ChatGPT and Generative AI

Just Hype, or an Internet and Tech Disruptor?

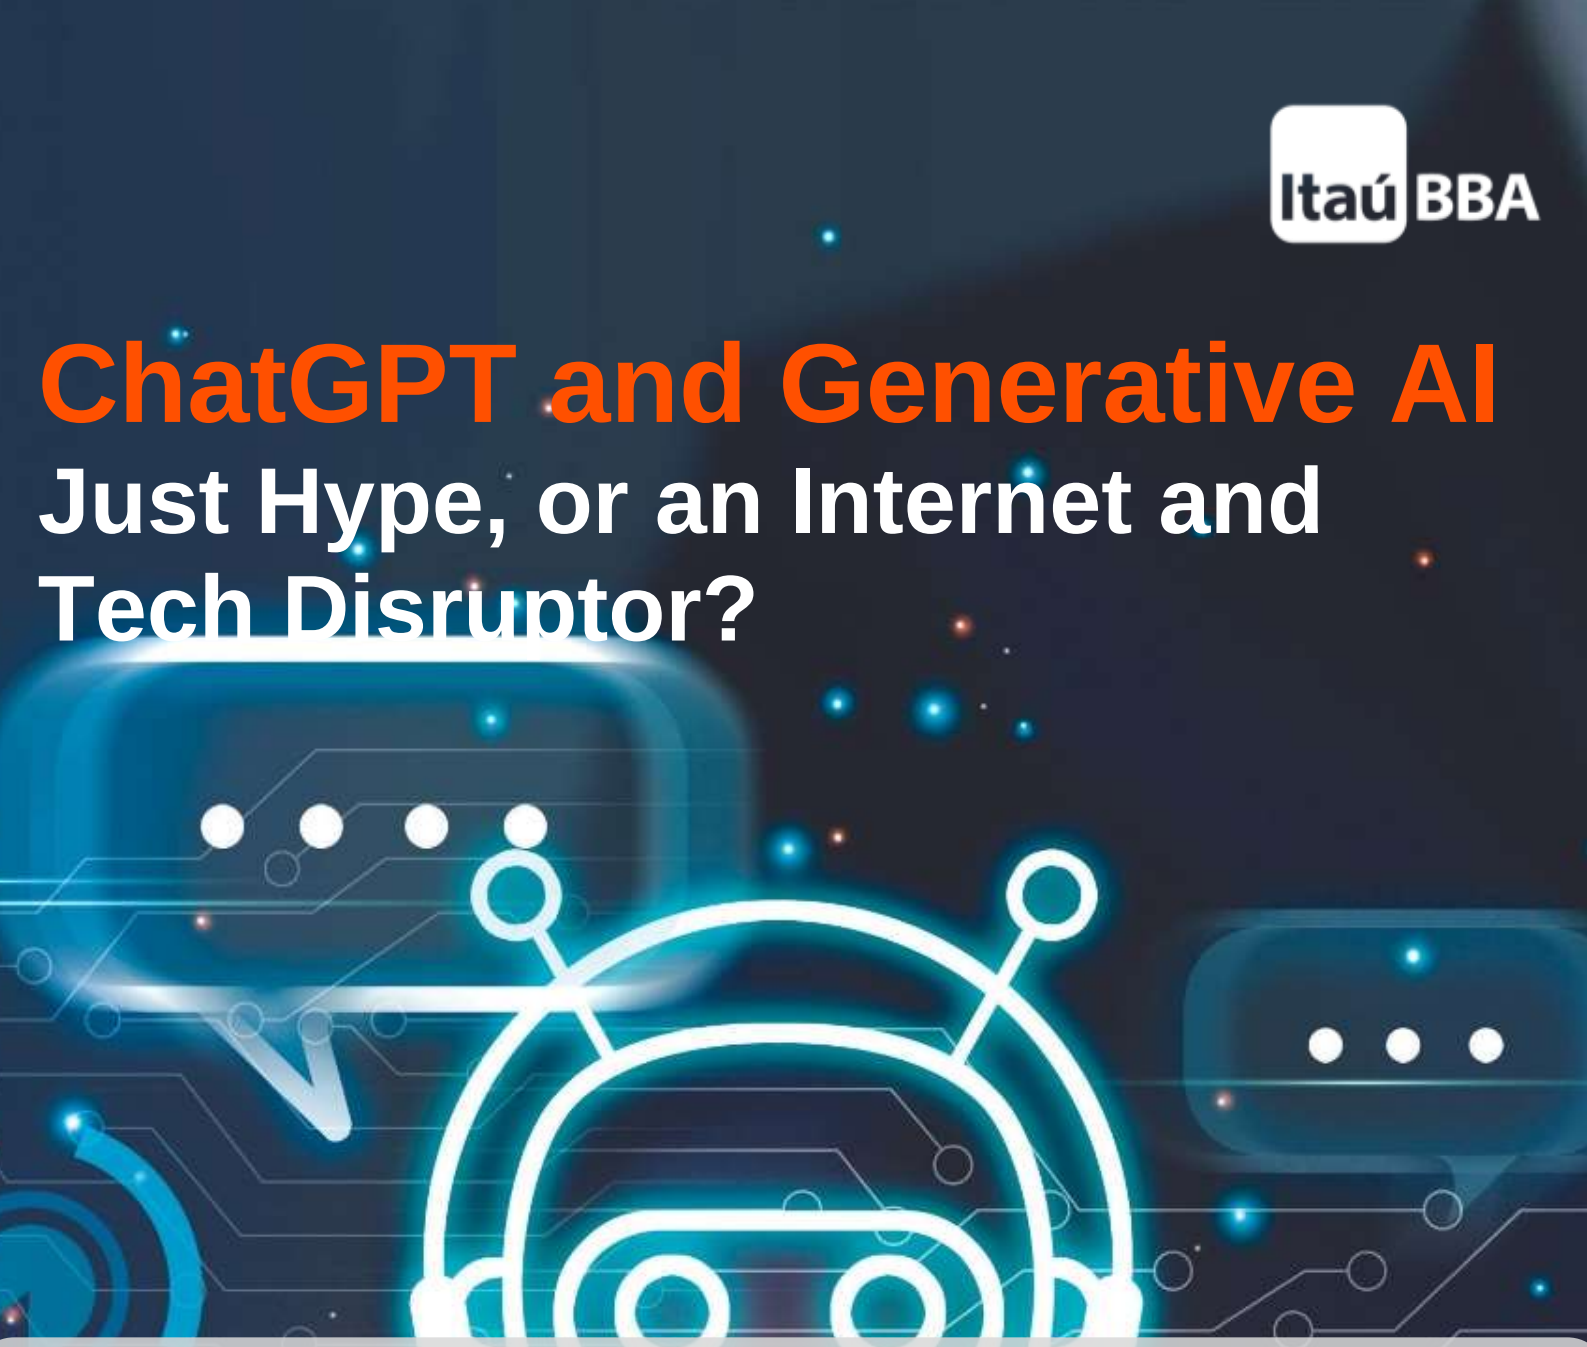

February 27, 2023

TECH TEAM

Thiago Alves Kapulskis, CNPI
+55-11-91585-0188
thiago.kapulskis@itaubba.com

Cristian Faria, CNPI
+55-11-93025-7484
cristian.faria@itaubba.com

ChatGPT is a large language model developed by OpenAI. It is designed to be a conversational agent that can interact with humans through natural language, like a chatbot. The "GPT" in its name stands for "Generative Pre-trained Transformer," which refers to the neural network architecture used to build the model. ChatGPT has been trained on a vast corpus of text data from the internet, including books, articles, and websites, allowing it to generate responses to a wide range of queries and questions. Its capabilities include answering questions, providing information, making recommendations, and engaging in natural language conversations with users.

Guess who wrote the above? Well, none of the humans signing this report. This definition of ChatGPT was written by ChatGPT itself. Its capacity to deliver responses like this – thorough, fluent, plausibly attributable to a live human – attracted not only us but more than 100 million other people to use ChatGPT in January alone. It was the fastest adoption of an internet tool in history. (It took 10 months for TikTok to get to 100 million MAUs!)

If this is so captivating for us, it could also be a very good business opportunity for companies. There is untapped potential to capture at least some of the world's daily 10 billion search queries, which could translate into annual revenues of USD ~200 billion in the search market alone. Given such huge potential, it comes as no surprise to learn that Microsoft has been working with OpenAI (ChatGPT's owner) in stealth mode for the past four years! And yet this still seems to be only the beginning.

With the world's imagination captured by this crazy phenomenon, we are getting tons of questions on the topic. We can't deny that we are glad about this – we take it as a good sign that the role of sell-side analyst hasn't been disrupted by ChatGPT (yet…).

After two months of intense work on generative AI, please see our main findings in the next two pages and much more inside this report.

- ChatGPT seems to be just a “Level 1” AI tool, which means that we are still in the very early innings of AI. In a nutshell, ChatGPT doesn’t “do” anything, it just “sees” – meaning, in this case, that it infers from very advanced statistics and mathematical calculations the probable “next word” in a sequence of words (written or being written).
- Many companies, including NVIDIA, Google, META, Baidu and now Microsoft (with OpenAI) have large language models (LLMs) with more than 100 billion parameters; these models are the basis for ChatGPT-like tools. This is another reason why we are likely to see more innovation and competition in this arena.
- We believe that Google has the technology to rapidly adapt to LLMs in terms of market share and products. Our impression, however, is that bureaucracy and “wrong” incentives (like stock-price-linked compensation) at Google led to overly risk-averse behavior, allowing ChatGPT to evolve faster than BARD or any of Google’s other solutions. (Google’s LaMDA LLM existed long before GPT-3 LLM, the model that is the basis for ChatGPT.)
- On the other hand, we believe that Google could be weighed down by much higher costs, which could reduce its EBIT by as much as 60% in an extreme scenario. The main factor determining this profitability erosion (its speed and scale) is likely to be the competition, in our view; the more companies try to capture search ad dollars via ChatGPT-like tools, the more queries Google will have to treat as ~10x-higher-cost LLM queries.

- **Who are the winners?**
  - **Microsoft**: While in this report we do not explore the many alternative platforms Microsoft has – PowerApps, Teams, GitHub and more – we believe that it could be a big winner in AI in the long term (even if Bing + ChatGPT is unsuccessful).
  - Among other big techs, both **Meta** (given its renewed tech stack) and **Apple** (given the dominance of iOS) could be secondary winners, via either disruptive products (Meta) or more revenues coming from Search Engine bids (Apple).
  - **Semis is definitely the winning sector in this whole story**, but not all semis are the same. Our main pick on the AI theme – by far – is **NVIDIA**, which is far ahead in adoption, even if for research purposes only, in a world where parallel computing will be subject to more and more demand.
  - We also believe that **data warehousing companies like SNOW** could benefit, as AI advances could untap even more demand for analytics.
  - Watch **Baidu** – it seems to the best (non-academic) Chinese player in this arena.
  - Within our Brazilian coverage universe, the only company that appears to be exposed to the theme is **Zenvia**, which could have more use for CPaaS tools. Outside our coverage universe, **Semantix** could benefit from higher demand for data warehousing (as AI enables larger-scale data analytics).
- **Who are the losers (besides Google)?**
  - **Internet companies that are heavily dependent on search traffic – the likes of Booking, Expedia and Trip.com** – could suffer as they try to adapt their services to a potentially new and more costly traffic engine.
  - **Online education companies (like Chegg)** could suffer, as ChatGPT is a cheap and fast alternative to low-cost databases or courses. Watch out for companies exposed to low-income student bases.
  - **ZM** (once again) looks like a loser, as MSFT is already adopting ChatGPT inside Teams.
- **Trading ideas specifically based on this theme:**
  - **Longs: NVDA, MSFT, SNOW, BIDU, ZENV, STIX**
- **Shorts: GOOGL, BKNG, EXPE, TRIP, CHGG, ZM**

Contents

HOW TO READ THIS REPORT

SECTION 1

Our Big Techs IoC started with a “how to read this report” section, given the large size of the report, the number of topics it covered and the likely diversity of readers’ interests related to the subject. We will do the same here, to help orient readers of this piece.

- If you are interested in learning how AI works and how ChatGPT was trained and created, we recommend that you read Section 2 (“AI 101: The Current State of Ai and LLMs, How They Work, and Safety Concerns”).
- If you want to learn about the LLM market (the models that are the basis for ChatGPT-like tools) and potential future players/competitors in this space, we recommend that you read Section 3 (“The LLM Market and Why We Think This Is Just The Beginning…”).
- If you are curious about which subsectors could be affected by ChatGPT, see Section 4 (“When Mammoths Fight, You’d Better Be Far Away”: Google Search Versus Microsoft Bing, and the Potential Consequences for Internet Subsectors”). In the same section, we also compare the experiences of Google Search, ChatGPT and BinGPT (ChatGPT inside Bing). We recommend this section to all investors – including those who cover only LatAm or emerging markets.
- If you are interested in an analysis of all the Big Techs (Google, Microsoft, Meta, Amazon and Apple) in relation to this topic, read Section 5 (“Thoughts on the Big Techs: The Winners, the Losers and Some Math). Our analysis includes some food for thought on competition, a quantification of the potential impact on Google’s margins, and more.
- Finally, if you cover Semis, we have written a brief analysis of that segment, focusing on NVIDIA. (See Section 6, “Semis Are the Clear Winners; Nvidia Would Be Our Pick”.)

AI 101: THE CURRENT STATE OF AI AND LLMS, HOW THEY WORK, AND SAFETY CONCERNS

SECTION 2

Back at the beginning of January, we were very lucky to meet a great friend from our MBA program who has worked with AI for about 20 years. He recommended an excellent book called “The Book of Why: The New Science of Cause and Effect ”, by Judea Pearl, as a reference source that can help non-developers understand what’s behind AI and catch up on its current state. And indeed, this masterpiece is a must – we highly recommend that you read it. The first section of this report is largely based on this book.

