Citi GPS: Global Perspectives & Solutions

# GENERATIVE AI

## ChatGPT and Search

With the announcement of OpenAI's ChatGPT being integrated into both search engines and the broader web, we look at what generative AI is and how models like ChatGPT could be transformational in how we search for things on the web, use information, and communicate with each other. Generative AI has the potential to change the business model of search and how we access content on the web.


February 2023

As our premier thought leadership product, Citi GPS is designed to help readers navigate the most demanding challenges and greatest opportunities of the 21st century. We access the best elements of our global conversations with senior Citi professionals, academics, and corporate leaders to anticipate themes and trends in today's fast-changing and interconnected world. This is not a research report and does not constitute advice on investments or a solicitations to buy or sell any financial instruments. For more information on Citi GPS, please visit our website at www.citi.com/citigps.

## Primary Authors


**Tyler Radke**
Global Software Analyst
Citi Research

+1-415-951-1660 | tyler.radke@citi.com

**Ronald Josey**
U.S. Internet Analyst
Citi Research

+1-212-816-4545 | ronald.josey@citi.com

**Atif Malik**
U.S. Semiconductor & Semiconductor Capital Equipment Analyst
Citi Research

+1-415-951-1892 | atif.malik@citi.com

**Alicia Yap, CFA**
Head of Pan-Asia Internet Research
Citi Research

+852-2501-2773 | alicia.yap@citi.com

**Amit B Harchandani**
Head of EMEA Technology Research
Citi Research

+44-20-7986-4246 | amit.harchandani@citi.com

**Robert Garlick**
Managing Director
Citi Global Insights

+44-20-7986-3547 | robert.j.garlick@citi.com

**Tahmid Quddus Islam**
Senior Associate
Citi Global Insights

tahmid.quddus.islam@citi.com

**Kathleen Boyle, CFA**
Managing Editor,
Citi GPS

+1-212-816-3608 | kathleen.boyle@citi.com

## What is Generative AI?

Generative Artificial Intelligence (AI) aims to understand and predict information from a particular data set. It is important to know that generative AI is not totally new. It is already used in applications like email via smart compose, which allows the email program to finish a sentence started by a user. In many ways it is an existing tool that has only recently started to ramp up.

### Generative AI at an Inflection Point

Deep learning and predictive AI have also been in existence for some time, however, recently there has been an incredible increase in model size and complexity. Large Language Models (LLMs) exist today with hundreds of gigabytes that can analyze huge amounts of data sets, although this analysis or "training" takes a lot of computing power. The increase in model size has been made possible through improvements in computing technique including Central Processing Units (CPUs) and cloud computing, which allow customers to use thousands of Graphics Processing Units (GPUs) from the cloud, as well as skyrocketing amounts of available data. Creators of the models have also made them more "human friendly" as they launch public applications, thereby making them more accessible.

### Why are Transformers an Inflection Point for LLMs?

Transformers are deep learning models that use self-attention mechanisms to weight the significance of each part of a given input data. Their use in LLMs launched a chain of development in Natural Language Processing (NLP) — a branch of AI aimed at aiding computers in understanding natural human language. The transformer model, when used for NLP, can more efficiently train AI GPUs, thus significantly driving down the costs of training versus alternative models. As the CEO of NVIDIA, Jensen Huang, put it in 2022: "Transformers made self-supervised learning possible, and AI jumped to warp speed."[1] For a more in-depth explanation of LLMs and Transformers, see commentary from Citi Global Insights on page 11.

### OpenAI and ChatGPT

OpenAI started as a research lab in 2015 and is the AI research and deployment company behind three generative AI models — ChatGPT, Codex, and DALL-E. These models are trained to understand the structure of human language to create text, code, and image content, as well as new types of data/insights from a training set. The release of the models has become an inflection point in generative AI due to improvements in compute, data availability, and public ability to test and further refine the model. The third iteration of ChatGPT (Generated Pre-trained Transformers or GPT-3) was launched in November 2022 as a human-like AI platform capable of solving/answering prompts. What is different about ChatGPT for search is it provides a conversational style response to an inquiry versus links to suggested sites. Since its launch, it has become the fastest growing consumer application in history, accumulating 100 million Monthly Average Users (MAUs) in January 2023. For context, the prior record keepers for fastest growing application are TikTok at 9 months and Instagram at 2.5 years.[2]

[1] GTC 2022 Spring Keynote with NVIDEO CEO Jensen Huang.

[2] Dylan Patel and Afzal Ahmad, "The Inference Cost of Search Disruption — Large Language Model Cost Analysis," SemiAnalysis, February 9, 2023.

© 2023 Citigroup

The ChatGPT announced in February with GPT-3.5 technology can use up to 175 billion parameters, representing more than 100x the number of parameters used in GPT-2, which only utilized around 1.5 billion. The increase in parameters comes with a significant uptick in computing power, but also gives the newer model the ability to perform tasks it was not trained on. While not exempt from incomplete answers, inconsistent feedback, or biased behavior, ChatGPS is a promising, user-friendly application of LLMs. Looking ahead, ChatGPT-4 is forecast to use 100 trillion parameters, showing that progress in LLM technology is potentially ramping up at an exponential pace.

## Generative AI Market Opportunity

AI is embedded into every layer of the technology stack. Generative AI is category of AI that is not only trained to recognize patterns and make predictions, but generates new outputs and novel insights from the datasets it was trained on. To size the market opportunity in the category, Citi Research took to the conversational AI market for similar growth comparisons, although noting they are not synonymous. Conversational AI is technology that mimics human communication and is not necessarily run via an LLM, though chatbots (and upgrades of traditional conversational AI) are more a natural initial use case for generative AI.

