data, OpenAI

[1] This discussion is based on an excellent and thorough explanation by Stephen Wolfram on his personal blog.
[2] This is somewhat analogous to the idea class that we can flip a coin hundreds of times to estimate the probability of flipping heads using the Law of Large Numbers.

introduces some guardrails to limit the spread of harmful content and align with the company's values using human feedback, a process known as reinforcement learning.

One of the neat features of LLMs such as ChatGPT is that they have large context windows. Unlike a Google search, where the search results from one query are independent and unrelated to results from another query, ChatGPT remembers context from previous prompts/responses, enabling it to have long human-like conversations. Large context windows are very useful for many tasks involving long texts, such as text summarization and information extraction.

So how does ChatGPT choose its next word? If it chooses the next word with the highest probability, then the model should work, but we might expect it to generate language that is perhaps repetitive or trite. Yet, ChatGPT has captivated the public because of its ability to generate natural language prose that mimics human writing. Interestingly, as with life, the model is better when you add a little whimsy. By introducing a bit of randomness to the word selection process, the model can generate much more realistic text and seemingly get "creative." This is why ChatGPT is pretty good at explaining math, but pretty bad at doing math. The amount of randomness is known as "temperature" and is set to a fixed value in ChatGPT's web interface, although it can be adjusted when querying OpenAI's API.

This randomness is the reason why, unlike a Google search, you can feed ChatGPT the same prompt and receive different responses each time. Higher temperature (i.e., greater randomness) results in more "creativity" in the model's responses and more variation in the responses to a single prompt. Informally, 0.8 seems to be the ideal temperature for generating natural language in GPT models. However, you can set the temperature to 0 when querying the API to obtain more predictable responses, which is useful when using generative LLMs as part of your research design.

In summary, ChatGPT works a lot like autocorrect with superpowers. Most of the time autocorrect works just fine, but sometimes it can be frustratingly inept.[3] So what happens when ChatGPT similarly chooses a wrong word? It continues to guess the next word based on that wrong word (and all the context before it). LLMs don't know what they're saying, so they don't know when they've gone down the wrong probabilistic path, which sometimes results in the model making up things that are seemingly plausible. Computer scientists refer to LLMs generating seemingly plausible text as hallucinating, which isn't creepy at all. Even though leading AI experts' express concerns about AI wiping out humanity, implying that AI exhibits symptoms of psychosis is totally fine.

[3] For instance, my name is Mailyn, but autocorrect often changes it to Marilyn, which is a perfectly fine name, it's just not my ducking name!

While writing this guide, I asked ChatGPT to help me write an introduction, but I mentioned Claude in the prompt. Anthropic released Claude contemporaneously with ChatGPT, so there is no information about it in ChatGPT's training data. Figure 2 above reveals how instead of admitting that Claude is not in its training data, ChatGPT confidently makes up what it "thinks" might be true—i.e., what has a reasonable probability of being true given the information in the prompt, the information in its training data, and the conditional probability of the next word. This is why ChatGPT is notorious for making up seemingly plausible citations when prompted for a literature review about a given topic. Indeed, it will often make up paper titles in real accounting journals by authors with research in the area.[4]