Association (Level 1), Intervention (Level 2) and Counterfactuals (Level 3)

The three “intelligence levels”. ChatGPT is Level 1 – the lowest!

Judea Pearl created a very interesting framework to define “levels of intelligence”. He describes three different levels (see figure below).

Level 1: Association. This is the lowest level, seen in animals and some robots. At this level, the brain merely makes associations with things it sees (“If I see X, then Y”). For example, if a springbok sees a lion, it runs. The springbok doesn’t really need to understand *why* the lion is there – it just needs to perceive that it *is* there. Associations can be well explained by statistics (i.e., correlations).

Level 2: Intervention. This is an intermediate level which is today seen in human babies and was (likely) prevalent among early hominids. It’s basically defined by the attaching of consequences to actions taken by the subject (“If I *do* X, then Y”). While this may look very similar to Level 1, it’s not, because intervention requires an understanding of causes and effects. Pearl argues that statistics is not enough to produce this level of intelligence – you need more mathematical tools, including probabilistic thinking.

Level 3: Counterfactuals. This is the “supreme” level of intelligence as we understand it today. It requires imagination – an ability currently only seen in humans, according to Yuval Harari’s book *Sapiens* – which is used to evaluate “alternative worlds” (other situations that could have happened). It is essentially defined by the questions “What would X have been if Y had not happened?” and “Was it Y that caused X?”.

Interestingly, ChatGPT (and its GPT-3 LLM) is still classified as a Level 1 intelligence, albeit one that is probably some distance down the road toward Level 2. It’s still a tool that works mostly on statistics (and “perceptions”), as we will describe further in the next section.

Judea Pearl's Three-Tier Framework for Intelligence Levels

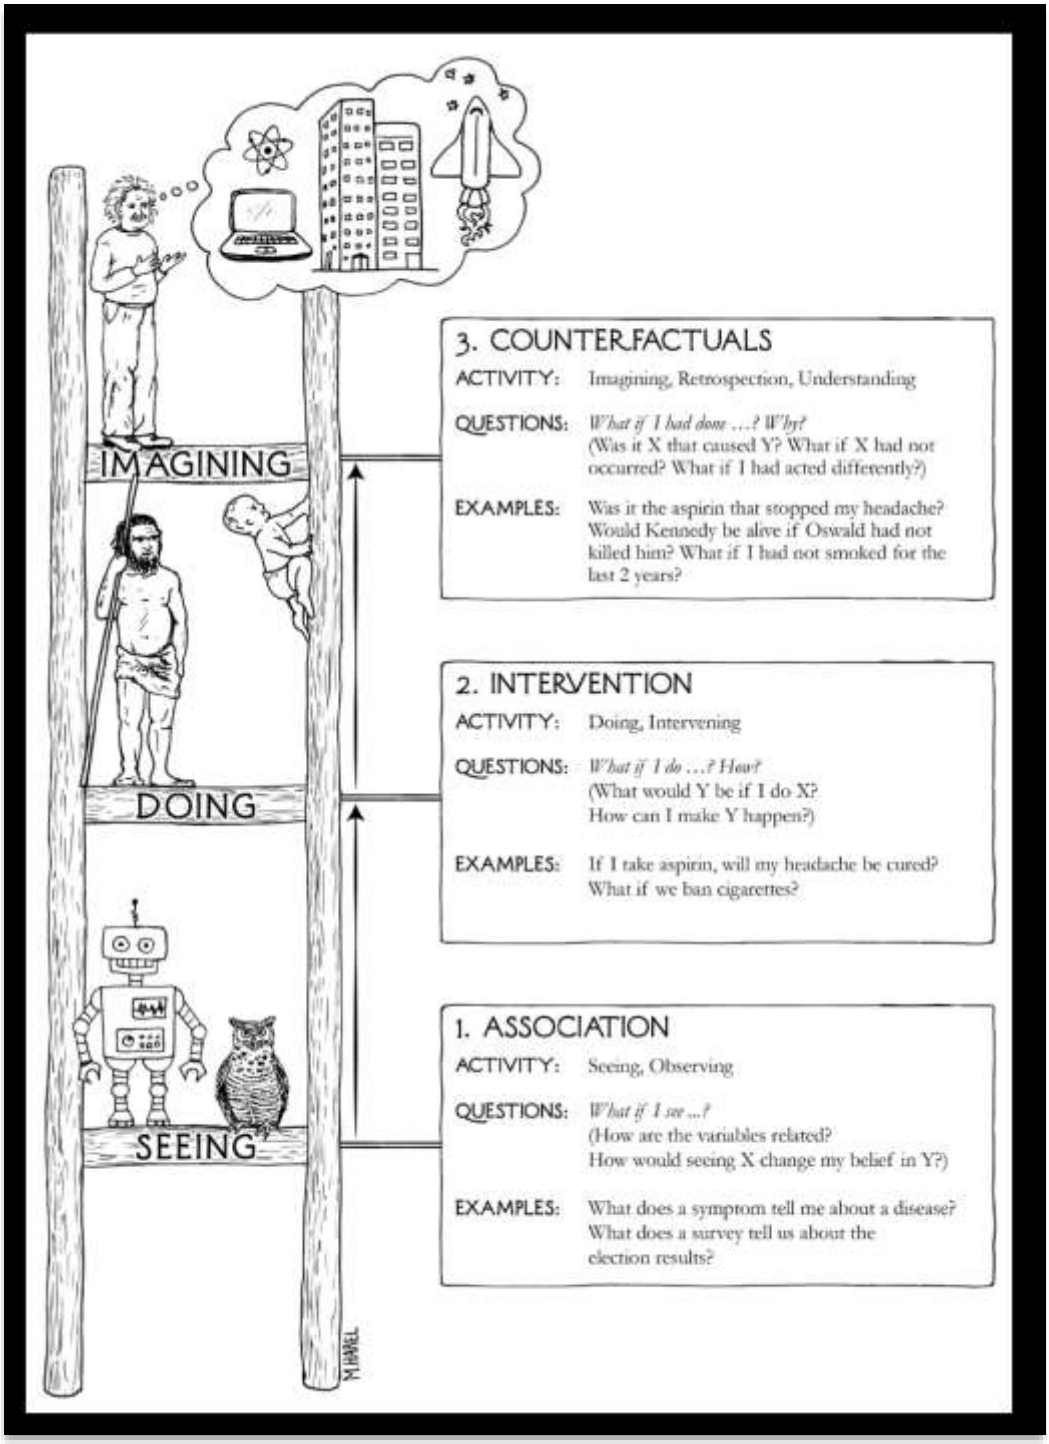

Source: The Book of Why

Inference Engines: The Framework Behind AI

Before we jump in into LLMs (large language models) and ChatGPT, let's first take a look at Judea Pearl's framework for AI systems, which is quite useful for understanding LLMs conceptually.

As can be seen in the figure below, Pearl's AI framework is quite complex, but the essence of how it works can be conveyed via an example.

Let's say a human wants to know if taking aspirin would alleviate a headache.

Question: *Does taking aspirin reduce headache?*

- **Step 1**: **Defining knowledge and assumptions.** In this first step, we define what is common knowledge and what the model's assumptions (boundaries) are, relative to what needs to be evaluated. For example, we could say that "taking an aspirin" means taking a pill, putting it in one's mouth and swallowing it. This is "knowledge". In our example, the model assumes that people already know this.
- **Step 2: Can the query be answered?** This is a very important step. It could be that the question was poorly written, or that the dataset we have to address the question is insufficient to answer it. For example, let's say that we asked the question above and tried to get an answer based on a database that contains only information about football games. These are completely unrelated areas of knowledge, so it would be impossible for an AI that was limited to this database to answer this question.
- **Step 3: Estimand.** This is a fascinating concept in Pearl's book. Basically, it refers to the process by which the query in human language (in this example, in English) is translated into "machine language". In Pearl's book, he demonstrates that standard statistical notation is insufficient for making such translations and shows how completely new notations have been created over the past 20 years to address this. Although the math behind this is extremely complex, one good simplification would be the creation of a probability equation.
- **Step 4: Statistical estimation.** Once the estimand has been defined, we can run statistical operations on the data to find a solution for it. In this example, we could run the operations on a database of experiments involving 2,000 people who took aspirin in order to determine whether there is a correlation between aspirin and headaches (or no headaches).

Answer. The answer is the best proxy for solving the problem based on the framework above.

Judea Pearl's Framework for AI Models

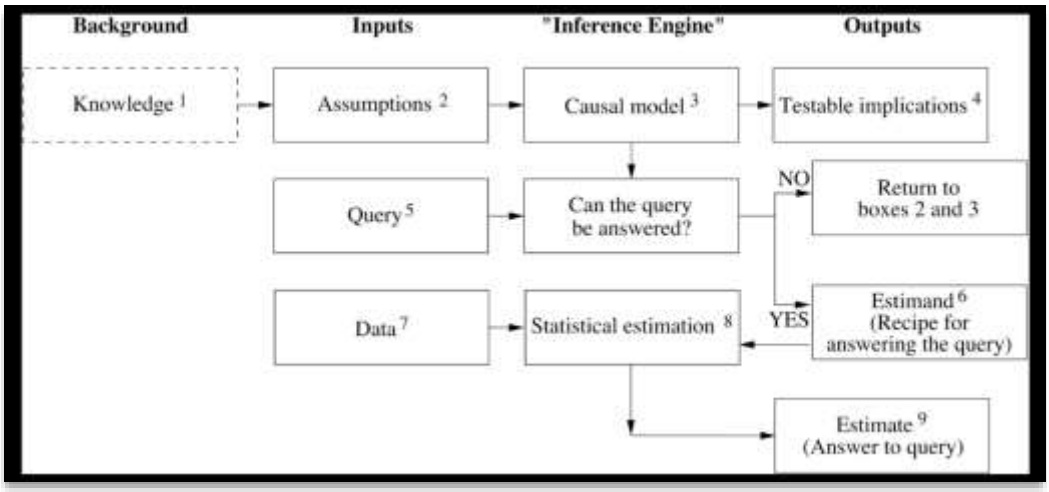

Source: The Book of Why

LLMs and ChatGPT: How ChatGPT's GPT-3 Was Trained

Now that we understand their frameworks, it's time to turn to ChatGPT and LLMs, which are a specific type of application for the process diagrammed in the box above.

ChatGPT is a chatbot based on a large language model (LLM) called GPT-3. Importantly, as we will show below in our explanation of how LLMs work, **ChatGPT doesn't understand language in the way that humans do; it operates in the aforementioned framework to capture regularities in data through statistical models (making it a Level 1 intelligence). In other words, ChatGPT "knows nothing".**

The best definition of LLMs that we've seen comes from NVIDIA:

A large language model, or LLM, is a deep learning algorithm that can recognize, summarize, translate, predict and generate text and other content based on knowledge gained from massive datasets.

Large language models are among the most successful applications of transformer models. They aren't just for teaching AIs human languages, but for understanding proteins, writing software code, and much, much more.

In addition to accelerating natural language processing applications — like translation, chatbots and AI assistants — large language models are used in healthcare, software development and use cases in many other fields.

The transformer models are the engines behind how LLMs work. Understanding these models is not easy, so we were grateful when another good friend of ours, who happens to have worked at Google AI in the past, helped us out by steering us to a couple of great articles at bdtechtalks.com which explain transformers. Here's a summary of the content of these articles:

I. In recent years, the transformer model has become a highlight of the advances in deep learning and deep neural networks. It is mainly used for advanced applications in natural language processing. Google is using it to enhance its search engine results. OpenAI used transformers to create its famous GPT-2 and GPT-3 models.

II. First, the input text must be processed and transformed into a unified format before being fed to the transformer.

- First, the text goes through a "tokenizer," which breaks it down into chunks of characters that can be processed separately (see figure below).
- In most cases, every word and punctuation mark roughly counts as one token. Some suffixes and prefixes count as separate tokens (e.g., "ize," "ly," and "pre").
- The tokenizer produces a list of numbers that represent the token IDs of the input text.

I. The tokens are then converted into "word embeddings." A word embedding is a vector that tries to capture the value of words in a multi-dimensional space.

- For example, the words "cat" and "dog" can have similar values across some dimensions because they are both used in sentences that are about animals and house pets. However, "cat" is closer to "lion" than to "wolf" across another dimension that separates felines from canids.
- Similarly, "Paris" and "London" might be close to each other because they are both cities. However, "London" is closer to "England", and "Paris" to "France", in another dimension that separates countries.
- Word embeddings usually have hundreds of dimensions. Word embeddings are created by embedding models, which are trained separately from the transformer. There are several pre-trained embedding models that are used for language tasks.

I. Once the sentence is transformed into a list of word embeddings, it is fed into the transformer's encoder module.

- The module can receive an entire sentence's worth of embedding values and process them in parallel.

  This makes transformers more compute-efficient than their predecessors and also enables them to examine the context of the text in both forward and backward sequences.
- To preserve the sequential nature of the words in the sentence, the transformer applies "positional encoding," which basically means that it modifies the values of each embedding vector to represent its location in the text.

I. Next, the input is passed to the first encoder block, which processes it through an “attention layer.” The attention layer tries to capture the relationships between the words in the sentence.

- For example, consider the sentence “The big black cat crossed the road after it dropped a bottle on its side.” Here, the model must associate “it” with “cat” and “its” with “bottle.” Accordingly, it should establish other associations such as “big” and “cat” or “crossed” and “cat.” In other words, the attention layer receives a list of word embeddings that represent the values of individual words and produces a list of vectors that represent both individual words and their relationships to each other. The attention layer contains multiple “attention heads,” each of which can capture different kinds of relationships between words.

I. The output of the attention layer is a vector representation, which is fed to a feed-forward neural network that is adjusted in the training process and sends it to the next attention layer in the next encoder. Transformers incorporate multiple attention blocks and feed-forward layers to gradually capture more complicated relationships.