While there are broad ranging definitions and sizing estimates for the general AI market, Citi Research believes IDC's Global Conversational AI Tools and Technologies market forecast may be among the most relevant, given its sizing lines up with some of Citi Research company analysis and the high growth rates (though off a reasonable base). According to IDC reports, the conversational AI market is expected to grow at a 37% compound annual growth rate (CAGR) from $3.3 billion in 2021 to just over $16 billion in 2026.[3] Within that, they forecast the public cloud component will grow at a 52% CAGR versus only 19% growth for on-premise solutions. Notably, IDC made this forecast in July 2022, four months before the launch of ChatGPT, so this could be conservative and underestimate the future growth potential.

With the significant growth expected in public cloud versus on-premise conversational AI use, as well as high profile use cases such as ChatGPT, which require large amounts of computing power, Citi Research believes the maturity and scalability of the cloud market could be validated.

[3] Hayley Sutherland and David Schubmehl, "Worldwide Conversational AI Tools and Technologies Forecast, 2022-2026," IDC, July 2022/

© 2023 Citigroup

Figure 1. Worldwide Artificial Intelligence Software Revenue by Deployment, 2021-26 ($mn)

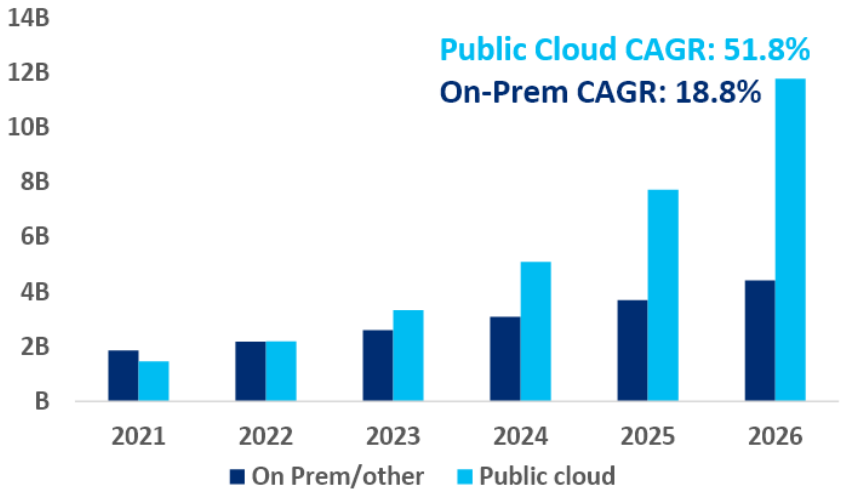


Source: IDC, Citi Research

## Key Questions

- **Whether generative AI offerings will shift search share:** While it is still early in the transition and consumers tend to be sticky (meaning changing their habits will take time), the roll-out of generative AI offerings could lead to market share shifts in search. However, it maybe some time before this shift is evident. Every one percentage point change in search share equates to around $2 billion of search revenue.

- **Whether fast followers compete in generative AI:** Despite the lead of early adopters of generative AI in search and browsers, most major players in the space are working on generative AI and LLM offerings. This means the landscape could change significantly over the near to medium term as new offerings come to the market. In the technology sector, there are multiple examples of fast followers disrupting first movers in areas such as email and short reels. A higher focus on user experience is a factor in fast follower success.

- **Potential costs to scale chat-based search:** Generative AI is expensive, both in terms of development and service. Queries using generative AI are several cents more expensive than traditional search results.

- **Reliability and responsibility of result content:** The reliability of ChatGPT results is an open question. Since it is a pre-trained model, the context/content is highly dependent on the input and the refinement of the model. For earlier iterations of ChatGPT, this meant the output was only as current as the data set the LLM was trained on. However, newer offerings pair web searches for the latest data with output from the training set. There could be potential risks/liabilities though if safety issues arise in the future, including the risk of misleading answers. This differs from traditional search, which directs users to the original source with limited liability on the accuracy of the content source.

- **Potential regulatory implications for chat-based search:** As search engine results pages give more relevant answers due to generative AI, users may spend more time on those pages. This could possible trigger “gatekeeper” issues under the EU’s Digital Markets Act, which is expected to take effect in May 2023.

© 2023 Citigroup

## What Has Been Announced

In February 2023, Microsoft announced it was integrating a new generative AI model — Open AI's ChatGPT, which was based on GPT-3.5 and optimized for search — within its Bing search engine and Edge browser. This integration would allow (1) the creation of content based on relevant/personalized results, (2) complete answers to questions summarized with links across the web, (3) chat for follow up and more specific queries, and (4) expanded search with a 1,000-character text box. In addition to creating a better user experience, Microsoft also noted that search relevancy increased when applying the ChatGPT model to its search engine. This is important as increased search relevancy can help address the 40%+ share of searches that fail to immediately return an answer.