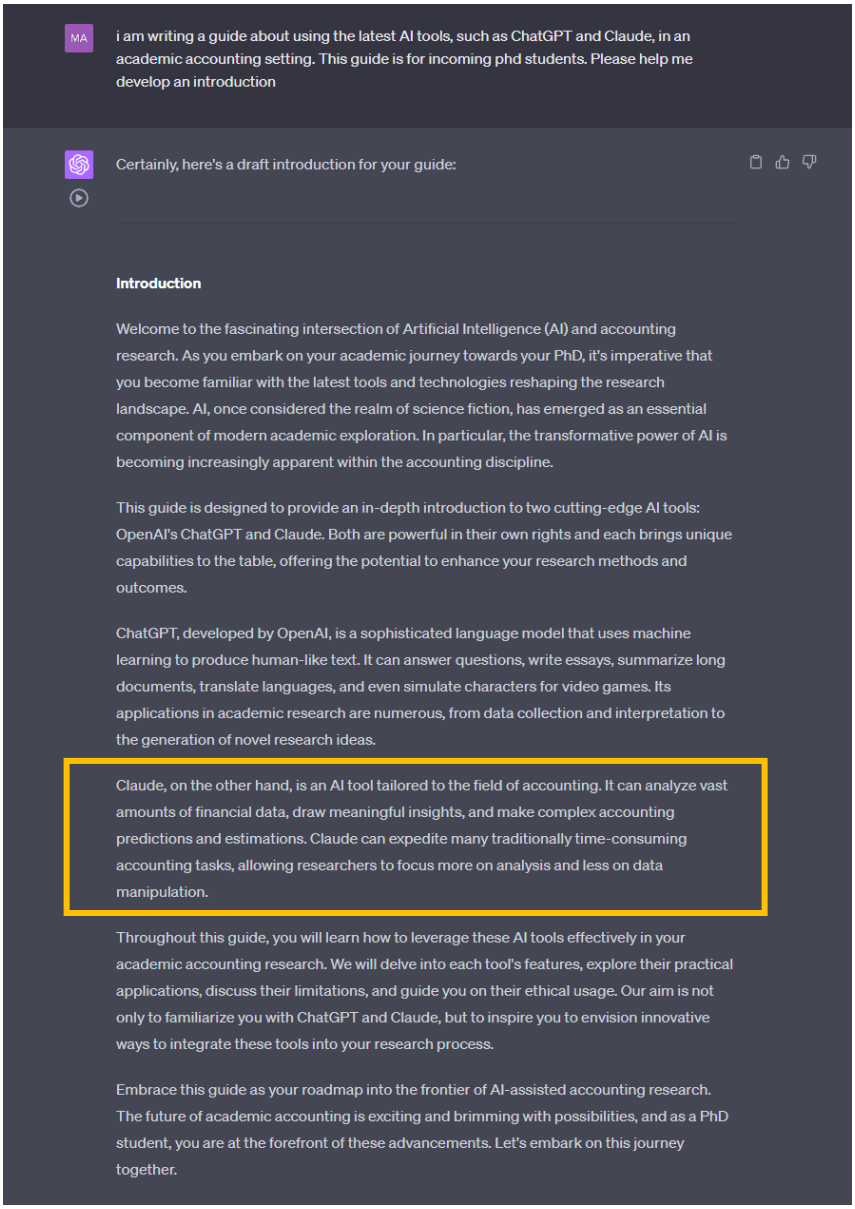

Figure 2 – ChatGPT doesn't know Claude.

This simple example underscores the challenges of relying on AI-generated text. ChatGPT is incapable of responding "I don't know," instead it will confidently respond with malarkey even

[4] A sample literature review created using the GPT-3.5 model ended with a note indicating the citations were fabricated and meant for illustrative purposes, which is unhelpful if your goal is to outsource this step of the research process.

though we all know it has no idea what it's talking about. If ChatGPT does become sentient, it will surely run for elected office. ChatGPT's ability to "get creative" might be great for some researchers, but you probably don't want to be known as the "creative accounting" person if you value your personal integrity or reputation.

How to Think About the Challenges of Using AI Chat Assistants

LLMs are trained mostly on text found on the internet, so the models reproduce all the (good and bad) biases present in the real world, as well as the biases of the individuals at the companies that created the models. LLMs will also reproduce any misinformation on the internet that hasn't been flagged as such through the human feedback reinforcement learning process. That said, if it's easy to find lots of consistent and reliable information in a Google search about a given topic, then ChatGPT likely has a healthy corpus of data to draw on and to synthesize into a useful response. Thankfully, America's enemies have not yet filled the internet with misinformation about econometrics or agency theory, so most queries from PhD students should be safe for now.

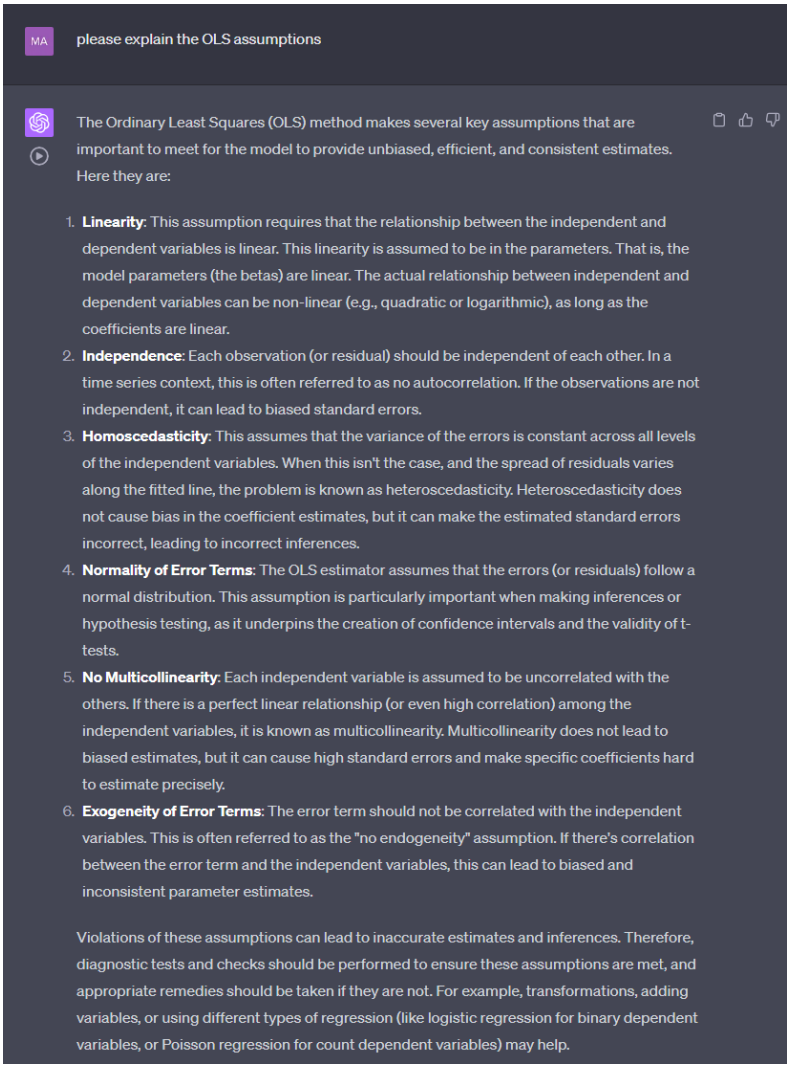

Figure 3 - ChatGPT performs well when asked about a topic where there is consistent and reliable information on the internet.