II. The task of the decoder module is to translate the encoder’s output vector into the output data (e.g., the translated version of the input text).

- During the training phase, the decoder has access to both the output vector produced by the encoder and the expected outcome (e.g., the translated string).

  a. The decoder uses the same tokenization, word embedding and attention mechanism to process the expected outcome and create output vectors. It then passes these encoder vectors through an attention layer to the encoder module, which establishes the relationships between the input and output values. In the translation application, this is the part where the words from the source and destination languages are mapped to each other.

  b. Like the encoder module, the decoder attention vector is passed through a feed-forward layer. Its result is then mapped to a very large vector which is the size of the target data (in the case of language translation, this can amount to tens of thousands of words).

I. This process ends when the very last encoding-decoding interaction is complete (training) or a final answer to the client (user) is generated.

Text Tokenization for GPT-3

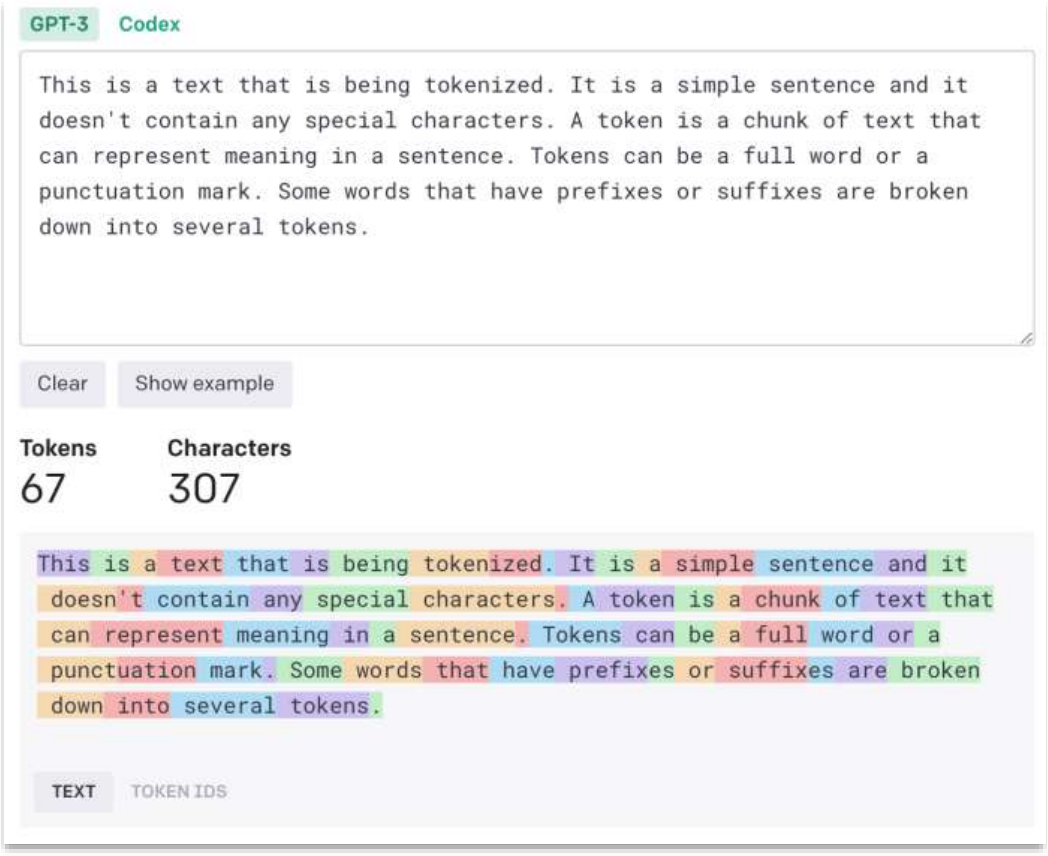

Source: BDtechtalks.com

Attention Layers: The Engines Behind Transformers (and LLMs)

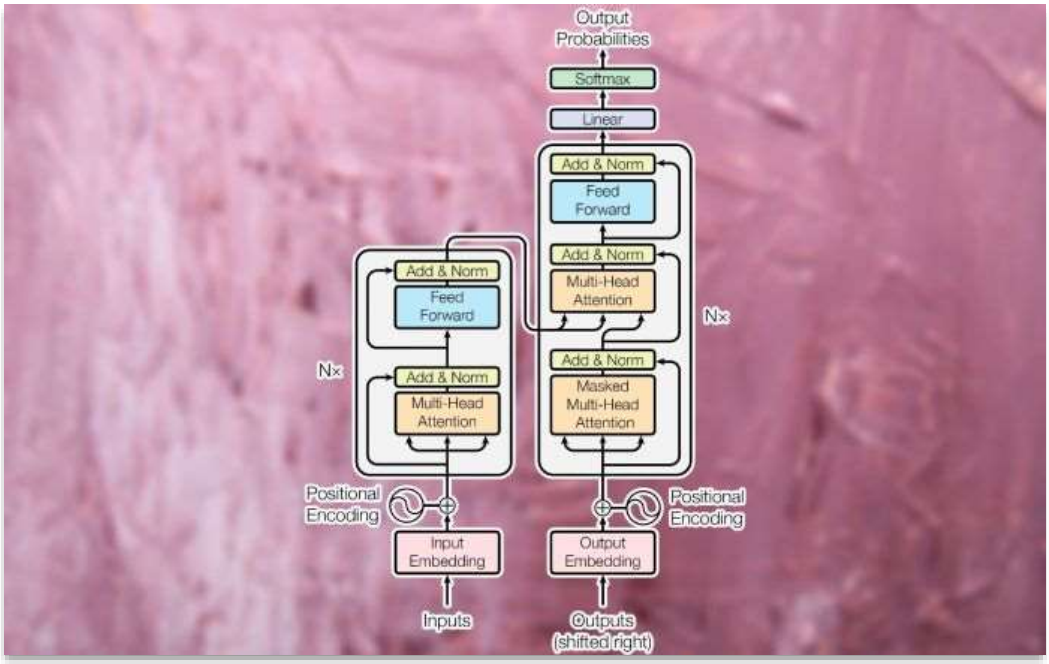

Source: BDtechtalks.com

In practice, as NVIDIA's description indicates, large language models learn from huge volumes of data. As the name suggests, one of the distinguishing features of an LLM is the size of the dataset it's trained on. But the definition of "large" is growing along with AI.

Large language models are typically trained on datasets large enough to include nearly everything that has been written on the internet over a large span of time. This is exactly how GPT-3 was trained – including Google's search data. Essentially, Google's search data (along with other datasets) were added to the framework above to create GPT-3 (and then ChatGPT).

Thus, a good description of ChatGPT is that it is a language processing tool that answers queries based on statistical inferences from all the language used on the internet up to 2020.

Could AI Have Unintended Consequences?

Yes, and it's no conspiracy theory – the UK government is already acknowledging it.

Although we've consistently heard that AI is far from causing damage to humanity, there is no consensus on this in the scientific community.

As AI investors Nathan Benaich and Ian Hogarth have written in their excellent State of AI Report (which we highly recommend reading!), early AI pioneers Alan Turing (1951), IJ Good (1965) and Marvin Minsky (1984) all sounded alerts about AI's potential risks to humanity.

In late 2021, the UK published a national strategy for AI in which it made multiple references to AI safety and the long-term risks posed by misaligned AI.

Further, a survey of the ML (machine learning) community found that 69% believe AI safety should be prioritized more than it currently is. In a different survey of the NLP (natural language processing) community, over 70% responded that they believe that AI will cause social change at the level of the Industrial Revolution and 40% responded that they believe AI could cause a nuclear war!

Surveys of the Scientific Community Raise Concerns About AI Safety

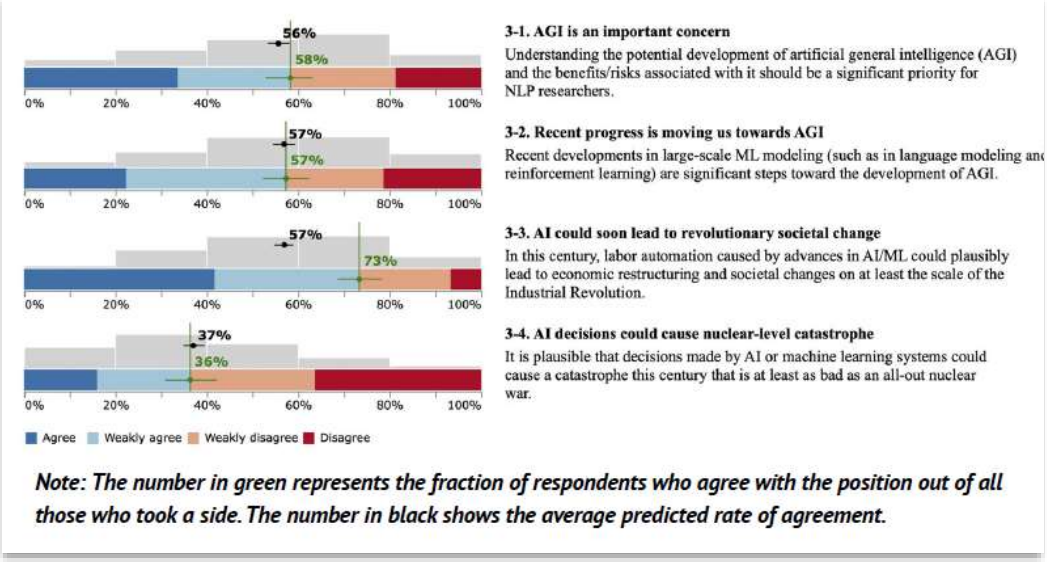

Source: State of AI 2022 Report

THE LLM MARKET AND WHY WE THINK THIS IS JUST THE BEGINNING…

SECTION 3

"BIDU was up 15% on a ChatGPT-like announcement yesterday and BABA is up on a similar thing today. Is this hype or is it for real?"

We got this question from an investor in one of our trips earlier this year. We strongly believe that we are just in the early innings of this "LLM game", and this section is intended to show evidence supporting this conclusion.

When we first studied LLMs, we were surprised by the number of models already existing today. According to Lifetime.ai, there are 13 models in the Western world with more than 100 billion parameters, as shown in the figure below. As a reminder, GPT-3, the model behind ChatGPT, has 175 billion parameters. There are two models (Google's PaLM and NVIDIA's MT-NLG) with more than 500 billion parameters. Notably, the State of AI report also showed that many tech players – including China's Huawei, Baidu and Tsinghua University – have been implementing/cloning/ improving GPT-like models.

Mapped LLMs in the Western World as of December 2022

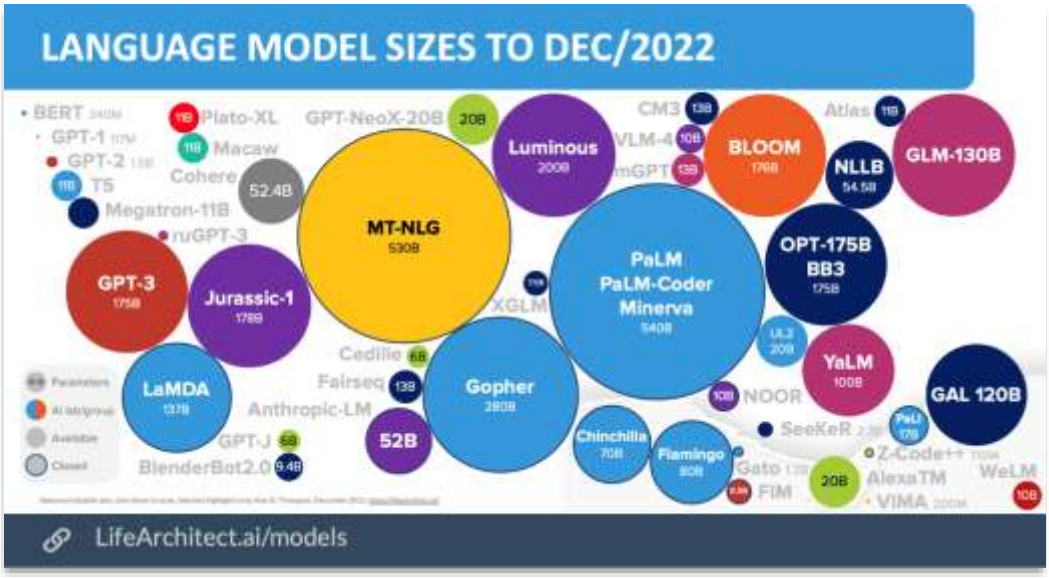

Source: LifeArchitect.ai

Many Players, Including Chinese Companies, Have Been Replicating GPT-3

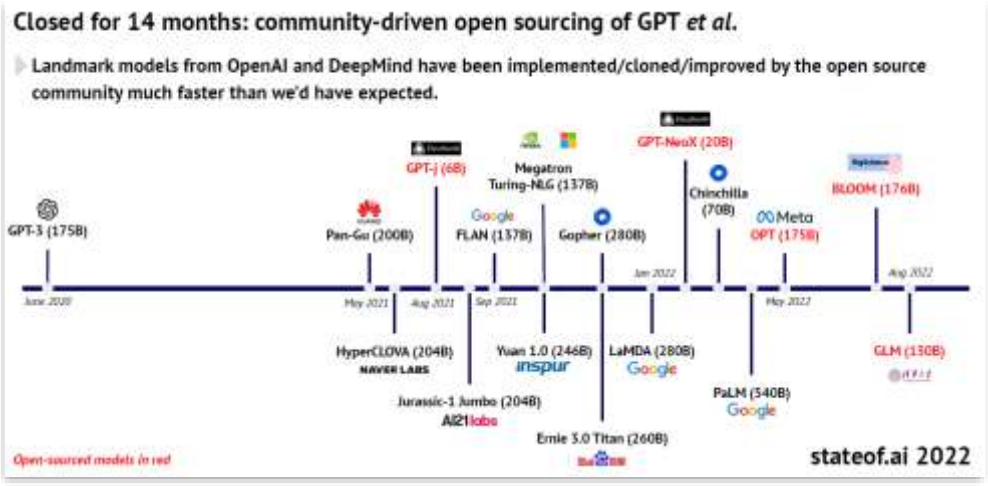

Source: State of AI Report

While the number of parameters is not the only variable that matters for an LLM, it is seen as one good guideline for assessing if one model is "better" than another, as we demonstrated in Section 2. The number of parameters refers to the size of the matrix that represents each work in LLM. Once we understand how transformer and attention layers work, we can see intuitively that the larger the dataset, the better the training, statistical inferences and (therefore) outcome.