**Figure 2. Bing's New Chat Tab for Open-Ended Questions Brings Search Results from Multiple Sources**

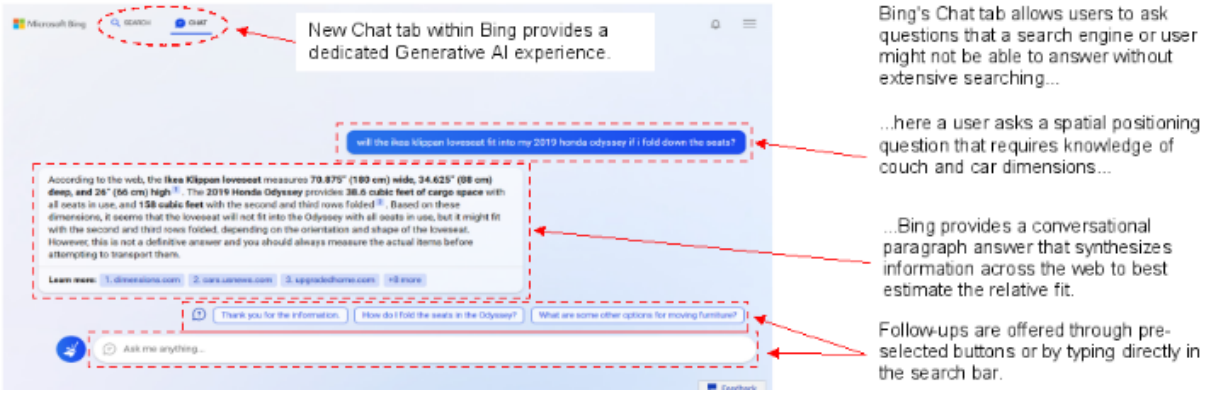


Source: Company Reports, Citi GPS

**Figure 3. Example of New Edge Sidebar with Generative AI Functionality to Create a LinkedIn Post**

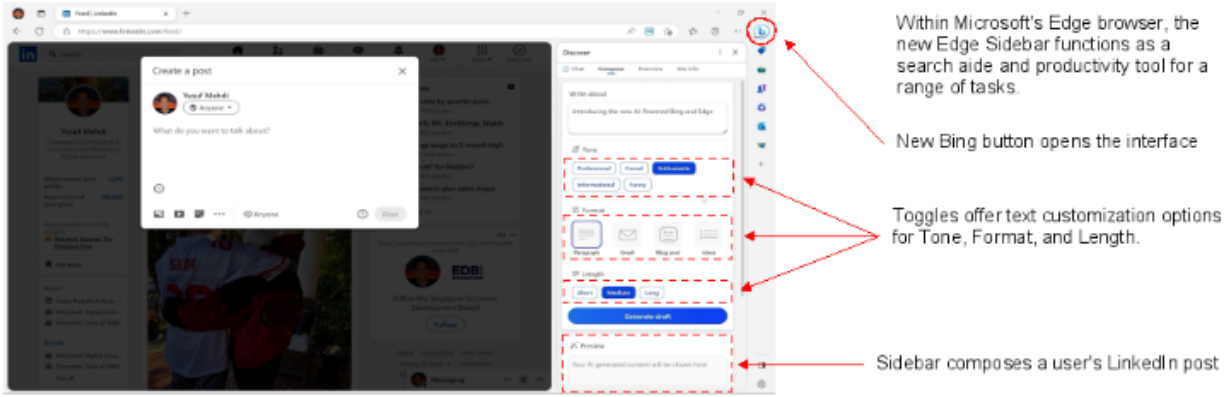


Source: Company Reports, Citi GPS

In early February, Google unveiled Bard, a generative AI tool which is its first conversation-based AI model. Bard is powered by a "lightweight" LaMDA model trained on 137 billion parameters that currently has limited features and functionality. Bard is a publicly facing large language model (LLM) based on several underlying AI tools including (1) LaMDA, Language Model for Dialogue Applications, which focuses on chat/conversation-based models; (2) PaLM, or Pathways Language Model, which is a more advanced multimodal successor that can cohesively process text, audio, and images; and (3) Sparrow, which emphasizes safe dialogues and is based on Google's original transformer research in 2017. Transformer still serves as the underlying basis for may LLMs like ChatGPT.