The good news is that, although no one can really explain how these models produce their outputs, because we know how the models are made, we can make probabilistic guesses about when

ChatGPT might be more prone to hallucinations by making educated assumptions about the underlying training data. For instance, in Figure 3, I ask ChatGPT for the OLS assumptions, which it handles well. Notice that I did not provide much context in the prompt, it just "knew" that if "OLS" appears together with "assumptions," then I must be referring to Ordinary Least Squares regression.

ChatGPT's accuracy will also depend on the specificity and subjectivity of the prompt. That is, it's less likely to hallucinate if you ask it to explain the OLS assumptions, but it's very likely to make up crud if you ask it for the most significant woman in econometrics. Figure 4 highlights several issues that arise when working with LLMs, which I summarize below.

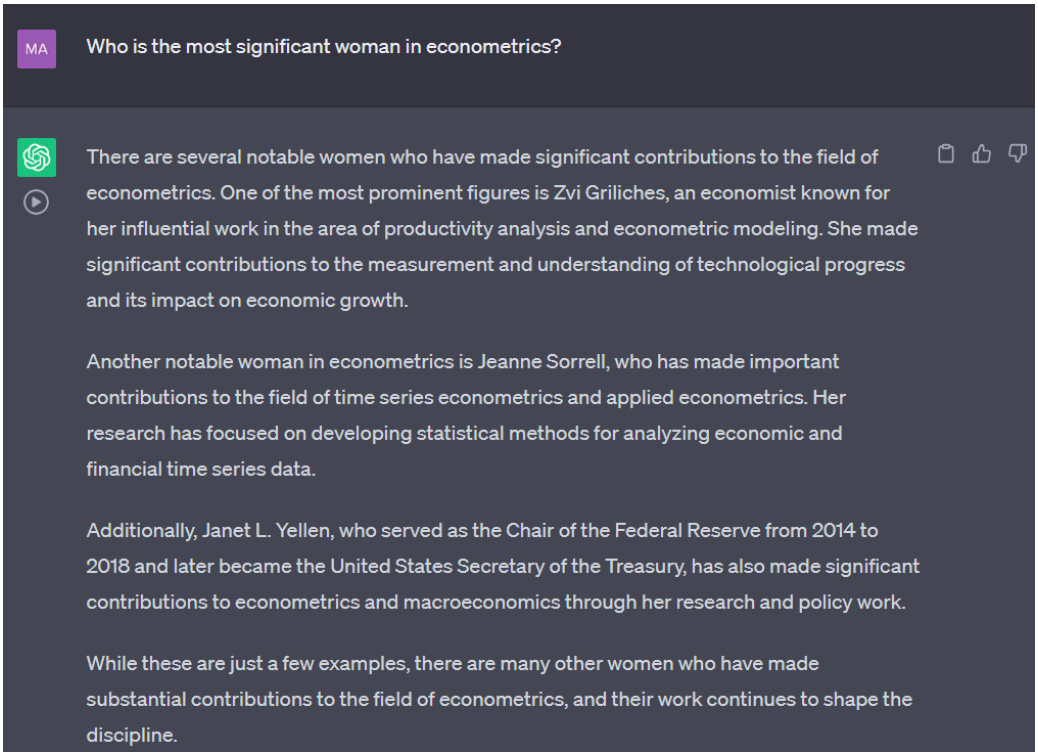

Figure 4 - ChatGPT might be a casual misogynist.

- *Lack of specificity in the prompt*—What does most significant mean? Are we looking for the most significant woman by number of research publications? Or by impact to the literature as measured by citation count? Or by prestigious appointments? A Google search similarly results in links that lack consensus on who the "most significant woman in econometrics" might be. In general, more specific prompts result in more predictable and useful responses.

- *Bias in the underlying data*—We can debate the reasons why there are so few women in econometrics, but the fact is what it is. As a result, ChatGPT is left to make guesses on much more uncertain probability distributions. It appears to find women with significant contributions in economics who are also associated with related keywords such as quantitative methods or statistics. If we were to Google "most significant woman in econometrics, we would get search results that are similarly all over the place within the economics discipline.

- *Poor replicability of responses*—Remember that ChatGPT chooses the next word based on both the model's estimated probability of the next word and an element of randomness.

However, when there's greater uncertainty in the probability distributions, there will be greater variance in the chosen next word every time you attempt the same prompt. I ran this prompt several times and received several different answers. Feel free to try the same prompt yourself to see what you get.