In practice, what matters here is the overall degree of magnitude and not the specific number. So GPT-3's 175 billion is not necessarily better than Google LaMDA's 137 billion, but it is certainly better than GPT-2's 1.5 billion.

Another interesting fact is that LLMs created for a certain application can be tweaked relatively quickly for a different use case. For example, Google's LaMDA was initially designed to be the basis for Google Assistant's chatbots, but it is now used in Google's BARD tool (to compete with ChatGPT within Bing – more on this in Section 4).

A Closer Look at Who Dominates LLM: How Did Google Let This Happen?

It's also interesting to see which company owns each of the largest models, which may be a good indication of the potential winners:

- PaLM (540 billion parameters): **Google**
- MT-NLG (530 billion): **NVIDIA**
- Gopher (280 billion): **DeepMind** (Alphabet's)
- Ernie 3.0 Titan (260 billion): **Baidu**
- Yuan 1.0 (246 billion): **Inspur** (a Chinese company)
- HyperClova (204 billion): **Naver Labs** (a South Korean company)
- Pan-Gu (200 billion): **Huawei**
- Jurassic (178 billion): **AI21 Labs** (a startup)
- Bloom (176 billion): **Hugging Face** (a startup)
- OPT-175BBB3 (175 billion): **Meta**
- GPT-3 (175 billion): **OpenAI** (Microsoft owns a stake)
- LaMDA (137 billion): **Google**

Someone looking at this table might well ask: Google and Alphabet have three of the top 12 LLMs in the world. How could it have allowed ChatGPT to happen?

While there is no obvious reason for Google's failure to lead in this technology, there is certainly an important conclusion to draw: just having the technology is not enough to win the game. In interviews published on various blogs, former Google AI employees have said that the culture there is getting more risk-averse. It's important to keep in mind that Google is no longer led by its founders and its management's compensation is largely tied to the stock price. It's even more shocking to see Microsoft potentially eating Google's lunch when you learn that Google's LaMDA was available a year *before* GPT-3.

At the same time, we believe that Google is very well fortified against new entrants. It's certainly possible that it will lose share in search, but it looks like it can adapt. The largest risk is in the area of costs, as we will demonstrate in Section 5.

We are far more worried about other internet companies, particularly the ones that depend heavily on search traffic. And this is an appropriate segue into our next section.

"WHEN MAMMOTHS FIGHT, YOU'D BETTER BE FAR AWAY": GOOGLE SEARCH VERSUS MICROSOFT BING, AND THE POTENTIAL CONSEQUENCES FOR INTERNET SUBSECTORS

SECTION 4

"It is not the strongest of the species that survives, nor the most intelligent that survives. It is the one that is most adaptable to change." – Charles Darwin

ChatGPT's AI has a lot of non-obvious potential consequences for internet companies, and in this section of our report, we leverage our Microsoft (MSFT) and Google (GOOGL) coverages to get a better grasp on them. As Microsoft sell-side analysts, we were granted access to ChatGPT + Bing (thank you, Microsoft!), which provided great insights. Let's start with advertising and e-commerce.

Advertising and E-Commerce

We would be more worried about sellers than about ad platforms; watch out for companies that depend heavily on search traffic.

In our view, Google Search is the internet's largest "eyeball workflow orchestrator". In other words, a large part of the entire internet's traffic originates from Google's search platform as it is today. However, if ChatGPT captures even part of this "initial flow", it could drastically alter this workflow and therefore traffic patterns. Companies that depend heavily on Google to build up their traffic (companies like BKNG, EXPE, TRIP, some e-commerce players, etc.) face the largest risk, in our view. Let's work through an example.

Let's say we are interested in taking a vacation in Europe but are undecided about where to go. We have three ways to begin exploring this question: going to Google Search, going to ChatGPT, or going to BinGPT (Bing powered by ChatGPT) – see the respective figures below.

ROUTE 1: Google Search

I. If we go to Google Search and type in "What's the best trip to do in Europe?", we get some websites telling us about destinations (all sponsored ads). We decide to click on the first one.

II. We are routed to a trip to Venice. If we scroll down, we are taken to some sightseeing portals with tickets available for purchase (this is how the website/advertiser makes money!).

ROUTE 2: ChatGPT…and then Google Search

I. If we go instead to ChatGPT and ask the same question, "What's the best trip to do in Europe?", ChatGPT returns a completely different answer. It basically gives us some options, grouped into categories with names such as the "The Classic European Tour" or "Mediterranean Adventure".

II. We decide that "The Classic European Tour" is our best option. But no links have been provided. It's not intuitive what we should do next.

III. We then decide to go to Google Search and run the query "The Classic European Tour". This query returns a completely different set of links – all from *different* advertisers than those we encountered in Route 1.

ROUTE 3: ChatGPT + Bing (which some are calling "BinGPT")

I. We ask Bing's Chat function (powered by ChatGPT) the exact same question ("What's the best trip to do in Europe?").

I. It provides a completely different answer (and in our case, the answer was in Portuguese; even after we had set Bing to English mode and asked the question in English, we got answers in Portuguese – probably based on a different default option we couldn't find).

II. Interestingly, BinGPT's answer is based on travel/vacation websites such as Earthtrekkers, US News Travel and TripAdvisor. The answer is basically a list of cities to visit in Europe – starting with Paris.

III. Also, the answer includes "travel to Europe on a budget", and offers some qualified alternatives (despite the fact that we never added "on a budget" to our search).

IV. BinGPT finishes up with a question for us ("What is your favorite destination in Europe?"), offering some default answers for me to click on (e.g., "Tell me more about London and Paris," or "What are some other places to visit in Europe?").

V. Finally, the quoted websites (Earthtrekkers, Travel.usnews.com and tripadvisor.com) are listed in a "learn more" row where we can click on and access these websites.

Route 1 – Google Search: "What is the best trip to do in Europe?"

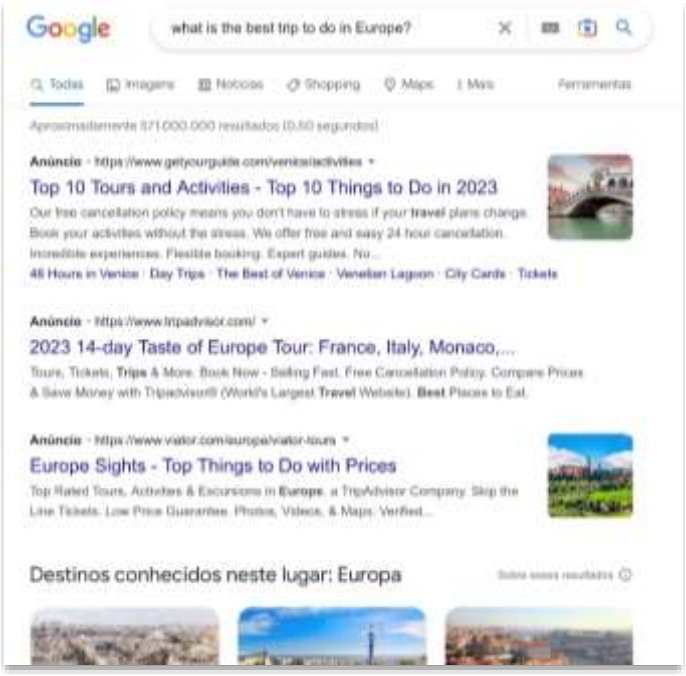

Source: Google Search

Route 1 – Google Search: If we click on the first link, we're directed to a trip to Venice and a choice of sightseeing destinations that we can buy tickets for

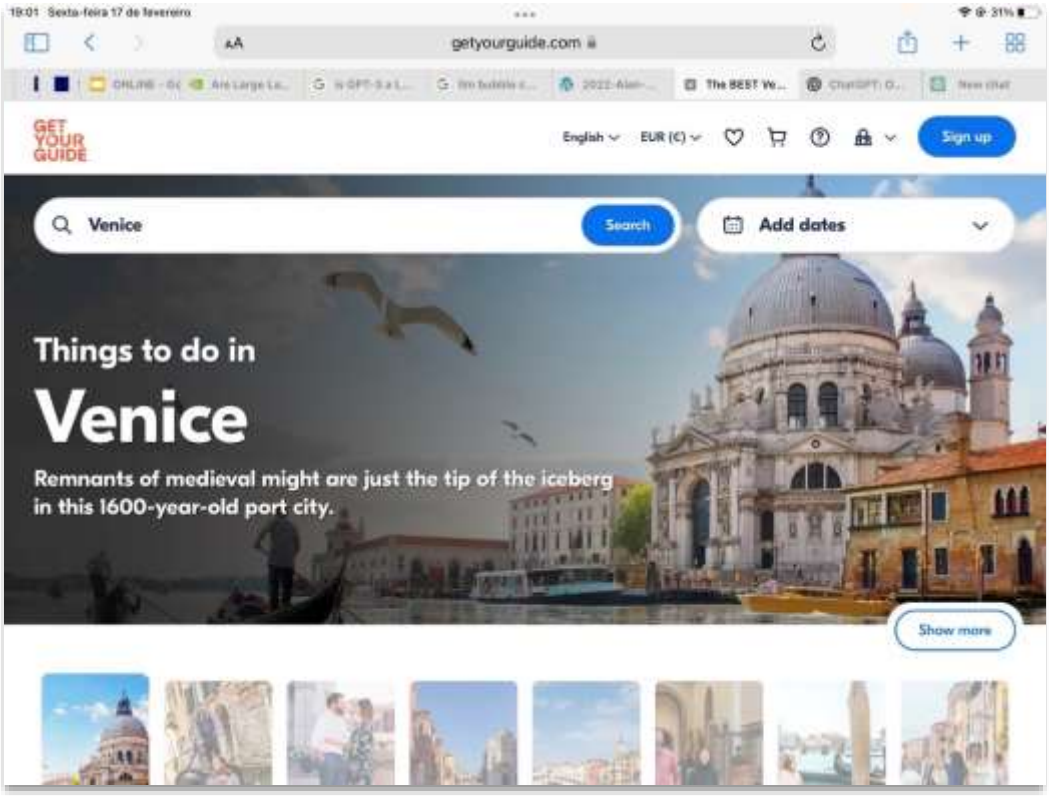

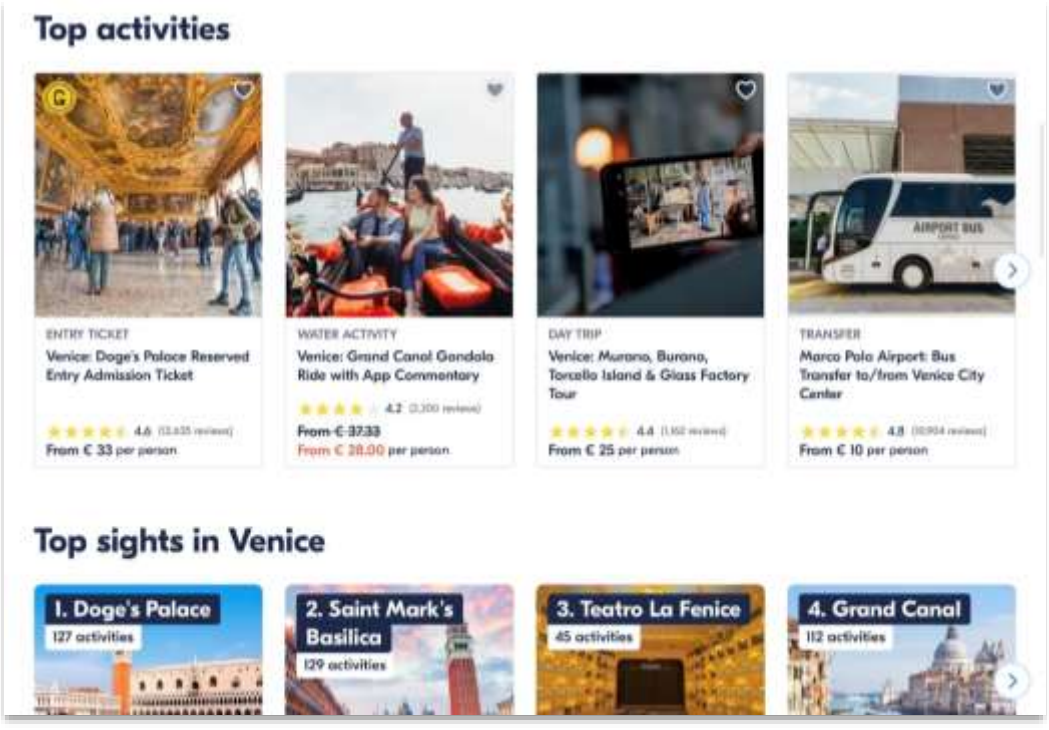

Source: Getyourguide.com

Route 2 – ChatGPT: “What is the best trip to do in Europe?”