© 2023 Citigroup

**Figure 4. New Bard AI Interface for User Queries**

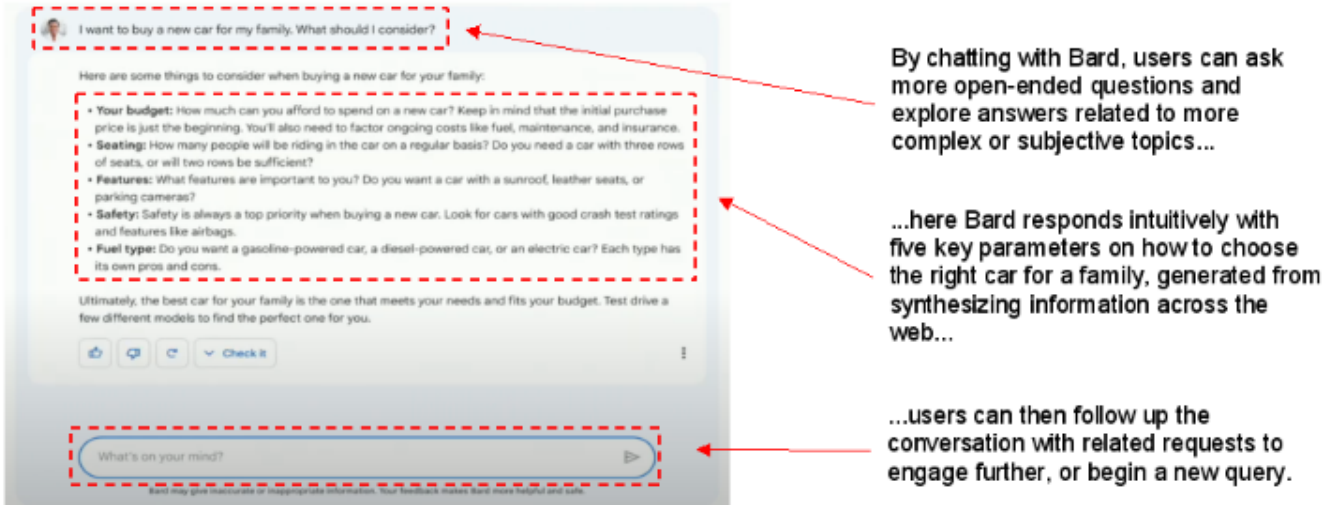


Source: Company Reports, Citi GPS

**Figure 5. Generative AI Integrated into Google Search**

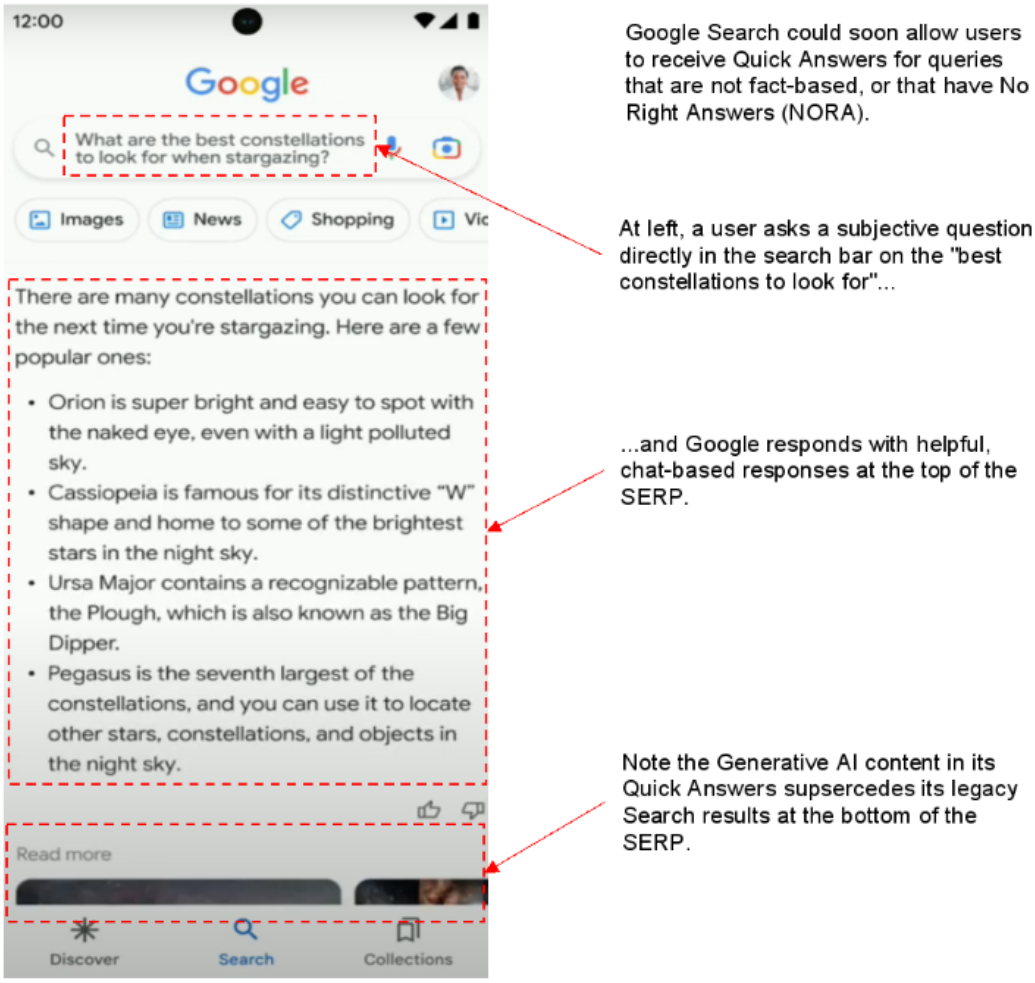


Source: Citi GPS, Company Reports

© 2023 Citigroup

## China

Given the strong excitement for ChatGPT and the release of Bard, Reuters and other news media have reported that Chinese internet companies are developing similar products to incorporate ChatGPT-like features into their services and products.

- February 7: Reuters reported that Baidu disclosed it would complete internal testing of a ChatGPT-style project called “ERNIE Bot” in March. Per the news report, Baidu aims to make the service available as a stand-alone application and gradually merge it into its search engine, incorporating the results when users perform search queries.
- February 8: Reuters reported, quoting 21st Century Herald newspaper, that Alibaba is also developing a ChatGPT-style AI tool and is currently undergoing internal testing.
- February 8: Reuters also reported that JD.Com intends to integrate ChatGPT methods and technical points into its product services.
- February 9: Mainland China media Cailian Press reported that NetEase’s education unit, Youdao, is researching to incorporate AI-generated content into its education service.

### ERNIE Model

ERNIE (or Wenxin in Chinese), which stands for, “Enhanced Representation through kNowledge IntEgration,” is a natural language processing, deep-learning model introduced by Baidu in March 2019. ERNIE is a continual pre-training framework that builds and learns incrementally by pre-training tasks through sequential multi-task learning.

**Figure 6. ERNIE Milestones**

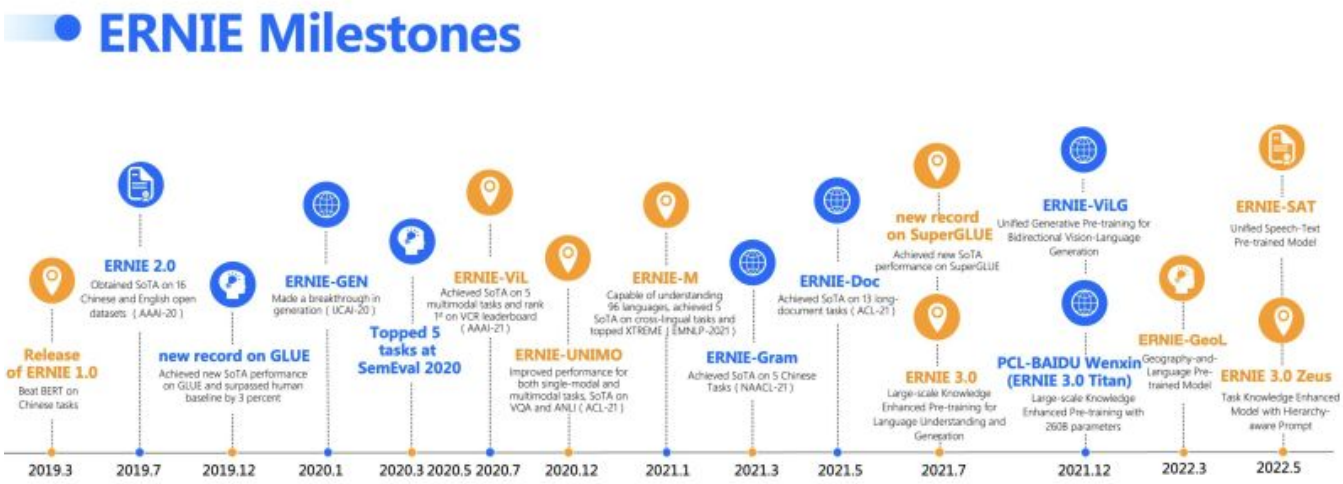


Source: Company Reports on Github, Citi GPS

© 2023 Citigroup

Multiple iterations of the ERNIE model have been developed since 2019:

