



Large Language Models and Generative AI, Oh My!

Michael Zyda , University of Southern California

Artificial intelligence (AI) has a miserable reputation regarding promises made and results delivered. Last century, there were many AI winters. The hype surrounding large language models and generative AI could join this trend, we do not hope.

INTRODUCTION

Artificial intelligence (AI) winters are periods of low to no funding by government agencies and the venture community. AI winters usually come when grand promises are made by researchers and startup founders to interest government funding agencies and investors. Grand

promises are made, funding is provided, and results delivered are less than stellar. New funding and interest dry up and blow away. Rinse and repeat at a later date. Figure 1 shows the various failures that led to AI winters in the last third of last century. That material is from Wikipedia¹ and redrawn here for your edification. The source¹ provides the right amount of detail should you need more.

THE AI HYPE CYCLE

The AI hype cycle is shown in Figure 2. It starts with hype, which is defined as extravagant or intensive publicity or promotion, according to the online Oxford Languages dictionary. In our case, with respect to AI, this starts with big researcher promises followed by press amplification of those promises. From industry, this all started with hiring people with the title “Evangelist,” which was the warning sign to us all that something was not quite right. Out of this, the AI hype cycle has historically ended with unrealistic

Digital Object Identifier 10.1109/MC.2024.3350290
Date of current version: 1 March 2024

expectations being held and disappointments to funders (research sponsors, venture capitalists, and corporate investments), with all of this resulting in funding cuts and researchers and investors abandoning the area. So, all of us on the sidelines are wondering, is there a large

language model (LLM) or generative AI winter roaring toward us, or is something useful going to come out of all the noise?

WHAT ARE LLMS?

“Large language models (LLMs) are deep learning models that are pre-

trained on massive sets of data, some as large as 50 billion web pages. A core part of the LLM is a transformer model. A transformer model is a set of neural networks that consist of an encoder and decoder with self-attention capabilities, meaning the models are capable of unsupervised training from the