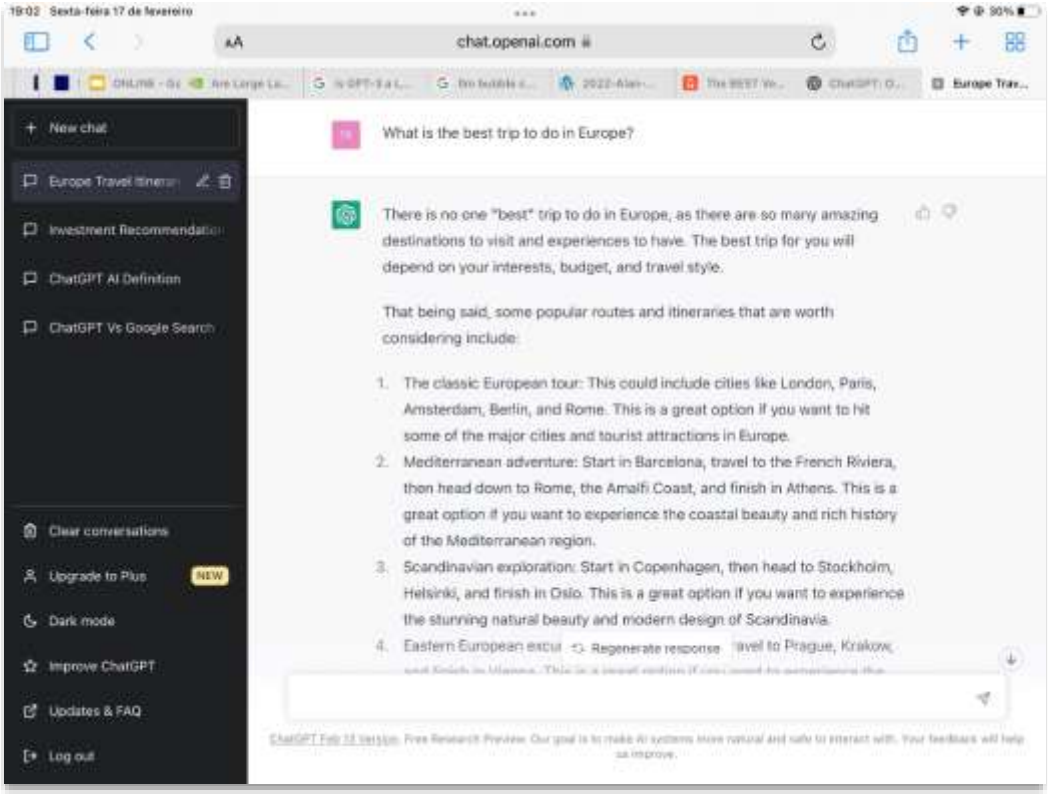

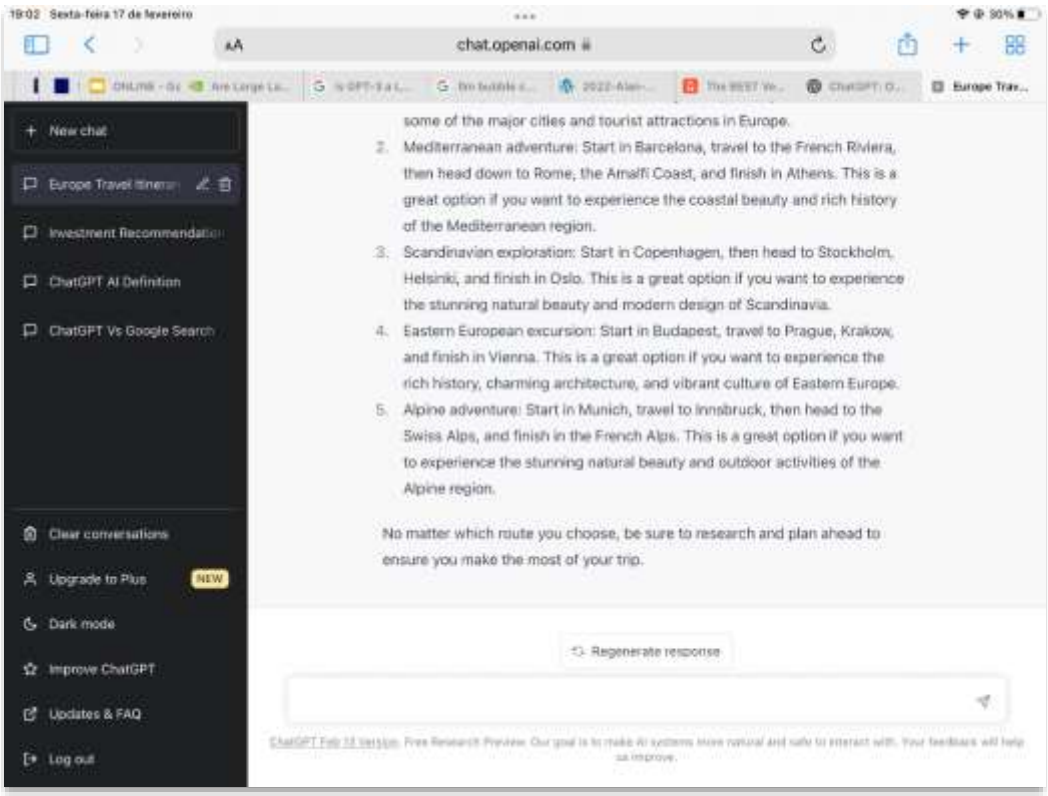

Source: ChatGPT

Route 2 – Google Search: if we query “The classic European tour”, we get completely different websites than we did in our Route 1 Google Search

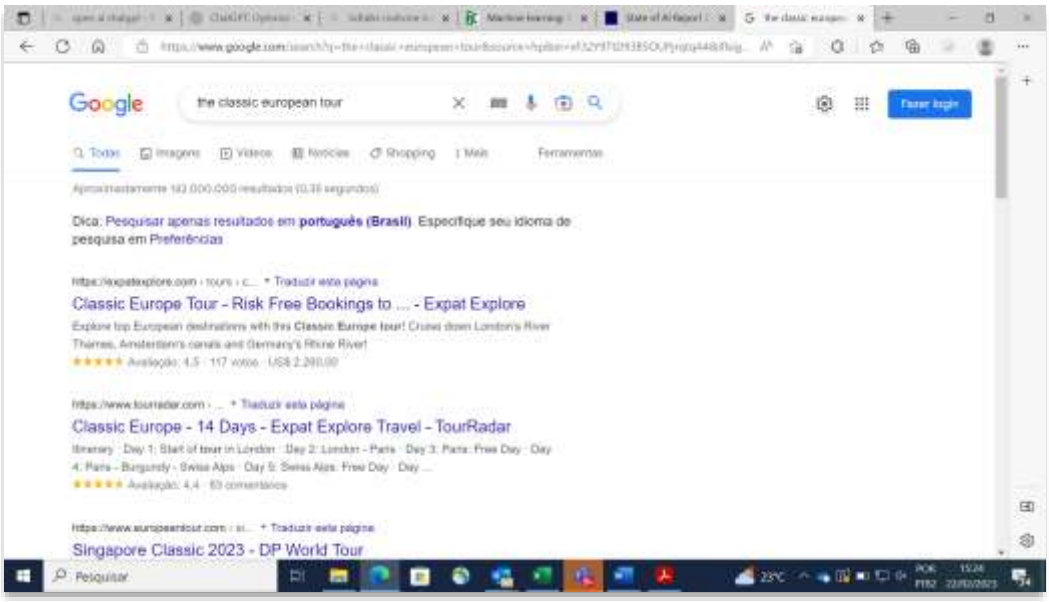

Source: Google Search

Route 3 – BinGPT: “What is the best trip to do in Europe?”

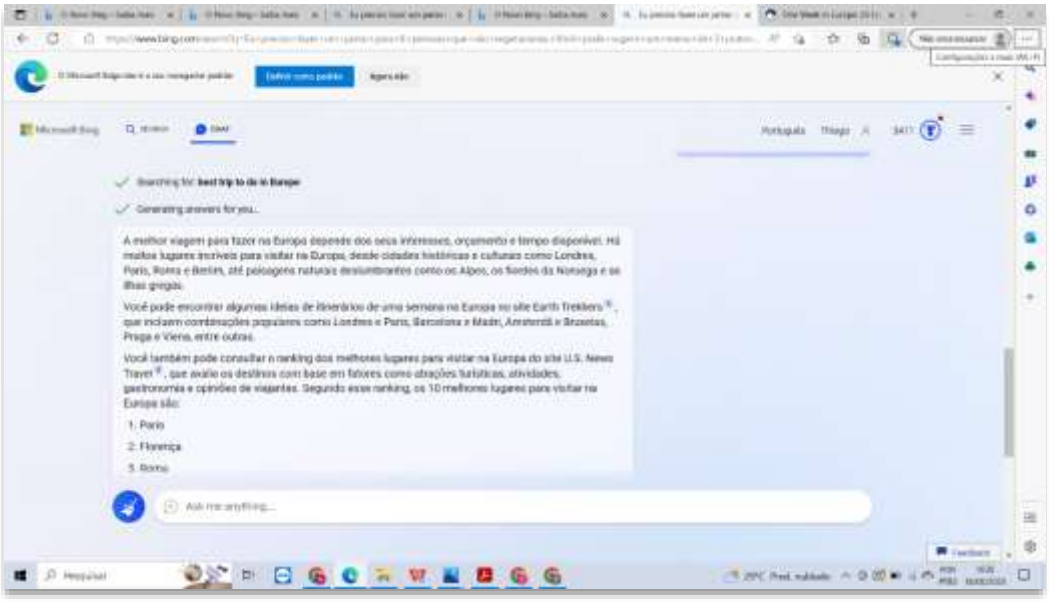

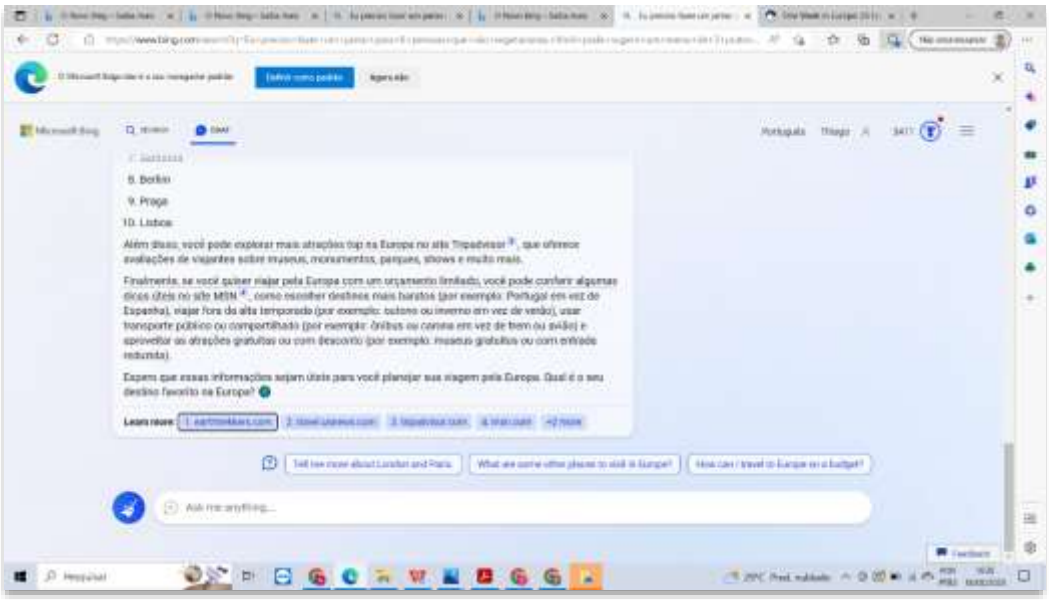

Source: Microsoft Bing (powered by ChatGPT)

Here are some of our takeaways from the experiment we’ve just done on the three routes.

I. **ChatGPT looks much less monetizable for advertising, at least for now.**

II. **Each of the three routes led to completely different websites**. At first, this may look useless, but think about it: companies spent years building SEO strategies to optimize traffic on their websites (backed on AdWords, for example). This means that if some people decide to take Route 2 rather than Route 1, for example, the entire pattern of internet traffic may change, which could lead to winners and losers.

- In our example, Earthtrekkers and Tripadvisor were winners and Getyourguide was a loser.

I. **We could argue that BinGPT’s answer is more “interactive”, which could lead to a higher CPM being paid by these advertisers (Earthtrekkers, Tripadvisor, etc.). Our understanding is that this is exactly MSFT’s strategy with this product.**

II. **As a consequence of the above, advertising market CPMs could change, which could have economic consequences for a number of companies that depend on search traffic.**

For advertising/e-commerce, our main conclusion from this work is that there could be winners and losers, with the most drastic changes (good or bad) affecting companies that depend heavily on search to build traffic. In our Brazilian coverage universe (Tech), there are no obvious winners and losers related to this specific issue, but there could be in other Brazilian sectors.

Online Education

Watch out for disruption in commoditized courses/databases for low-income students.

Maybe the best example of the potential disruption ChatGPT could cause is in the education sector. We've heard shocking news reports about the capabilities of ChatGPT in academic environments:

- ChatGPT passed a Wharton Business School MBA exam (Wharton is the business school at the University of Pennsylvania, an Ivy League university).
- Teenagers built a machine, based on a 3D printer, that could transcribe ChatGPT answers onto paper using human cursive writing, meaning that they could cheat on school tests.