- **ERNIE 3.0 TITAN:** a pre-training language model with 260 billion parameters that was trained on massive unstructured data and a knowledge graph that excels at both natural language understanding and generation.
- **ERNIE-ViL;** a knowledge-enhanced approach to learn joint representations of vision and language using structure to enhance vision-language pre-training.
- **ERNIE-ViLG 2.0:** a text-to-image diffusion model with a knowledge-enhanced mixture of de-noising experts that improve fine-grained semantic control and alleviates the problem of object-attribute mismatching in generated images.
- **Big Model ERNIE family**: a series of Big Models in collaboration with multiple industries and companies that apply a full-featured AI development platform to a specific industry application scenario. Eleven new models under the Big Model ERNIE family were unveiled in December 2022, bringing the total number to 36.

© 2023 Citigroup

# Generative AI and Search

Generative AI is redefining search and represents one of the most significant evolutions on the internet. This is because generative AI is redefining the overall search and browsing experience by making it more natural and intuitive. With generative AI integrated into the search experience, e.g., Microsoft's AI-powered ChatGPT for Bing and Edge and Google's Bard, search and the search engine results page (SERP) are becoming more conversational, more personal, and in many ways more like a personal concierge that could change how we search for travel, buy goods, and research products.

Over the past few years, search has evolved from its original "10 blue links" to a more visual experience. Google's Multisearch (text, images, and video), Lens (image recognition technology that brings up relevant information related to objects it identifies through visual analysis), and QuickAnswers (presented at the top of a search engine results page) are examples here. But with generative AI, the search experience takes a significant step forward as it evolves from just presenting relevant links to users, to delivering relevant answers directly in the search engine results page. Generative AI solves for the almost 40% of click-backs in search (per Microsoft) when users do not always find the most relevant answer. This should improve the user experience, while also potentially further consolidating the market.

**Figure 7. Search Engine Worldwide Share (Jan 2023)**

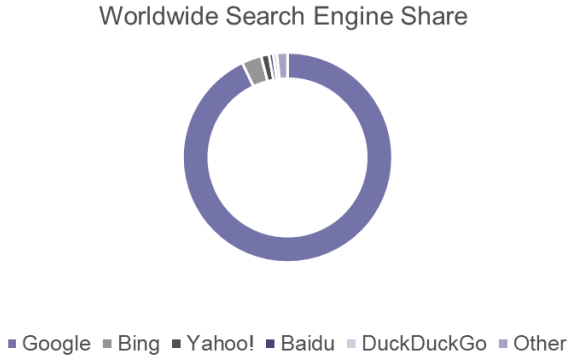


Source: Statcounter, Citi GPS

**Figure 8. Search Engine U.S. Share (Jan 2023)**

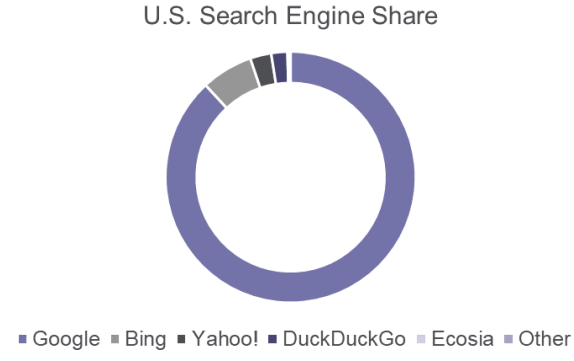


Source: Statcounter, Citi GPS

**Figure 9. Browser Worldwide Share (2022)**

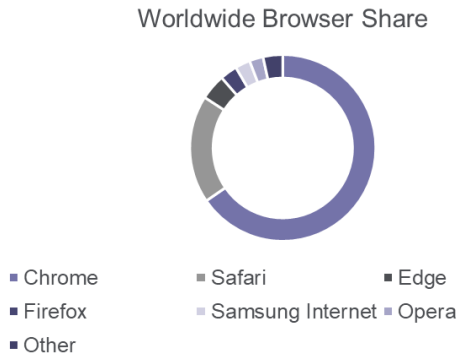


Source: Citi GPS, Statcounter

**Figure 10. Browser U.S. Share (2022)**

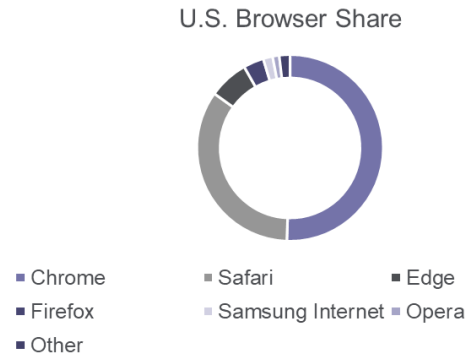


Source: Statcounter, Citi GPS

© 2023 Citigroup

At stake is the $225 global search advertising market (per eMarketer), with the introduction of generative AI in search engines raising the competition in search to a level not seen in some time. Also at stake is the broader digital advertising market, which has a combined total addressable market in 2022 of $570 billion.

**Figure 11. Worldwide Search Ad Share (2022)**

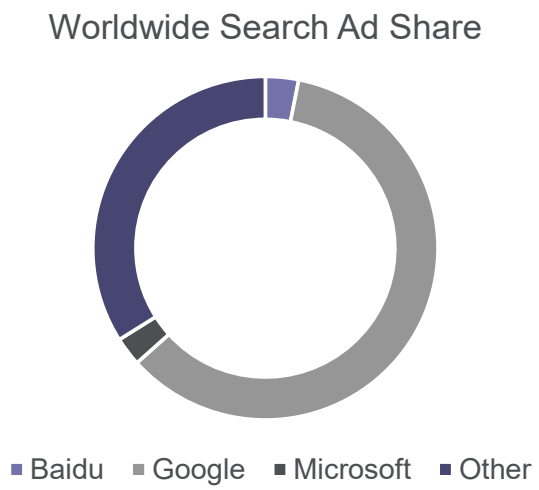


eMarketer broadly defines Search Ad revenue as net ad revenues after companies pay traffic acquisition costs (TAC) to partner sites and includes contextual text links, paid inclusion, paid listings (paid search), and search engine optimization. Hence, market share % in chart above does not reflect market share dynamics based on search engine web traffic.