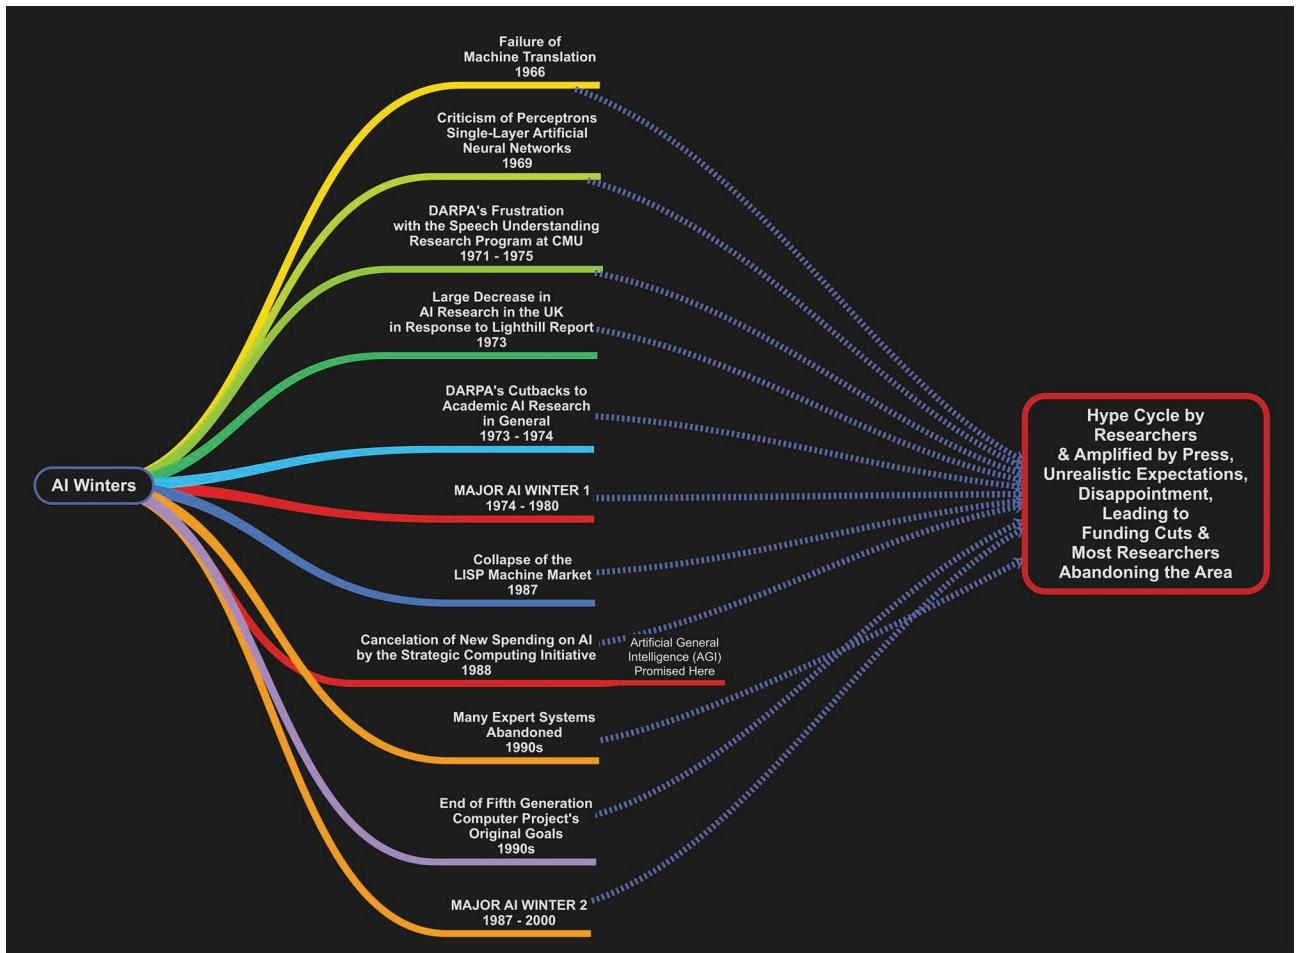


FIGURE 1. Early examples of periods of AI winters.¹




FIGURE 2. The hype cycle.

data fed them. The transformer models learn to understand the sentence, paragraph, and complete article of the data presented. Transformers process entire sequences of data in parallel, meaning their computation can be performed on a GPU (graphics processing unit) dramatically reducing the training time for the LLM.”^{2,3,4} Note that the above

paragraph is a rewrite and compression of the cited sources.^{2,3,4}

WHAT IS GENERATIVE AI?

Fully pretrained LLMs are called foundation models. Foundation models can then be specialized (built on top of) to create generative AI applications, applications that can generate new text,

new images, new audio, language interpretations, and other new synthetic items.⁵ The generative AI applications, usually shortened to just generative AI, that people have been most excited about are ChatGPT, DALL-E, and Midjourney.

ChatGPT lets you provide it text, and it will give you new text in response. For



FIGURE 3. A turkey at the beach in Carmel by the Sea.



FIGURE 5. A turkey playing the accordion.



FIGURE 4. A turkey cooking dinner.



FIGURE 6. A turkey wearing a Mickey Mouse hat.

example, almost all my friends that tried ChatGPT 3.5 first asked it to write their personal biography. When I did this, it was sort of accurate but had me dead at the end of the written bio. There were several mistakes in the bio, and there is no way anyone could ever really use the generated text. This changed with ChatGPT 4.0, and I did the same test and

got a reasonably excellent generated bio back: I only modified one sentence, and that was more for style than anything factual. In a previous “Games” column,⁶ I tried to hand ChatGPT 3.5 a full screenplay, which it rejected, and had to settle for just a single scene from my favorite noir film, *The Long Goodbye*. ChatGPT 3.5 gave me a new scene that seemed

pretty reasonable for a scene in a noir film not yet made, but I did not provide ChatGPT 3.5 any guidance as to how that scene fit into the larger story of the complete future film.