ChatGPT Passed a Wharton Business School Exam

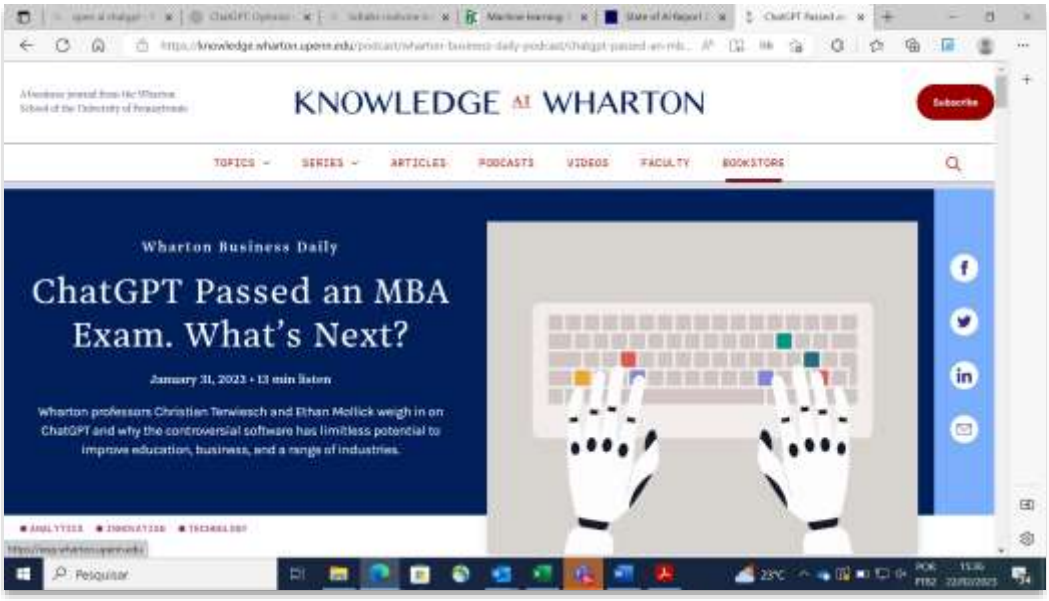

Source: Knowledge At Wharton

ChatGPT + 3D Printer Writing Someone's Homework

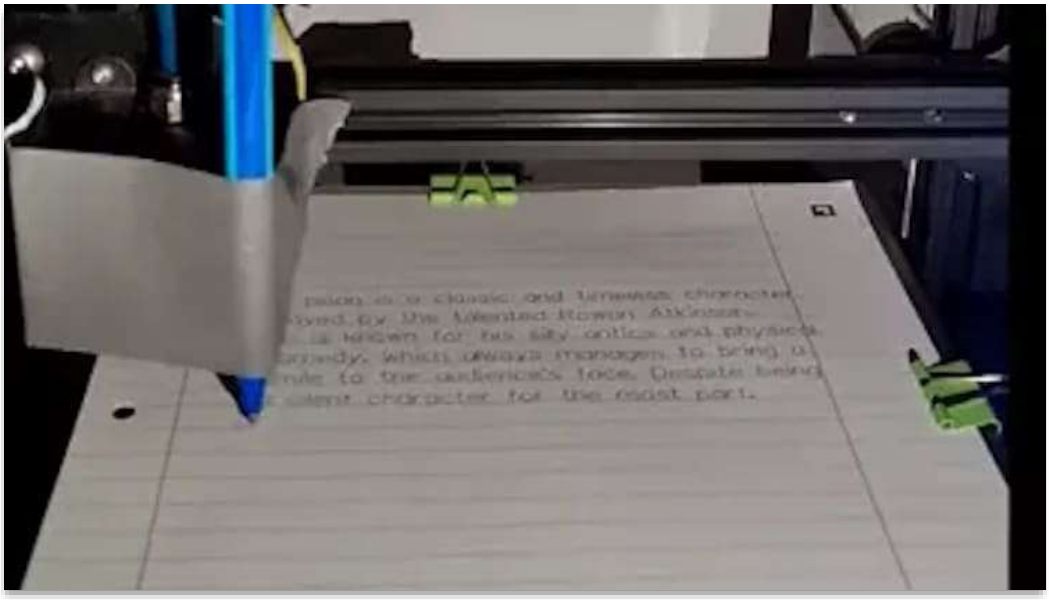

Source: My Modern Met

While the business responses to these new phenomena are still being devised, we already saw an impact in some education companies' 4Q22 results. Chegg, an online education program that offers solutions to help college students struggling with homework, tests, etc. (through a massive database of homework and exams), provided weak guidance for 2023, and some analysts blamed that on ChatGPT. While management said that ChatGPT didn't have any "noticeable" impact, they acknowledged that there is a 6% overlap between what ChatGPT and Chegg can provide. CHGG tumbled 20% on this earnings outlook.

CHGG Earnings: 6% Overlap Between ChatGPT's and Chegg's Offerings; Stock Tumbled 20% on Weak Guidance

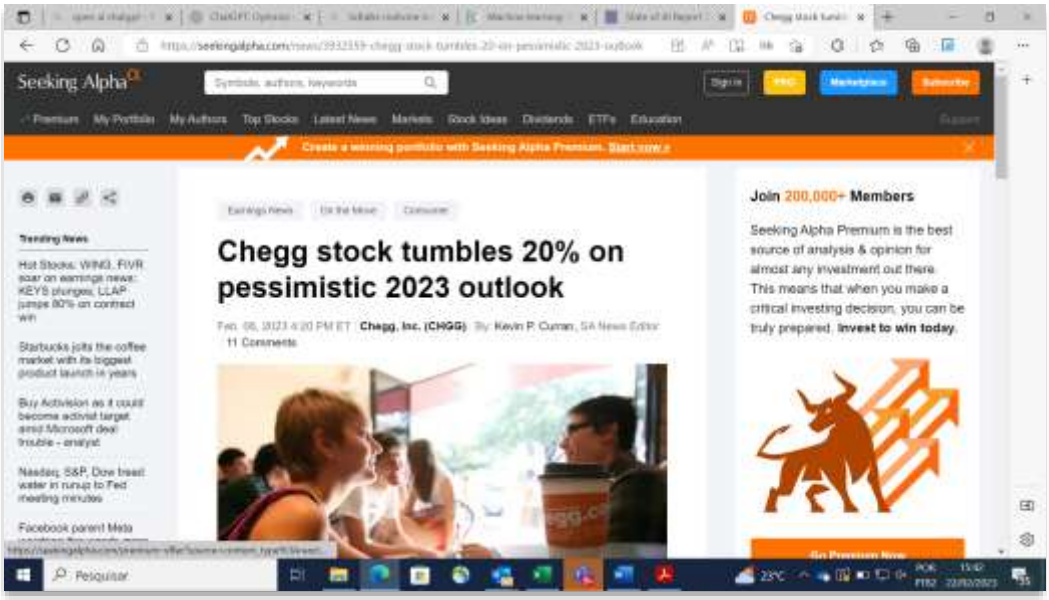

Source: Seeking Alpha

THOUGHTS ON THE BIG TECHS: THE WINNERS, THE LOSERS AND SOME MATH

SECTION 5

In our U.S. coverage universe, ChatGPT certainly has important implications for the big techs. In this section, we analyze each case.

Microsoft and Google (1): A Market Share Battle

One has a lot to win, and the other a lot to lose.

Given that MSFT and GOOGL are competing directly for search ad dollars, it's hard to work on one and not the other.

We believe that the best takeaways for both companies come from some insights we had during our trip to Europe a few weeks ago, when we visited a few long-only (LO) investors. In general, MSFT and GOOGL are the two big techs that LOs see as compounders/GARP (META looks like a trading call, AMZN is hard to invest in given its very high P/E, and AAPL doesn't grow much). With MSFT attacking GOOGL, this is how some investors have been simplifying the stories:

- MSFT is a consolidated cloud player which is now entering the search ad market (with decent odds of success). It should continue to grow strongly in cloud, and it has a lot to gain if it's also successful in search. There are other investment themes for MSFT, but they are less relevant.
- GOOGL is a consolidated search ad player which has entered the cloud market. While it should grow a lot in cloud, it will be hard for it to become larger than MSFT and AWS. Also, it has a lot to lose in search if it's unsuccessful in cloud. There are also other investment themes for GOOGL (YT, Waymo, etc.), but they are less relevant.

While long-term investors see these two stories as much more similar now, their pre-ChatGPT valuations were not too different. As shown in the table below, the three-year IRR for both MSFT and GOOGL was ~10% on February 7 (the day before ChatGPT launched in Bing).

The rationale was very simple: if the two companies were becoming more similar, if their IRRs were not too different, and if one had a lot to lose and the other a lot to win on the same theme, then investors should SELL GOOGL and BUY MSFT (all else being equal). This impression was augmented by a lack of understanding of and expertise in ChatGPT – there are so many generalists in these stocks that there's just too many investors with a limited capacity to evaluate the odds of success for ChatGPT + Bing versus Google Search.

Indeed, we believe that this outcome (selling GOOGL and buying MSFT) is exactly what happened, and it led to MSFT outperforming GOOGL by 11% over a mere 20 days.

MSFT's and GOOGL's Three-Year IRRs Were Not So Different Prior to ChatGPT + Bing

Ticker	Metric	Entry Multiple	CAGR	Exit Multiple	IRR
MSFT	P/E	23.4x	20.1%	18.0x	10.0%
GOOGL	P/E	23.8x	20.3%	18.0x	9.6%

Source: Bloomberg, Itaú BBA

This Led to a Sell-Off of GOOGL, With Investors Buying MSFT; a MSFT vs. GOOGL Call Would Have Outperformed by ~11% From February 7 to February 27

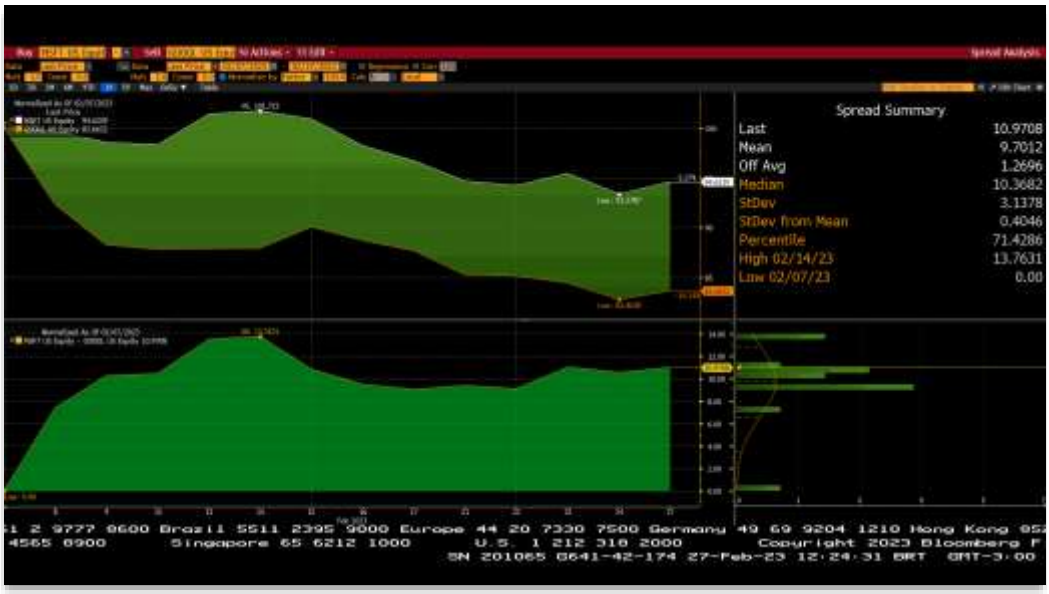

Source: Bloomberg, Itaú BBA

Another exercise we can conduct is to evaluate how much market share GOOGL will lose and MSFT will win, according to the market. Assuming that 100% of the stock price change was related to a market share loss for GOOGL as a result of ChatGPT (and BinGPT), we find that the market is pricing in a 23% market share loss to other players. Here are the assumptions behind this number:

- EBIT margin of 37% (same for Google Services in 2022)
- Tax rate of 20%
- Fair P/E of 18x

If we were to judge the impact on Google Search's market share over the past two months, we would call the stock price movement an overreaction. According to Similarweb, Google Search traffic has remained strong, hitting more than 88 billion visits per month in January (the highest rate in at least six months). But the problem with this analysis is that ChatGPT had a *very* strong entry – more than 100 million MAUs in less than a month.

Assuming That 100% of GOOGL's Underperformance From February 7 to February 27 Is Related to Market Share Losses, Market Values Imply That It Will Lose 25% of the Market to MSFT or ChatGPT

USD Mn	
Google Search Revs	162,450
Market Share	90%
Total Market Size	180,500
Market Share Lost	23%
Revenues Lost	40,726
EBIT Margin	37%
Potential EBIT	15,069
Tax Rate	20%
Net Income	12,055
P/E	18.0x
Market Cap Lost	216,989
GOOGL Market Cap	
Feb 7th	1,393,745
Feb 23rd	1,176,753
Difference	216,992

Source: Company Reports, Itaú BBA

GOOGL's Search Traffic Remains Strong, While ChatGPT is Growing Very Fast and Bing Remains Weak

Source: Similarweb

ChatGPT Reached 100 Million MAUs in a Mere Month – Far Faster Than Any Other Internet Company in History

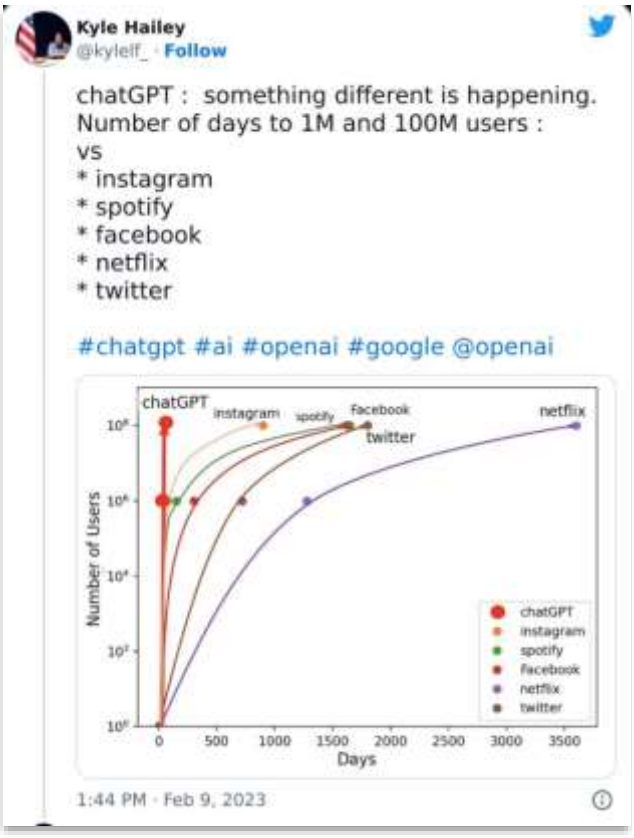

Source: Kyle Hailey (AIwave.art founder)

Microsoft and Google (2): The Cost Question

The extent of Google's profitability erosion will depend on the how aggressively competitors move to enter/gain share on search. We see a 5%-60% impact on EBIT, depending on the share of queries that move into LLM, which in turn depends on the competition.