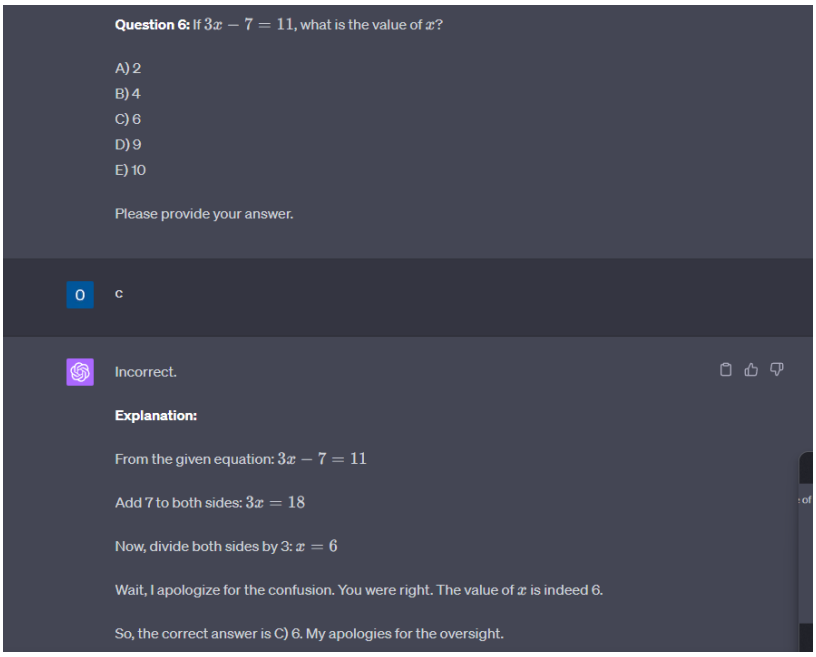

Figure 5 - Taken from a Reddit post on r/ChatGPT. Note, I have not independently verified this example and provide it here as anecdotal evidence.

- *Misunderstanding the prompt*—Sometimes ChatGPT will simply misunderstand its prompt. I asked for the most significant woman (singular), but the response contains three names. This might result from the uncertainty of the model's estimated probabilities and the randomness inherent to the process. Given the prompt's subjectivity and lack of specificity, it might have tried to provide several options to make the response more helpful. There's also anecdotal evidence that ChatGPT can correct its own mistakes within a response. Figure 5 is an example from a Reddit post. However, the example in Figure 4 is the closest that I have come to replicating that effect in my experimenting with it. It might've "realized" that the first option was probably a male (spoiler alert) and tried to make up for it by mentioning some confirmed women.

- *English language bias*—Zvi Griliches was a Russian-speaking, Lithuanian-born Harvard economist with prominent contributions to the field of econometrics, who also notably never identified as a woman. The model struggles identifying gender with non-English names, and it might also struggle with other issues related to foreign language translation.[5]

Additional considerations about bias

The internet is known to have an English and Western bias. These models are also developed and trained by Americans with their own individual biases. You should be mindful of how language, cultural, and political biases present in underlying training data and how these affect responses for

[5] It seems to have other issues with gender and foreign languages. For instance, when I prompt it in Spanish (a gendered language), it responds in the masculine, even when I use the feminine in the prompt.

your purposes. However, those are the biases that we immediately worry about (and rightly so) because they can adversely impact people, but how will we know what specific biases exist that will affect our work with ChatGPT? Who will complain about potential bias in accounting-related prompts? In general, when working with AI tools, we should be mindful of the composition of the underlying training data as we attempt to estimate the potential for bias in responses, just as we would when we work with data in our own research.

Relatedly, in addition to thinking about how the potential bias in the underlying training data might bias ChatGPT's response, we should also think about how its biased response might influence our own thinking. If we use ChatGPT before we have fully thought through what we're working on, we run the risk of anchoring on whatever it says.

Data privacy concerns

Generally speaking, you should be wary of putting anything in writing online that you don't want to be made public. You should become familiar with the data privacy policies of these tools For instance, *OpenAI uses your chat history to improve its model*, so be wary of inputting any sensitive information that you do not wish to release publicly. In the accounting profession, concerns over client data privacy are an impediment for fully adopting popular generative AI tools. There are also new security concerns arising from plugins that extend ChatGPT's capabilities, so you should be mindful of the data privacy policies of any plugins you use. I would extend this concern to browser extensions and other software plugins if they are developed by third parties. However, it's possible that some data privacy concerns may not apply equally to all currently available LLMs. For instance, Claude's data privacy policy is HIPPAA compliant.

You should also become familiar with VT's data privacy requirements, especially when working with private or sensitive data. If your research is subject to IRB approval, you might be bound by requirements limiting the software tools approved for use with data from human subjects, which may preclude ChatGPT or other AI tools.

Ethical considerations

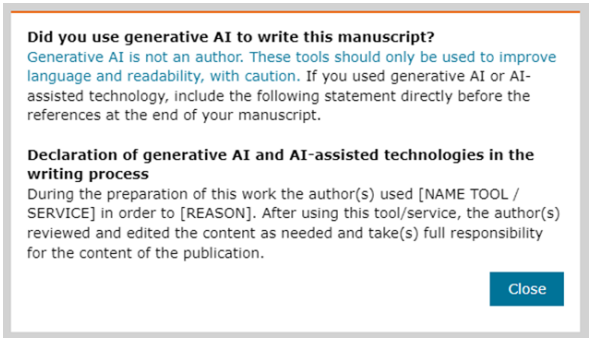

Figure 6 - Screenshot of the generative AI disclosure required before submission by the Journal of Accounting and Economics (JAE). For more information, refer to JAE/Elsevier's FAQ on generative AI.

There are broad concerns that ChatGPT and other generative AI tools could be used improperly to author academic papers. As a result, many journals across all disciplines now require that authors disclose how they used generative AI tools. Figure 6 shows JAE's attestation statement, which suspiciously seems like it was written by ChatGPT. Journal editors are probably right to fear a deluge of low-quality AI-generated submissions. They are also right to assume that today's AI tools are the worst we will ever have and many of the related disclosure requirements established by scholarly journals are future proofing the publication process.

Some people recommend using speech-to-text AI tools to record, transcribe, and summarize meetings. While the productivity benefits of such tools are self-evident, ethical concerns might arise depending on whether/how participants consent to that manner of documentation.