Source: eMarketer, Citi GPS

© 2023 Citigroup

# Large Language Models and Transformers

Citi Global Insights

Language models are in essence models that predict what word or series of words come next in a sentence. Large Language Models (LLMs) use large neural networks to determine the context of words and provide probability distributions as to which words go next.[4] Almost all of us have encountered Large Language Models at some point — they are in effect what underline the predictive text models we encounter when we search for something on an online search engine. LLMs also power the digital assistants in use today, from Apple's Siri to Amazon's Alexa, and in that sense have become a ubiquitous part of our day-to-day lives.

Large Language Models are a very versatile form of AI with many use cases. In addition to being used for predictive text, they can be also used to improve speech recognition software by reducing the likelihood that unusual or illogical words (within the context of the overall sentence) are transcribed by the software. The ability of LLMs to accurately predict text and understand the context of conversations has lent itself to them being used as online support chatbots, saving companies considerable time and money from employing individuals

**Figure 12. Large Language Models**

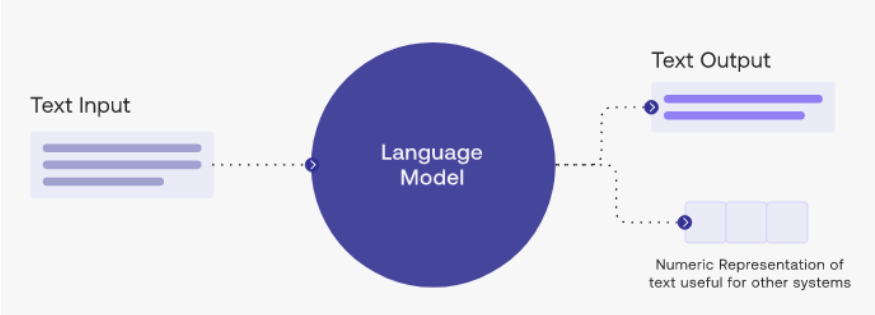


Source: co:here, "Introduction to Large Language Models"

## Breakthroughs in the Development of LLMs

The use of LLMs have heralded many developments in Natural Language Processing (NLP), a branch of AI that aims to enable computers to understand natural human language in the same way as humans.[5] One of the key breakthrough moments in NLP was the introduction of a transformer model by Google in 2017. Before this, Recurrent Neural Networks (RNNs) and Convolutional Neural Networks (CNNs) were two popular and dominant models of understanding human language until Google demonstrated that its transformer model was better able to understand human language.

[4] Shalini Urs, "The Power and the Pitfalls of Large Language Models: A Fireside Chat with Ricardo Baeza-Yates," Information Matters, May 4, 2022.
[5] IBM, "What is Natural Language Processing," Accessed February 14, 2023.

© 2023 Citigroup

Google's transformer model included a new novel "attention" mechanism, which took into account the relationship between all words in a sentence, and then weighted them appropriately.[6] One of the key takeaways was not just that the Transformer model was able to perform better at NLP, but that it was more efficiently able to use the computing resources available to it, training on eight NVIDIA P100 graphic processing units (GPUs) for just 3.5 days. That resulted in training costs that were a small fraction of the next-best models.[7] Ultimately, these transformer neural networks effectively learn context, and thus meaning, through the tracking of relationships in sequential data such as the words that make up a sentence.[8]

The importance of transformer model LLMs understanding context cannot be overstated. One desirable behavior of LLMs is for them to give precedence to context in a situation when it contains task-relevant information that may conflict with any of the model's memorized knowledge — the idea being that predictions are grounded in context, which can then go on to be used to correct future predictions without regular re-training.[9] In 2019, Fellow and Vice President of Search at Google Pandu Nayak, described the addition of machine learning language models to Google Search as, "the biggest leap forward in the past five years, and one of the biggest leaps forward in the history of Search."[10]

Program Manager of the Bing platform at Microsoft, Jeffrey Zhu, similarly described transformers as a "breakthrough in natural language understanding" and that "unlike previous Deep Neural Network (DNN) architectures that processed words individually in order, transformers understand the context and relationships between each word and all the words around it in a sentence."[11]

## A New Era for AI

Although transformers offer many benefits to LLMs, it is important to note that they can be applied to any situation that uses sequential text, images, or even videos. This is because transformers eliminate the need for large, labeled datasets to train models. Rather than the expensive training datasets used previously, the ability of transformer models to identify relationships between different elements in a dataset meant that it could in effect learn from the masses of unlabeled data on the internet.[12] As the CEO of NVIDIA, Jensen Huang put it in 2022, "transformers made self-supervised training possible, and AI jumped to warp speed."[13]

Transformers have been dubbed by researchers at Stanford as a "foundational model" of AI. That is, "any model that is trained on broad data (generally using self-supervision at scale) that can be adapted (e.g., fine-tuned) to a wide range of

---

[6] Diego Negri, "Transformer NLP & Machine Learnings: Size Does Count, But Size Isn't Everything!", Eidosmedia, March 15, 2021.
[7] Ashish Vaswani et al., "Attention is All You Need," Part of Advances in Neural Information Processing Systems 20 (NIPS), 2017.
[8] Rick Merritt, "What Is a Transformer Model," NVIDIA, March 25, 2022.
[9] Daliang Li, et al., "Large Language Models with Controllable Working Memory," DeepAI, November 2022.
[10] Pandu Nayak, "Understanding Searches Better than Ever Before," Google: The Keyword blog, October 25, 2019.
[11] Jeffrey Zhu, 'Bing Delivers Its Largest Improvement in Search Experience Using Azure GPUs," Azure blog, November 18, 2019.
[12] Rick Merritt, "What Is a Transformer Model," NVIDIA, March 25, 2022.
[13] YouTube, "GTC 2022 Spring Keynote with NVIDIA CEO Jensen Huang, Accessed February 14, 2023.