The other debate on MSFT and GOOGL concerns margins. Alphabet's Chairman, John Hennessy, said that search queries running in LLMs cost ~10x more than standard queries (For tech giants, AI like Bing and Bard poses billion-dollar search problem | Reuters). Many analysts have tried to run exercises to forecast the impact of this development on the margins of both MSFT and GOOGL.

We believe that the most important discussion concerns GOOGL. We ran our exercise using the following steps:

I. Sam Altman (the founder of Open AI, which owns ChatGPT) has tweeted that the average cost per query for ChatGPT is "in the single-digit cents". Therefore, we initially narrowed the range to USD 0.01-0.09.

II. We assumed that LLM search costs are 10x higher, as GOOGL's chairman said, which implies that the incremental difference in cost per query would remain in a single-digit order of magnitude.

III. However, the USD 0.01-0.09 range is too large. We had to narrow this query down.

IV. What we did then was to ballpark the cost per query based on some assumptions:

a) According to Open AI, the number of ChatGPT words per token is 0.75.

b) According to the University of Maryland, it takes 6 milliseconds per GPU (A100) to process one token in a 3-billion-parameter model (or 166 tokens processed per second). If we

linearize this (an assumption), there would be three tokens processed per second per GPU, or 23 tokens per 8x A100 (a typical NVIDIA processor). This translates to 17 words per second per processor.

c) If we assume 75 words per query, 0.2 queries would be processed per second per processor.

d) Finally, the Azure prices per 8xA100 vary between USD 5 and USD 15 per hour. Applying this to the numbers above, the cost would be USD 0.02/query.

e) We applied this number to the number of queries per day. According to MSFT, there are 10 billion queries per day in the search market. We then assumed the different shares of queries that will run on LLMs versus traditional platforms.

The conclusion for GOOGL is that we could see an EBIT impact of up to 60%! And even if the price per 8xA100 is at the lower bound, there could be a material change (~20%) for 100% of the queries that have migrated. While many analysts have dismissed the possibility of a sudden impact, we believe that the main driver of both the magnitude and the timing of the coming drop is competition. If all the companies that have LLMs with parameters exceeding 100 billion (MSFT, META, NVIDIA, BIDU, Hugging Face, etc.) decide to try their luck in the search market, we could see GOOGL forced to move queries into LLMs sooner than later!

For MSFT, we believe that the impact is minimal, given that advertising contributes only 6% of its revenues. If MSFT's margins fall more than expected, it likely means that the company was successful in its battle with Google.

Cost-per-Query Analysis: We Believe It Costs USD 0.02/Query to Process Using LLMs

Cost per Query Analysis			
ChatGPT Words / Token	0.75	a	Open AI
Tokens / Second per GPU	3	b	U of Maryland
Tokens / Second per 8xA100	23	c = 8 x b	
Words / Second per 8xA100	17	d = a x c	
Words / Query	75	e	
Queries / Second per 8xA100	0.2	f = d / e	
8xA100 Pricing (USD/Hour)	15	g	Avg Azure
8xA100 Pricing (USD/Second)	0.004	h = g / 3600	
Cost Per Query (USD/Query)	**0.02**		

Source: Open AI, University of Maryland, Azure, Itaú BBA

GOOGL Margin Impact – Incremental Costs

GOOGL Incremental Costs (USDmn)	
# of Daily Queries (Total)	10,000
GOOGL Market Share	90%
# of Daily Queries	9,000
# of Yearly Queries	3,285,000
Share of Queries Through LLM	100%
# of Queries Under LLM	3,285,000
Incremental Cost per Query	0.02
Incremental Costs	**59,883**

Source: Company Reports, Itaú BBA

Impact on GOOGL: 5%-60% of 2024 EBIT, Depending Mainly on the Share of Queries That Move to LLM

Share of Queries Into LLM	8xA100 Pricing (USD/Hour)		
	5	10	15
20%	4%	8%	12%
40%	8%	16%	24%
60%	12%	24%	37%
80%	16%	33%	49%
100%	20%	41%	61%

Source: Itaú BBA

META

The AI revamp could pay off more than expected; could META join the generative AI battle?

While many analysts and investors have criticized META's focus on the metaverse, with its significant associated costs, we believe that it could bring positive secondary benefits that may become more important than our initial expectations.

In the State of AI report, Benaich and Hogarth show that META's data centers now have the largest number of AI chips in the world. Indeed, the company even referred to "AI and non-AI" workloads in its latest conference call. While the metaverse may have been the greatest reason (or excuse) for heavy investments in AI, which hit META's P&L (and stock) during 2021-22, there are certainly many potential benefits of having such advanced technology and thus the capacity to use generative AI.

While it's hard to think of potential benefits for the P&L, we wouldn't be surprised to see META competing with TikTok with some sort of generative-AI-based feature, or even the creation of search alternatives within META's properties. We note that some TikTokers prefer to run searches within TikTok (rather than Google), depending on the use case, as the number of clicks/steps needed to get results is lower in the former.

It's also important to remember that META already has an LLM (see Section 3), which positions it well for a move in this direction. Interestingly, META announced just last Friday that it is launching a new LLM called Llama (Meta Introducing AI Large Language Model Called 'Llama' | Bloomberg)

META Has the Largest NVIDIA A100 GPU Count

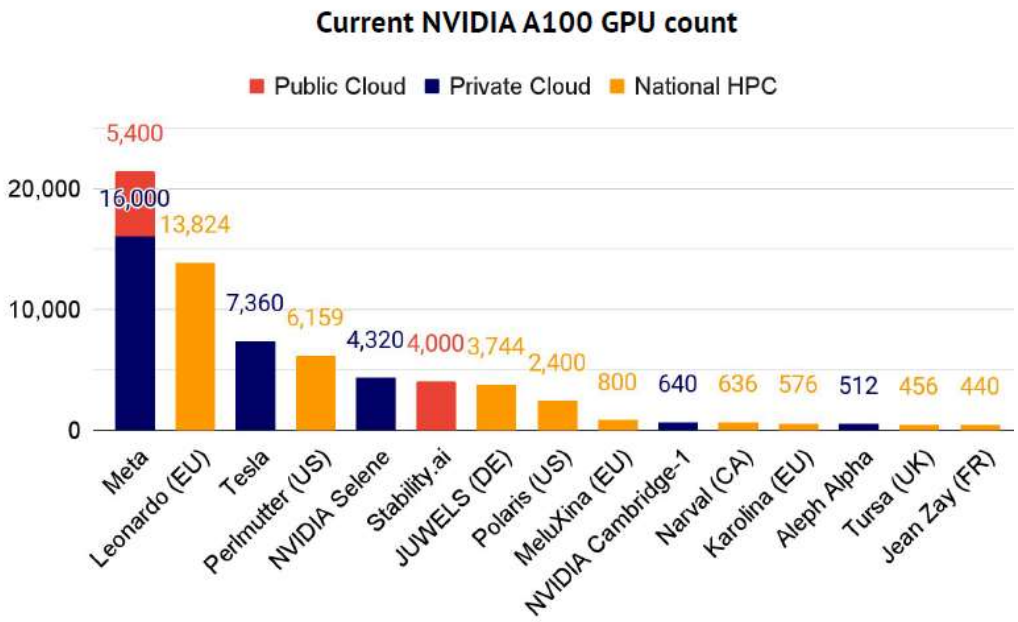

Source: State of AI Report 2022

Amazon

As 1P advertising gets more important than ever, AMZN looks behind its peers on Generative AI.

Many investors were disappointed that AMZN appears to have fallen behind in the generative AI battle. Even META appears to be ahead of AMZN (after the former's massive investments in AI). We wouldn't be surprised to see AMZN getting into this business at some point in the future.

Having said this, we believe that the same risks described in Section 4 apply to Amazon. The company depends on search traffic, and categories in which this is important could be at higher risk.

What mitigates these risks is that AMZN fields a large share of initial queries, particularly in the United States. The latest data we could find for this is from 2018, when AMZN had a market share of 47%, compared with GOOGL's 35%, for initial product search queries.

AMZN Has the Largest Share of Initial Product Searches vs. GOOGL

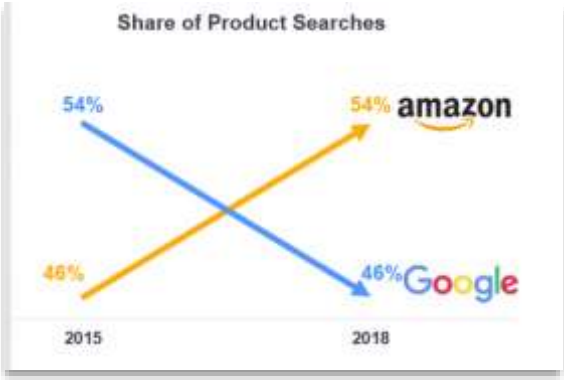

Source: Jumpstart

Apple: A Potential Long-Term Winner

The battle for default search could get more intense, but it hasn't happened yet. Whenever it happens, Apple could eventually monetize even better demand for its properties.

We believe that Apple could be one of the main winners in this story. If Microsoft focuses even more on getting a share in search, we could see Bing bidding for the iOS search default engine, potentially further increasing the USD 15-20 billion that AAPL earns from Google for these rights today. And even other LLM owners might make such a bid in the future.

While we believe that this is a real possibility, we are skeptical about such a move in the short term. We believe that MSFT is now 100%-focused on proof-testing Bing + ChatGPT on desktops, where MSFT has 1.4 billion users in total. There is such a critical mass in this channel that it doesn't really make sense to aggressively move into mobile when the product is not proven yet. And MSFT seems far ahead of its peers on this front, at least for now.

As a reminder, we estimate that every additional USD 1 billion in default search revenues should translate into a 1% higher EBIT for Apple (assuming zero incremental costs).

SEMIS ARE THE CLEAR WINNERS; NVIDIA WOULD BE OUR PICK

SECTION 6

In this last section, we will briefly address what, in our view, are likely to be the biggest winners from the generative AI story: semis. As the world increasingly digitizes, demand for semis should continue to be strong. Indeed, on its earnings call last week, NVIDIA mentioned that demand for chips has increased notably in the two months since ChatGPT was released (leading to a 14% pop in the stock the day after).

Still, we are skeptical that all types of semis will benefit from generative AI. The most obvious winner to us, once again, is GPU, and inside this technology, we would favor NVIDIA. Let's go over some reasons why we believe that this is the right call.

I. NVIDIA chips are the most popular in AI research papers, by a massive and increasing margin (source: State of Ai report).

a) As shown in the figures below, GPUs are 131x more commonly used than ASICs; 90x more than chips from the startups Graphcore, Habana, Cerebras, Sambanova and Cambricon combined(!); 78x more than Google's TPUs; and 23x more than FPGAs.

I. NVIDIA has over 3 million developers on its platform (it's hard to think of what a second-place player could be, when it comes to GPUs).

II. Going forward, we believe that CUDA, GPU's proxy for an OS/coding framework in CPU, will continue attracting more developers to NVIDIA's ecosystem than to its competitors (like AMD).