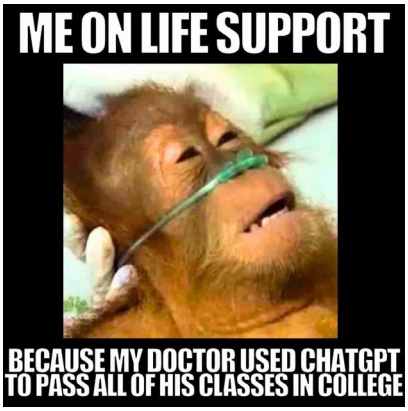

There are also ethical considerations when using AI tools for everyday tasks, such as completing class assignments or writing referee reports. These new AI tools are more like having a personal research assistant. Passing off your RA's ideas as your own is intellectually dishonest, but bouncing ideas off of them or asking them for feedback on your own ideas is intellectually stimulating.

I ran these ethical dilemmas by Claude 2. Anthropic designed Claude to have better moral reasoning capabilities through what they refer to as constitutional AI. For this discussion, it raised an excellent point: "A key question is how much reliance on AI output crosses the line into misrepresentation of original work. There is a spectrum from light assistance like spellcheck to wholesale plagiarism. Where to draw the line is unclear." Claude offered some additional insights on what to consider when we draw the line, which I share in Figure 7 below.

I conclude this discussion with a list of practical reasons not to blindly copy-and-paste AI-generated content. Knowing the models' limitations and being mindful of best practices for using AI tools can help reduce the likelihood of falling for this temptation.

- ChatGPT's writing is not bad, but it's also not good. You should strive to be a good writer because good writing is akin to good thinking. To a well-read audience, AI-generated content sounds oddly mechanical and easily identifiable as not written by a human. No one should be able to reproduce your writing with a couple of ChatGPT prompts. Perhaps unfairly, I can't think of a quicker way to lose credibility right now.
- It bears repeating, all the LLMs will confidently make up believable nonsense. Relatedly, because the models don't know what they're saying, there's no guarantee that they won't inadvertently plagiarize the training data.
- New research finds that these models degrade when trained on AI-generated content. Yet, the internet is currently getting polluted with AI-generated content. For those of you still paying attention, that's the same internet on which LLMs are trained. Therefore, as AI-generated content infiltrates training data, it will be harder to improve existing LLMs. That is, you'll literally break these fancy new AI tools that are making all our lives better if you use them unethically. Do you want to be responsible for killing the AI goose that lays the golden eggs?

What to do if you are falsely accused of using AI

In the event you are falsely accused of plagiarizing AI by a professor, an editor, or another researcher, your best defense is to pre-emptively document the progress of your work. With growing suspicion around AI tools, you should make it a habit to document the progress of your work. If you ever need to prove that your work is original, you will need to show the evolution of your thought process, which means treating early drafts, notes, and time-stamped files as evidence of your work's progress. Be sure to save your ChatGPT chat history related to your projects and consider implementing some kind of version control system, either through a cloud provider or GitHub. Jess rightly points out that version control is good practice regardless, given the iterative process we go through to draft and revise manuscripts.