© 2023 Citigroup

downstream tasks," and went on to provide examples of transformer models such as BERT and GPT-3.[14]

**Figure 13. The Progress of Foundational Models**

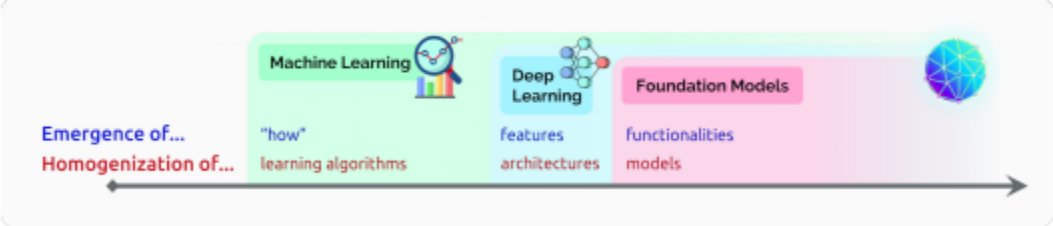


Source: Rishi Bommasani et al., "On the Opportunities and Risks of Foundation Models"

This is likely because one of the greatest leaps in the abilities of LLMs in recent years came from OpenAI's launch of GPT-3, an autoregressive language model with 175 billion parameters (10x more than any previous non-sparse language model).[15] For context, GPT-3 replaced GPT-2, which had only 1.5 billion parameters, i.e., more than 100 times smaller. Unsurprisingly, GPT-3 has been reported to perform considerably better at some tasks it was explicitly trained for. What has been a surprise to many, however, is the way that a relatively simple scaling of the learning dataset and computational power of the model, has resulted in GPT-3 being able to perform better at tasks it was not explicitly trained on.[16] This has led to excitement in the community of the potentially unforeseen flexibility of such models.

[14] Rishi Bommasani et al., "On the Opportunities and Risks of Foundation Models," Center for Research on Foundational Models, Stanford Institute for Human-Centered Artificial Intelligence, Stanford University, 2021.
[15] Tom B. Brown et al., "Language Models are Few-Shot Learners," Advances in Neural Information Processing Systems 33 (NeurIPs), 2020.
[16] Alex Tamkin, Miles Brundage, Jack Clark, and Deep Ganguli, "Understanding the Capabilities, Limitations, and Societal Impact of Large Language Models," February 2021.

© 2023 Citigroup

**Figure 14. The Exponential Increase in the Number of Parameters in NLP models**

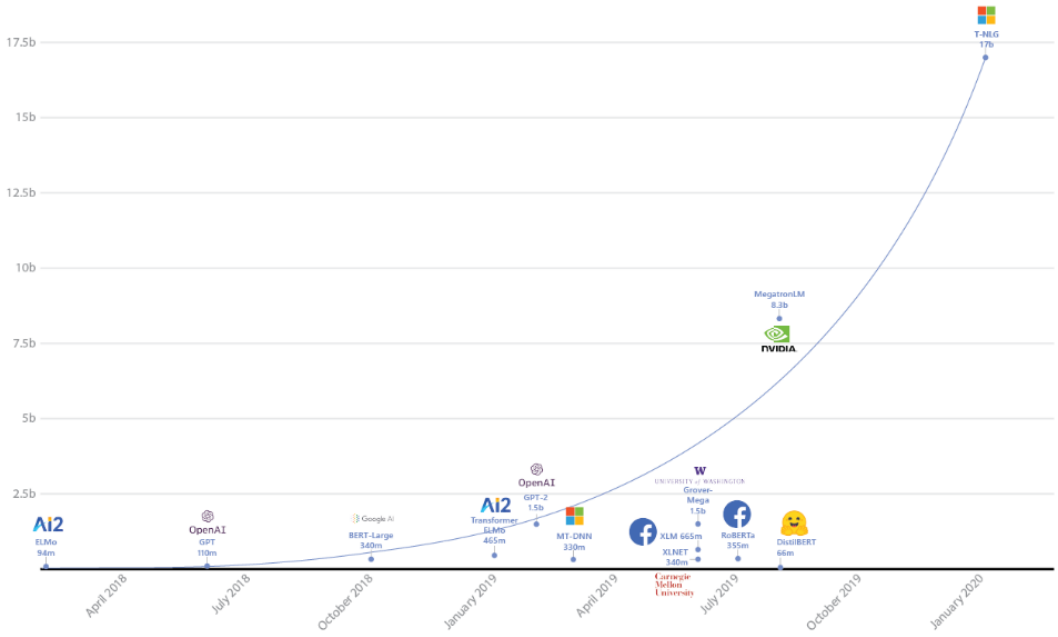


Source: Microsoft

The ultimate form of AI, artificial general intelligence, should, in principle, be able to learn general problems.[17] In recent years, LLMs have increased exponentially in size, with some describing them as increasing by ten-fold every year.[18] This has led some to question whether the development of LLMs could be considered as a new Moore's Law.[19] Saying that, with the increase in the size of LLMs, there has been an increase in computing power needed to train them, so much so that the capabilities of AI are beginning to be largely restricted by the computing power currently available. With a single NVIDIA V100 (a GPU specially designed for AI training), it would take 355 years and a $4.6 million electricity bill to train GPT-3 to produce human-like text.[20] The next iteration, GPT-4, will allegedly have 100 trillion parameters, 500 times more then GPT-3. With Moore's Law under pressure in recent years, training GPT-4 models will likely be an extremely difficult task for even our best computers.[21]

While there are a number of challenges ahead, modern LLMs (and the transformer model upon which they are based) are not just held in high regard for their excellent ability to understand and generate text, but also their ability to internalize massive amounts of real-world knowledge during initial training.