NVIDIA's Chips Are the Most Popular in AI Research Papers by a Massive Margin

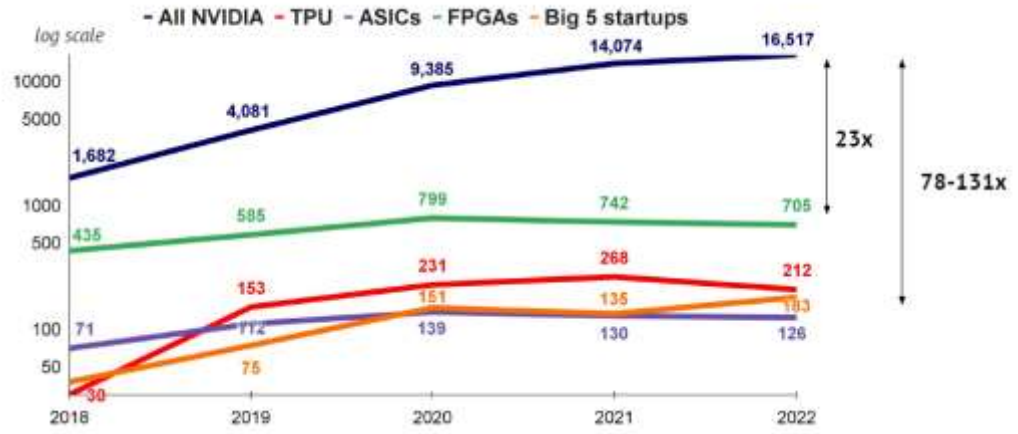

Source: State of Ai Report 2022

Number of AI Papers Citing Use of Specific Nvidia Cards

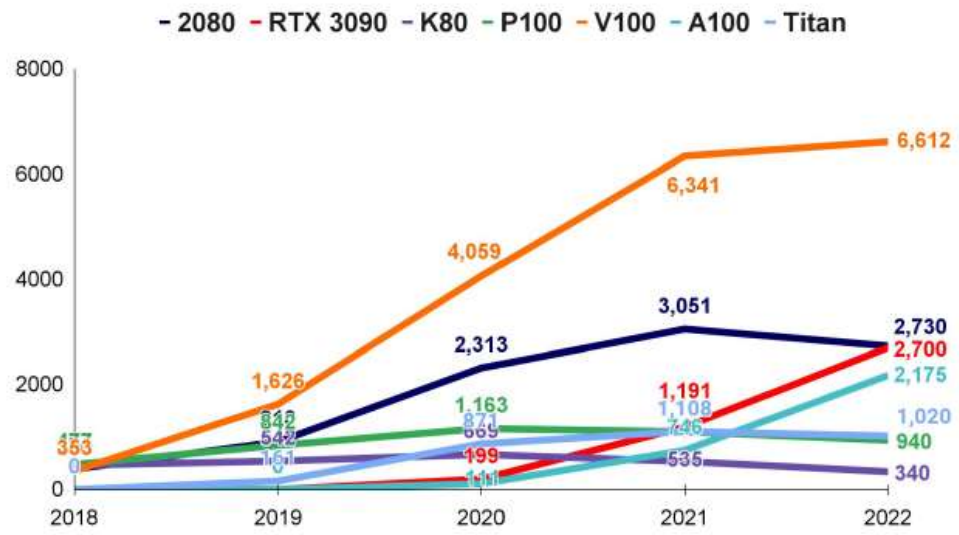

Source: State of Ai Report 2022

Number of AI Papers Citing Use of Specific AI Chip Startups

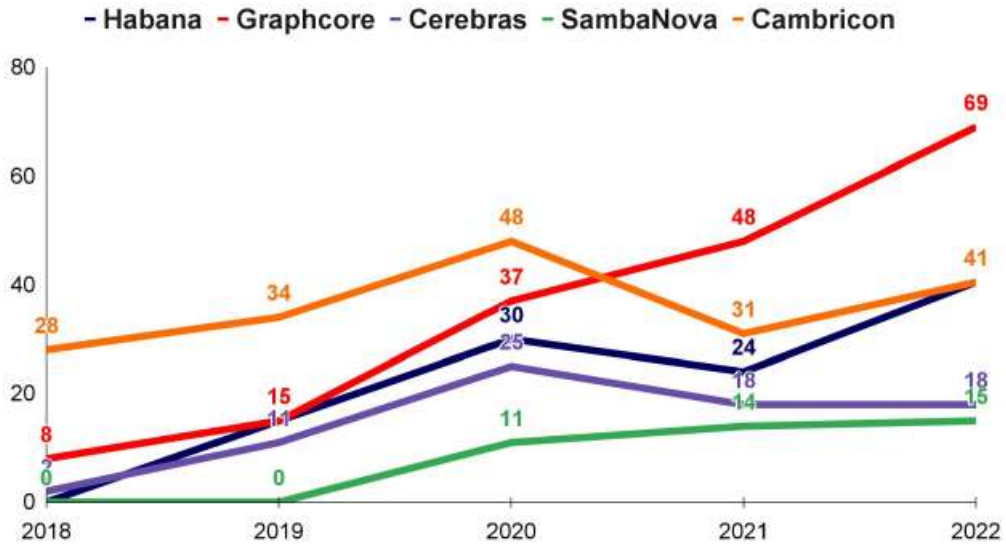

Source: State of Ai Report 2022

DISCLAIMER

Itaú BBA is a registered trademark used by Itaú Unibanco S.A.

Ratings: Definitions, Dispersion and Banking Relations

Ratings [1]	Definition [2]	Coverage [3]	Banking Relation [4]
Outperform	The analyst expects the stock to perform better than the sector average.	62%	62%
Market Perform	The analyst expects the stock to perform in line with the market average.	33%	31%
Underperform	The analyst expects the stock to perform below the sector average.	5%	5%

1. The ratings employed in this document (Outperform, Market Perform, and Underperform) basically correspond to Purchase, Hold, and Sell, respectively.
2. The ratings represent the analyst's assessment of the medium-term share price performance relative to the market average. These ratings may be reviewed by the analyst based on new developments or simply due to variations in share prices (such changes may occur at any time). Companies are grouped into sectors, based on similar characteristics. Sectors: (i) Banks and Financial Services; (ii) Consumer Goods & Retail + Food & Beverage (iii) Healthcare + Education; (iv) Steel & Mining + Pulp & Paper; (v) Oil, Gas & Petrochemicals + Agribusiness; (vi) Real Estate & Construction; (vii) Telecommunications, Media and Technology; (viii) Transportation, Capital Goods and Logistics; (ix) Public Utility Services; and (x) Strategy.
3. Percentage of companies covered by Itaú Unibanco S.A. in this rating category.
4. Percentage of companies included in this rating category that were provided services by Itaú Unibanco S.A. or any of its affiliated companies.

Material Information

1. This report was prepared by Itaú Unibanco, a company regulated by the Brazilian Securities and Exchange Commission (CVM), and distributed by Itaú Unibanco, Itaú Corretora de Valores S.A. and Itaú BBA. Itaú BBA is a registered trademark used by Itaú Unibanco.
2. This report does not constitute an offer to buy or sell any financial instrument or to participate in any particular trading strategy under any jurisdiction. The information herein was deemed reasonable on the date of publication and was obtained from reliable public sources. Neither Itaú Unibanco nor any of its affiliated companies ("Grupo Itaú Unibanco") ensure or guarantee, either expressly or implicitly, that the information contained herein is accurate or complete. Furthermore, this report is not a complete or summarized representation of investment strategies, markets or developments thereof. The opinions, estimates and/or forecasts provided in this report reflect the current opinion of the analyst responsible for the content herein on the date of publication and are, therefore, subject to change without notice. The prices and availability of the financial instruments are merely indicative and subject to change beyond the control of Itaú Unibanco. The prices used for the production of this report were obtained as of the close of the relevant market. Neither Itaú Unibanco nor any of its affiliated companies are obligated to update, amend or otherwise alter this report, or to inform readers of any changes in its content, except upon termination of coverage of the securities issuers discussed herein.
3. **The analyst responsible for preparing this document, highlighted in bold, hereby certifies that all opinions expressed in this report accurately, solely and exclusively reflect his/her personal views and opinions regarding all of the issuers and securities analyzed herein, provided in this document independently and autonomously, including from Itaú Unibanco, Itaú Corretora, Itaú BBA, and any of their affiliated companies. Whereas the personal opinions of the analyst may diverge, Itaú Unibanco, Itaú Corretora, Itaú BBA and/or any of their affiliated companies may have published or eventually publish other reports that are inconsistent and/or reach different conclusions than those provided herein.**
4. The analyst responsible for preparing this report is not registered and/or not qualified as a survey analyst at the NYSE or FINRA, nor is such analyst in any way associated with Itaú BBA USA Securities, Inc. and is, therefore, not subject to the provisions of Rule 2242 on communications with researched companies, public appearances and transactions involving securities held in a research analyst account.
5. Part of the analysts' compensation is determined based on the total top line of Itaú Unibanco and some of its affiliated companies, including revenues arising from investment services and investment bank services. Nonetheless, the analyst responsible for the content in this report hereby certifies that no part of his/her compensation is or will be directly or indirectly conditioned to any specific recommendations or opinions herein.
6. The financial instruments and strategies discussed in this document may not be suitable for all investors and certain investors may not be eligible to purchase or participate in some or all of them. This report does not take under consideration the specific investment objectives, financial situation or needs of any particular investor. Investors who intend to purchase or trade the securities covered in this report must seek out the applicable information and documents in order to decide whether to invest in such securities. Investors must independently seek out financial, accounting, legal, economic and market guidance, based on their personal characteristics, before making any investment decision regarding the securities of the issuers analyzed in this report. Each investor must make the final investment decision based on a range of related risks, fees and commissions. In the event that a financial instrument is expressed in currencies other than the one used by the investor, exchange rate variations may adversely affect the price, value or profitability. Financial instrument yields may vary, ultimately increasing or decreasing the price or value of financial instruments, either directly or indirectly. Past performance is not necessarily indicative of future results, and this report does not ensure or guarantee, either expressly or implicitly, any possible future performance or any other aspect thereof. Itaú Unibanco and its affiliated companies may not be held liable for eventual losses, either direct or indirect, arising from the use of this report or its content. Upon using the content herein, investors undertake to irrevocably and irreparably hold Itaú Unibanco and/or any of its affiliated companies harmless against any claims, complaints, and/or losses.
7. This report may not be reproduced or redistributed to third parties, either in whole or in part, for any reason whatsoever, without prior written consent from Itaú Unibanco. This report and its contents are confidential.
8. In accordance with the rules of the Brazilian Securities and Exchange Commission, the analysts responsible for this report have described scenarios of potential conflict of interest in the "Relevant Information" chart below.
9. Itaú Unibanco is not party to an agreement with the issuer(s) relating to the production of this report.
10. The description of the internal organizational and administrative arrangements and of the information barriers in place for the prevention and avoidance of conflicts of interest with respect to the production of research reports are available at www.itaucorretora.com.br.
11. Material information regarding the proprietary models used by Itaú Unibanco are available at Itaú BBA Bloomberg page.
12. Other important regulatory information is available at https://disclosure-research.cloud.itau.com.br/home

Additional note for reports distributed in the (i) United Kingdom: This material has been prepared by Itaú Unibanco S.A. for informational purposes only and it does not constitute or should be construed as a proposal for request to enter into any financial instrument or to participate in any specific business trading strategy. The financial instruments discussed in this material may not be suitable for all investors and are directed solely at Eligible Counterparties and Professional Clients, as defined by the Financial Conduct Authority. This material does not take into consideration the objectives, financial situation or specific needs of any particular client. Clients must independently obtain financial, legal, tax, accounting, economic, credit, and market advice on an individual basis, based on their personal characteristics and objectives, prior to making any decision based on the information contained herein. By accessing this material, you confirm that you are aware of the laws in your jurisdiction relating to the provision and sale of financial service products. You acknowledge that this material contains proprietary information and you are required to keep this information confidential. Itaú BBA International plc (IBBAInt) exempts itself from any liability for any losses, whether direct or indirect, which may arise from the use of this material and is under no obligation to update the information contained in this material. Additionally, you confirm that you understand the risks related to the financial instruments discussed in this material. Due to international regulations, not all financial services/instruments may be available to all clients. You should be aware of and observe any such restrictions when considering a potential investment decisions. Past performance and forecast are not a reliable indicator of future results. The information contained herein has been obtained from internal and external sources, and is believed to be reliable as of the date in which this material was issued, however IBBAInt does not make any representation or warranty as to the completeness, reliability or accuracy of information obtained by third parties or public sources. Additional information relative to the financial products discussed in this material is available upon request. Itaú BBA International plc has its registered offices at 100 Liverpool Street, Level 2, London, EC2M 2AT, United Kingdom, and is authorized by the Prudential Regulation Authority and regulated by the Financial Conduct Authority and the Prudential Regulation Authority (FRN 575225). For any queries, please contact your relationship manager; **(ii) European Economic Area:** this report is distributed in the EEA by Itaú BBA Europe, S.A. ("IBBA Europe"), an affiliate of Itaú Unibanco with registered offices at Rua Tierno Galvan, Torre III, 11 Piso, Lisbon, Portugal, authorized as a credit institution by the Bank of Portugal (registration number 195), and as a financial intermediary by the Portuguese Securities Market Commission – CMVM (registration number 396). This material is informational only and shall not be construed as a proposal to trade with financial instruments, nor any kind of personal recommendation of advice. This material is directed solely at Professional Clients and Eligible Counterparties (as defined by Article 4, 1, (10) and Article 30 of Directive 2014/65/EU - MiFID II), who possess the experience, knowledge and expertise to make its own investment decisions and properly assess the risks that they incur. Relevant information regarding conflicts of interest contained in this material apply to IBBA Europe as an affiliate of Itaú Unibanco. IBBA Europe does not exercise discretion as to the selection of the research reports to disseminate and exempts itself from any liability for any losses, whether direct or indirect, which may arise from the use of this material and/or its content and has no obligation to update the information contained in this material. The financial instruments/services discussed in this material may not be available to all clients and/or in their respective jurisdictions. Additional information about the financial instruments discussed in this material are available upon request **(iii) USA:** Itaú BBA USA Securities, Inc., a FINRA registered Broker Dealer and an SIPC member, is responsible for distributing this report and the content herein. Any U.S. investor who receives this report and intends to trade any of the securities addressed herein must do so through Itaú BBA USA Securities, Inc. at 540 Madison Avenue, 24th Floor, New York, NY 10022; **(iv) Brazil:** Itaú Corretora de Valores S.A., a subsidiary of Itaú Unibanco S.A., is authorized by the Central Bank of Brazil and sanctioned by the CVM to distribute this report. Contact the Customer Service channel for more information: 4004-3131* (capital and metropolitan regions) or 0800-722-3131 (other locations) during business hours (BRT). If additional service is required after contacting the aforementioned channels, contact: 0800-570-0011 (business hours) or write to P.O. Box 67.600, São Paulo – SP, CEP (Postal Code) 03162-971. *(Local costs apply.)

Relevant Information – Analysts

Analysts	Disclosure Items				Sig
	1	2	3	4	
Thiago Alves Kapulskis					OK
Cristian Faria					OK

1. The securities analyst(s) involved in preparing this report are associated with individuals who work for the issuers addressed herein.
2. The securities analyst(s) spouse(s) or partner(s) hold, either directly or indirectly, on their own behalf or on behalf of third parties, stock and/or other securities discussed in this report.
3. The securities analyst(s), spouse(s) or partner(s) are directly or indirectly involved in the purchase, sale or intermediation of the securities discussed in this report.
4. The securities analyst(s), respective spouse(s) or partner(s) hold, either directly or indirectly, any financial interest related to the securities issuers analyzed in this report.