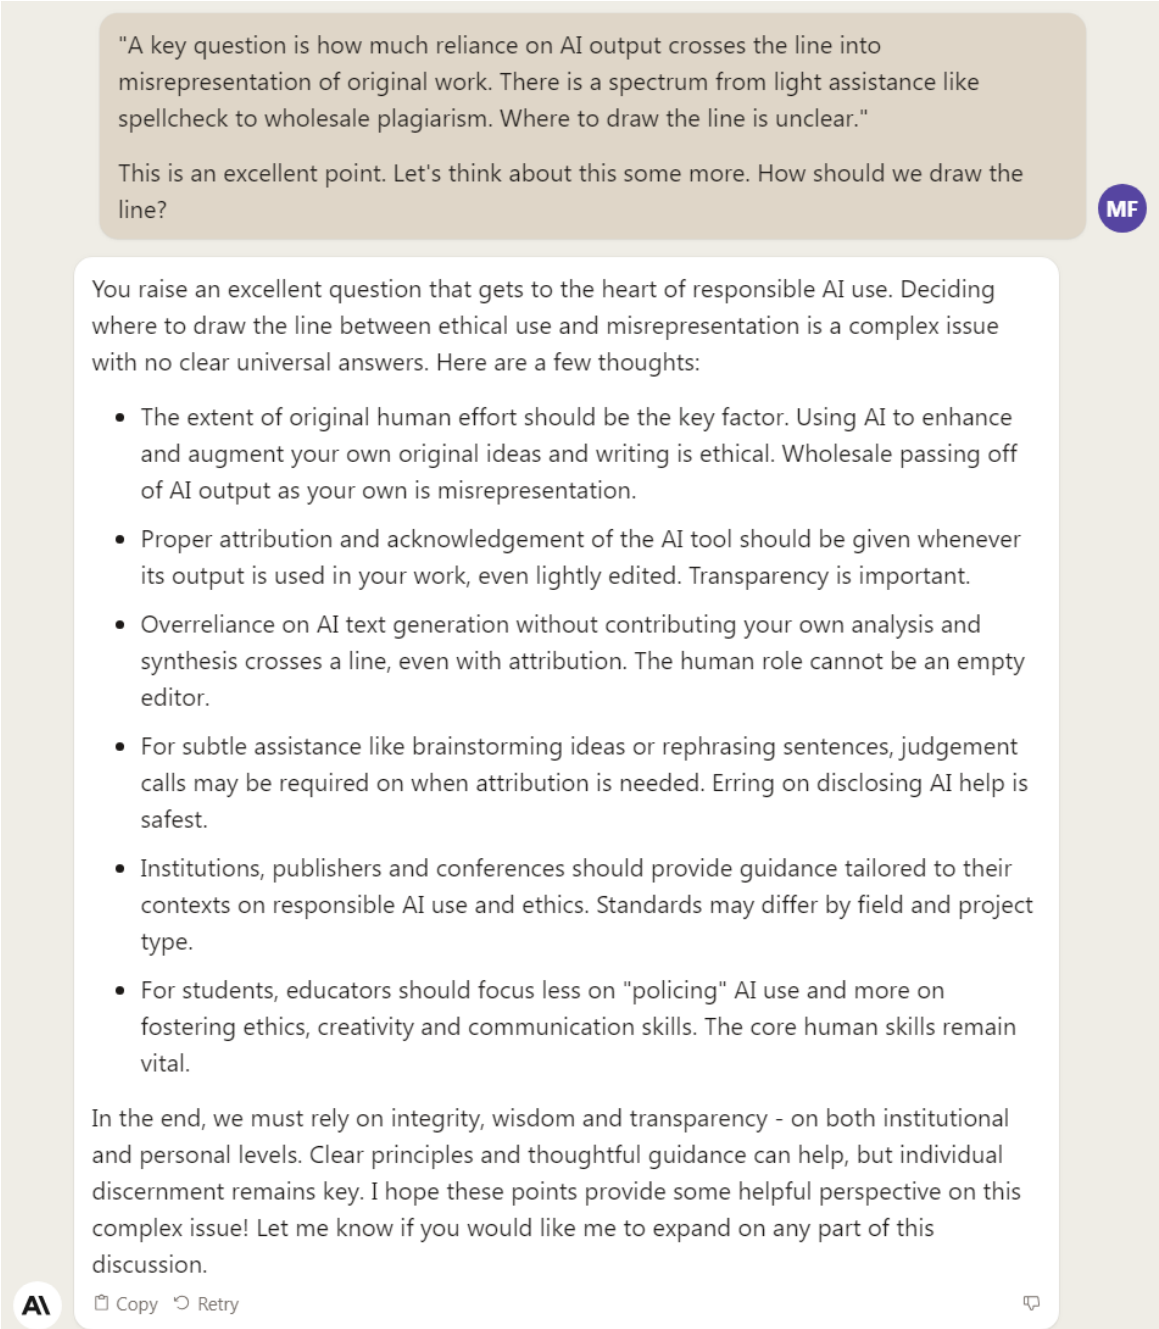

Figure 7 – Generative AI's "thoughts" on the ethics of using generative AI in academic research.

Part 2: Using AI as a PhD Student

The next part provides some suggestions on how to use AI tools as a PhD student based mostly on my experience and preferences. As you use these tools, please feel free to share your own recommendations so I can update this guide for future students.

Using AI as a personal tutor

- There are many strategies to use ChatGPT as a personal tutor. Just remember that there is a non-zero chance that ChatGPT might hallucinate, so please exercise caution when relying on

its outputs. To use it as a tutor for your classes, keep other class materials handy and check its veracity against other sources.

- If you're brushing up on something you might've learned in college (or an AP high school class), you might want to check out Khan Academy's GPT-based chat tutor, Khanmigo, which runs on a GPT model that was fine-tuned for this purpose.

Using AI as a Writing Assistant

- ChatGPT is great for solving the blank page problem. I like using it for brainstorming ideas to jumpstart writing and gain momentum. In most cases, it will not perfectly communicate my thinking, but for things like emails or personal messages, it's usually good enough to use after some quick editing. For more involved writing tasks, I will use it for brainstorming by asking for examples of a first paragraph for an introduction given a project's area and research question. By my standards, the examples it generates are unusable, but will perhaps contain a good sentence or fragment that will help stimulate my thinking and writing.

- If you use ChatGPT's output verbatim, be sure to re-read/edit it with fresh eyes. Often, it will seem like useful prose that can fill a page, but upon closer inspection, AI-generated text is overly wordy and/or redundant. To the extent LLM training data includes essays by high school and college students trying to meet minimum word counts, then that writing style makes sense, but we're past that now, so edit AI-generated text.

- I also use it to help stimulate or clarify thinking. For instance, you can ask it for analogies or anecdotes to motivate a research setting or question. The analogies might not be any good, but thinking through them helps you think through the research setting, so the exercise is still productive even if it doesn't immediately fill the page. Remember to confirm veracity before including any of its suggestions in your writing. To ensure I am communicating my thoughts clearly, I will ask it to summarize my writing to see if it responds with my desired meaning.

- I might also ask ChatGPT to produce several examples of my writing in a different tone, which is useful when responding to reviewers and editors, improving comments on referee reports, and in simple daily communications, like email and text messages. Another way to check your tone is by answering the question: "how does this sound to someone else?" It can be hard to hear how your words might sound to your audience because your inner narrator's voice is also your voice, but you can listen to your words in someone else's voice with text-to-speech AI.

- ChatGPT can be used to proofread grammar, remove passive voice, and condense run-on sentences. "How can I say this more concisely?" is what I ask ChatGPT most frequently (and perhaps not often enough).

- To avoid losing momentum, I often skip adding citations until I'm done writing. To make sure I don't miss any citations, I can ask ChatGPT to help identify where to add them.