DALL-E⁷ and Midjourney⁸ let you provide their generative AI applications text that describes an image you would like created. I haven’t tried



FIGURE 7. A turkey wearing a space helmet.



FIGURE 9. A turkey wearing running shoes.



FIGURE 8. A turkey wearing a tool belt.



FIGURE 10. A turkey working in a bank.

DALL-E, but I have spent a lot of time playing with Midjourney. Let's talk about that experience in the next section.

A MIDJOURNEY THANKSGIVING

One of the things to know about me is that when my wife decides to have a large dinner party, I get tasked with creating place cards for each person attending. I personalize these cards with respect to each person and with respect to the chosen theme for the dinner party. This year's theme was "A Midjourney Thanksgiving." Here is my list of very short texts provided to Midjourney:

- › a turkey at the beach in Carmel by the Sea
- › a turkey cooking dinner
- › a turkey playing an accordion
- › a turkey wearing a Mickey Mouse hat
- › a turkey wearing a space helmet
- › a turkey wearing a tool belt
- › a turkey with running shoes
- › a turkey working in a bank
- › happy Thanksgiving turkey apocalypse.

The pictures Midjourney created are shown in Figures 3–11.

So, this is the artwork for each place card, and I would paste that onto a card individualized with the name of each person using Photoshop. So, we are just talking about Midjourney, so the individuals are not that important except that the topic chosen for each image was turkey plus something to do with their occupation or life. I did not tell Midjourney to use a particular art style, yet it seemed to create each of these with a similar style! Midjourney did not provide me a reference or credit for that art style.

Now, we know that Midjourney's foundation model was trained by scraping imagery and text off of the Internet. And I did "a turkey wearing a Mickey Mouse hat" test to see if it would give me an image that would make Disney happy, and the turkey pics seem OK, but

it gave me two Donald Duck pics: Donald Duck is a cartoon character created by the Walt Disney Company! So, we have failed the Walt Disney test. How can we use generative AI without some way of assuring us and our attorneys

generative AI architecture doesn't just use a training set built by scraping the Internet but rather is provided a tool that can scrape the internal art archives of the particular client. This means if Walt Disney wants to use

Once we have all of the copyright and trademark issues resolved for generative AI, there is then the potential for a bright future in entertainment, especially for the games and film industries.

that we are not going to get sued for copyright infringement? This is already an issue as the *New York Times* sues Microsoft and ChatGPT maker OpenAI.⁹ Apple is taking a different path with exploring AI training data licensing with a variety of news publishers.¹⁰

So, what maybe really needs to be done is the user of a particular

generative AI to create new characters, then the training set built for the generative AI tool is made by scraping the art archives of Disney's internal historical art productions instead of the Internet. I am not sure that is an option offered by anyone right now but maybe soon as issues of trademark and copyright become familiar to the tech industry.



FIGURE 11. Happy Thanksgiving turkey apocalypse.

COMMENTS?

If you have comments about this article, or topics or references I should have cited or you want to rant back to me on why what I say is nonsense, I want to hear. Every time we finish one of these columns, and it goes to print, what I'm going to do is get it up online and maybe point to it at my Facebook (mikezyda) and my LinkedIn (mikezyda) pages so that I can receive comments from you. Maybe we'll react to some of those comments in future columns or online to enlighten you in real time! This is the "Games" column. You have a wonderful day.

GENERATIVE AI AND THE GAMES INDUSTRY


Once we have all of the copyright and trademark issues resolved for generative AI, there is then the potential for a bright future in entertainment, especially for the games and film industries. In the games and entertainment industry, there is always a demand for concept art, and the technology embedded in generative AI looks quite interesting for rapidly producing concepts in near-real time while the development crew is still sitting around the discussion table. With respect to things beyond 2D concepts, there are generative AI startups focused on creating 3D assets and worlds rapidly.¹¹ Generative AI has a great future for game characters and nonplayer characters as well.¹¹ We will see.