[17] Peter Voss, *Essentials of General Intelligence: The Direct Path to Artificial General Intelligence*," In B.Goertzel and C. Pennachin (Eds.) Artificial General Intelligence, (Springer Berlin Heidelberg, 2007).
[18] NVIDIA, "Codify Intelligence with Large Language Models," Accessed February 14, 2023.
[19] Julien Simon, "Large Language Models: A New Moore's Law?", Hugging Face, October 26, 2021.
[20] Chuan Li, "OpenAI's GPT-3 Language Model: A Technical Overview," Lambda Labs, June 3, 2022.
[21] Alberto Romero, "GPT-4 Will Have 100 Trillion Parameters — 500x the Size of GPT-2," Towards Data Science, September 22, 2021.

© 2023 Citigroup

© 2023 Citigroup

If you are visually impaired and would like to speak to a Citi representative regarding the details of the graphics in this document, please call USA 1-888-800-5008 (TTY: 711), from outside the US +1-210-677-3788

## IMPORTANT DISCLOSURES

This communication has been prepared by Citigroup Global Markets Inc. and is distributed by or through its locally authorised affiliates (collectively, the "Firm") [E6GYB6412478]. This communication is not intended to constitute "research" as that term is defined by applicable regulations. Unless otherwise indicated, any reference to a research report or research recommendation is not intended to represent the whole report and is not in itself considered a recommendation or research report. The views expressed by each author herein are his/ her personal views and do not necessarily reflect the views of his/ her employer or any affiliated entity or the other authors, may differ from the views of other personnel at such entities, and may change without notice.

You should assume the following: The Firm may be the issuer of, or may trade as principal in, the financial instruments referred to in this communication or other related financial instruments. The author of this communication may have discussed the information contained herein with others within the Firm and the author and such other Firm personnel may have already acted on the basis of this information (including by trading for the Firm's proprietary accounts or communicating the information contained herein to other customers of the Firm). The Firm performs or seeks to perform investment banking and other services for the issuer of any such financial instruments. The Firm, the Firm's personnel (including those with whom the author may have consulted in the preparation of this communication), and other customers of the Firm may be long or short the financial instruments referred to herein, may have acquired such positions at prices and market conditions that are no longer available, and may have interests different or adverse to your interests.

This communication is provided for information and discussion purposes only. It does not constitute an offer or solicitation to purchase or sell any financial instruments. The information contained in this communication is based on generally available information and, although obtained from sources believed by the Firm to be reliable, its accuracy and completeness is not guaranteed. Certain personnel or business areas of the Firm may have access to or have acquired material non-public information that may have an impact (positive or negative) on the information contained herein, but that is not available to or known by the author of this communication.

The Firm shall have no liability to the user or to third parties, for the quality, accuracy, timeliness, continued availability or completeness of the data nor for any special, direct, indirect, incidental or consequential loss or damage which may be sustained because of the use of the information in this communication or otherwise arising in connection with this communication, provided that this exclusion of liability shall not exclude or limit any liability under any law or regulation applicable to the Firm that may not be excluded or restricted.

The provision of information is not based on your individual circumstances and should not be relied upon as an assessment of suitability for you of a particular product or transaction. Even if we possess information as to your objectives in relation to any transaction, series of transactions or trading strategy, this will not be deemed sufficient for any assessment of suitability for you of any transaction, series of transactions or trading strategy.

The Firm is not acting as your advisor, fiduciary or agent and is not managing your account. The information herein does not constitute investment advice and the Firm makes no recommendation as to the suitability of any of the products or transactions mentioned. Any trading or investment decisions you take are in reliance on your own analysis and judgment and/or that of your advisors and not in reliance on us. Therefore, prior to entering into any transaction, you should determine, without reliance on the Firm, the economic risks or merits, as well as the legal, tax and accounting characteristics and consequences of the transaction and that you are able to assume these risks.

Financial instruments denominated in a foreign currency are subject to exchange rate fluctuations, which may have an adverse effect on the price or value of an investment in such products. Investments in financial instruments carry significant risk, including the possible loss of the principal amount invested. Investors should obtain advice from their own tax, financial, legal and other advisors, and only make investment decisions on the basis of the investor's own objectives, experience and resources.

This communication is not intended to forecast or predict future events. Past performance is not a guarantee or indication of future results. Any prices provided herein (other than those that are identified as being historical) are indicative only and do not represent firm quotes as to either price or size. You should contact your local representative directly if you are interested in buying or selling any financial instrument, or pursuing any trading strategy, mentioned herein. No liability is accepted by the Firm for any loss (whether direct, indirect or consequential) that may arise from any use of the information contained herein or derived herefrom.

Although the Firm is affiliated with Citibank, N.A. (together with its subsidiaries and branches worldwide, "Citibank"), you should be aware that none of the other financial instruments mentioned in this communication (unless expressly stated otherwise) are (i) insured by the Federal Deposit Insurance Corporation or any other governmental authority, or (ii) deposits or other obligations of, or guaranteed by, Citibank or any other insured depository institution. This communication contains data compilations, writings and information that are proprietary to the Firm and protected under copyright and other intellectual property laws, and may not be redistributed or otherwise transmitted by you to any other person for any purpose.

**IRS Circular 230 Disclosure:** Citi and its employees are not in the business of providing, and do not provide, tax or legal advice to any taxpayer outside of Citi. Any statements in this Communication to tax matters were not intended or written to be used, and cannot be used or relied upon, by any taxpayer for the purpose of avoiding tax penalties. Any such taxpayer should seek advice based on the taxpayer's particular circumstances from an independent tax advisor.

© 2023 Citigroup Global Markets Inc. Member SIPC. All rights reserved. Citi and Citi and Arc Design are trademarks and service marks of Citigroup Inc. or its affiliates and are used and registered throughout the world.