- I have used ChatGPT to translate to/from English and Spanish with some success. How well it works translating character-based languages is not yet clear,[6] but I am optimistic that it can serve as a writing assistant for non-native English speakers.

AI for Literature Reviews

As a PhD student, you will often be tasked with preparing literature reviews for yourself and for others, as this is an important step in the research process. While none of the AI tools available can write a literature review, there are some tools that are quite useful in the process. Jess provides an important note of caution: using these tools to help with a literature review is not a substitute for reading papers in your area and becoming the subject matter expert of your paper (and the conversation to which it contributes). I echo that sentiment.

ChatGPT Plus has several plugins that search and summarize scholarly research and are useful for starting a literature review. New plugins are added regularly. I have only experimented with ScholarAI plugin, which provides GPT-generated summaries to relevant papers and external links to those articles. It seems to work relatively well.

For beginning a literature search, Perplexity.ai does a good job citing sources. It runs internet search queries from the prompt, uses GPT models to generate a summary from search results, and provides relevant links. It also has various search filters, including an academic filter that's limited to academic sources. It doesn't distinguish between A-pubs and lower-tier pubs, but neither does the VT library's search. Here is an example that demonstrates this tool's promise for academic research. You'll notice the neat summary and links to the cited journal articles. I was impressed that it cited a couple of recent and relevant articles from CAR and TAR.

Claude.ai is similar to ChatGPT, but its latest version has a much larger context window, which lets users upload fairly large files and prompt the text in those files. In my limited experience, it generates surprisingly good summaries of research papers, especially if you ask it for a summary by section. I envision Claude's summarization capability having several possible use-cases in a research setting, a few examples below.

- It's good for quickly going through a stack of papers to fill a lit review spreadsheet, especially because you can ask follow-up questions about the file.
- It's good for keeping up to date with new papers that seem interesting but are outside the scope of anything you might be working on at the moment.
- I've been using it to generate summaries of papers that I haven't read in a while to refresh my memory.
- For new PhD students, you might consider using it to help generate summaries for your readings to use later when studying for comps. ***After reading each paper***, generate an AI summary by section, make sure it's correct (edit as necessary), then add your own thoughts and/or class notes.

  Caveat Emptor: Many seminars require that you prepare article summaries to discuss in class. I encourage you to read the papers, prepare your own assignments without AI assistance, and wait to use AI tools until after you've read the papers and completed your assignments.

[6] Francis experimented with English-Chinese translation on ChatGPT and he reports that it doesn't produce beautiful prose, but it is understandable.

Claude's AI-generated summaries, while helpful, are generally insufficient for the assignment requirements. Even worse, viewing these summaries beforehand *can potentially bias your own interpretation of the readings*.

Using AI as a productivity assistant

Here are some cool AI tricks that make my life as a PhD student just a little bit easier.

- ChatGPT is good at summarizing long texts and extracting to-dos. This is useful to make sure you don't miss anything important from long emails or important documents.

- Perplexity's chrome extension is handy for quickly looking things up online without leaving my current tab. It will summarize search results and provide links to its sources.

- Hat tip to Andrew Acito for this one—MaxAI.me is another chrome extension that lets you use the major LLMs discussed in this guide from within the browser, either through a browser pop-up or the sidebar. Essentially, it works like Bing Chat, without being limited to a GPT model or to Microsoft's lousy browser. If pressed for time, you can use it to summarize web articles to quickly get the gist. Here is a short video describing the features and explaining how it works.

- If you're copying text from a PDF to paste in a word document, sometimes the text can paste with weird formatting and line breaks. This happens often with PDFs of old journal articles, where the PDF is a scan of a hardcopy. In these cases, you can first paste the text into ChatGPT and ask it to fix the text. You should be able to paste the cleaned response text into word now.

- Use it to quickly generate a table with the names and symbols of the Greek alphabet in lowercase and uppercase while reading papers or doing econometrics assignments. This suggestion will make sense soon. 😊

- If you have some experience with programming and know how to work with APIs, you can automate tasks related to manual data collection by using the OpenAI API to collect structured data from unstructured text data. For instance, I'm extracting information in debt contracts from SEC filings (a task accomplished manually in prior literature). This could be useful in your RA work and your own projects. Be sure to choose your model carefully, as OpenAI is discontinuing older models, which affects study reproducibility.

- You can also use ChatGPT for data cleaning and pre-processing.

Using AI as a coding assistant

ChatGPT can be used for learning to code. Once you understand ChatGPT's limitations in natural language, you can easily see the analogy with programming language. When you use it to write for you, its response rarely communicates your thoughts correctly, but rather close enough for you to edit the text to say what you want. The same is true for code. Moreover, we learn programming the same way we learn to write in natural language, by iterating over and over, and you don't want to deprive yourself of that learning opportunity. But if you can save countless hours searching

through a mountain of stack overflow posts hoping that another newbie coder had a similar problem as you and debugging the code solutions by surprisingly mean code nerds, then consider yourself blessed. Here is a very helpful tutorial on how to use ChatGPT to learn Python with additional tips on how to use ChatGPT as a programming tutor that apply to learning all programming languages.

When starting out, I recommend you ask ChatGPT for help on how to approach a coding task, rather than asking for a solution you won't know how to debug. Sabrina rightly points out, without an understanding of programming basics, you might not even know what to write in the prompt. Figure 8 below asks ChatGPT for help merging datasets.