EVERYONE WANTS TO INVEST IN A GENERATIVE AI SUPERSTAR STARTUP

The number of friends of mine that have asked if I know anyone at Anthropic is

quite large. Everyone seems to want to invest in a generative AI startup, and the word on the street is that is the one. Who knows? There is a large list of generative AI startups here.¹¹ If some of the fundamental issues of copyright and trademark can be rectified, then maybe some of these startups will sail past the culling soon to come. And maybe, if all the beans line up, we will mostly miss another AI winter...

ACKNOWLEDGMENT

I acknowledge the contributions of my many friends in the games and computing industries, especially those who have been in AI since near the beginning...

REFERENCES

1. "AI winter." Wikipedia. Accessed: Dec. 25, 2023. [Online]. Available: https://en.wikipedia.org/wiki/AI_winter
2. "What are large language models (LLM)?" Amazon. Accessed: Dec. 28, 2023. [Online]. Available: [https://aws.amazon.com/what-is/large-language-model/#:~:text=Large%20language%20models%20\(LLM\)%20are,decoder%20with%20self%2Dattention%20capabilities](https://aws.amazon.com/what-is/large-language-model/#:~:text=Large%20language%20models%20(LLM)%20are,decoder%20with%20self%2Dattention%20capabilities)
3. "Generative AI exists because of the transformer," *Financial Times*, Sep. 2023. Accessed: Dec. 28, 2023. [Online]. Available: <https://ig.ft.com/generative-ai/#:~:text=Transformers%20process%20an%20entire%20sequence,or%20generate%20—%20text%20more%20accurately>
4. J. Uszkoreit. "Transformer: A novel neural network architecture for language understanding." Google Research Blog. Accessed: Dec. 28, 2023. [Online]. Available: <https://blog.research.google/2017/08/transformer-novel-neural-network.html>
5. "Generative AI." Accenture. Accessed: Dec. 28, 2023. [Online]. Available: https://www.accenture.com/us-en/insights/generative-ai?c=acn_glb_generativeai-lagoogle_13770004&n=psgs_0623&gclid=Cj0KCQiA1rSsBhDHARIsANB4EJZcm1xXZdUjHrcieDnW12Y1Xwu94k_Veg_fuGeOXLrDiH1z6_LzwwIaAnJIEALw_wcB&glsrsrc=aw.ds
6. M. Zyda, "Hubris—From the metaverse to ChatGPT to TikTok's time of tribulation," *Computer*, vol. 56, no. 7, pp. 131-136, Jul. 2023, doi: 10.1109/MC.2023.3267986.
7. "DALL-E." OpenAI. Accessed: Dec. 28, 2023. [Online]. Available: <https://openai.com>
8. Midjourney. Accessed: Dec. 28, 2023. [Online]. Available: <https://midjourney.com>
9. R. Browne, "New York Times Sues Microsoft, ChatGPT maker OpenAI over copyright infringement," *CNBC*, Dec. 2023. Accessed: Dec. 28, 2023. [Online]. Available: https://www.cnbcm.com/2023/12/27/new-york-times-sues-microsoft-chatgpt-maker-openai-over-copyright-infringement.html?__source=sharebar%7Cfacebook&par=sharebar&fbclid=IwAR2mn-MxwwNWP02xPsFZ1BEqmfGCxFC0tfzG1j1_SIMy1fdOF96qbussy1I
10. B. Mullin and T. Mickle, "Apple explores A.I. deals with news publishers," *NY Times*, Dec. 2023. Accessed: Dec. 28, 2023. [Online]. Available: <https://www.nytimes.com/2023/12/22/technology/apple-ai-news-publishers.html>
11. "Generative AI startups." Dealroom.co. Accessed: Dec. 28, 2023. [Online]. Available: <https://app.dealroom.co/lists/33530>

MICHAEL ZYDA is the founding director of the Computer Science Games Program and a professor emeritus of engineering practice in the Department of Computer Science, University of Southern California, Los Angeles, CA 90089 USA. Contact him at zyda@mikezyda.com.