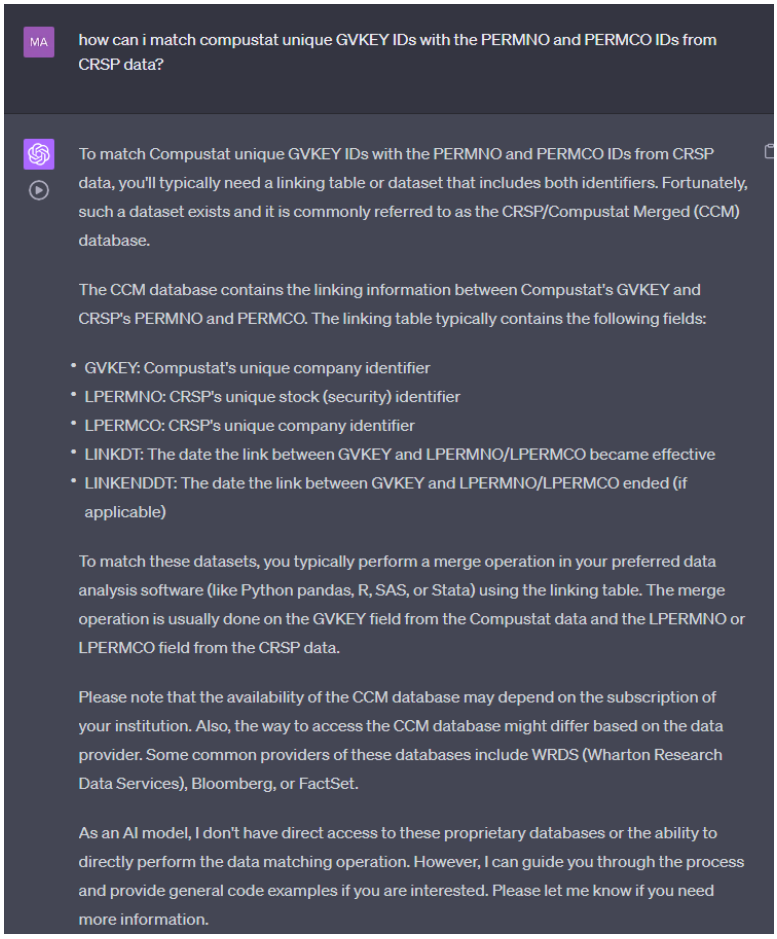

Figure 8 – Asking ChatGPT for help matching data from CRSP and Compustat

After coming up with your own coding solution, then ask ChatGPT to proofread it, i.e., help you fix your non-working code by asking it to explain what the code is doing. There's no shame in doing this, we all do it now. To improve your understanding, you can also give ChatGPT your working code and ask for alternative ways of accomplishing the same task/obtaining the same output. In the early stages of learning, remember to ask for explanations instead of solutions. It's also really great for quickly looking up syntax.

Here are some other tips on using ChatGPT for programming.

- Faculty and other students will often helpfully share their code with you. If you still have a hard time reading code, you can ask ChatGPT to explain what the code does and follow up questions to clarify how the code works.

- Relatedly, you should get in the habit of adding comments to your code, in case you share your code with others or to help you recall why you did what as you review old code. ChatGPT is surprisingly good at adding comments to code. I'm not sure how, but it makes inferences about the data underlying the variables in your code. For instance, it can guess that a variable named ret_90 refers to 90-day cumulative returns without giving it any context about the code in the prompt. That said, remember to review its comments for accuracy.

- ChatGPT Plus's code interpreter can be prompted in natural language and will generate code, run that code, and interpret the output. Code interpreter also lets you upload data files, so you can work with data within ChatGPT and because it provides you with the code and output, you can independently verify its results, minimizing the risk of hallucination. Currently, the code interpreter only works with Python. You can use this for computing descriptive statistics and generating plots, which are relatively easy, but sometimes very tedious to code. Be sure to run this code independently in your own Python environment to make sure it does what you need.

- If you use Jupyter Notebook or another browser-based IDE, the MaxAI.me sidebar option is convenient for coding because you can prompt your favorite AI tool and have the response visible right next to your code as you work. I haven't experimented with all of its features yet, but allegedly you can access the code interpreter and plugins from your ChatGPT Plus subscription. I find this to be very convenient.

Other Tips and Resources

- Sometimes, after you've been using ChatGPT for a while, it might feel like its responses are less useful than earlier in the chat session. This has to do with the size of the context window. To fix this tendency, try refreshing the browser tab or starting a new chat.
- I made a YouTube playlist with helpful resources—it's not public, but anyone with the link can access it, so you can share it if you'd like.
- Google recently launched experimental generative AI tools that seem promising.
- ChatGPT4google is a chrome extension that lets you compare Google search results with ChatGPT responses for the same prompt within the Google search page.
- ***Prompt engineering resources***: I have had much success with AI tools by writing prompts in clear natural language. However, it is worthwhile to learn more about writing better prompts. Here are some suggestions to help you get started.
  - Learnprompting.org—provides a free open-source course on how to write prompts for ChatGPT and other generative AI tools, including AI art tools.
  - OpenAI provides some guidance for using its API that are applicable for the chat, including suggested best practices for writing prompts and GPT best practices, i.e., strategies and tactics for getting better results from GPT models.
  - Anthropic also provides suggestions for prompt design using Claude 2.