

Understanding Generative AI

A Glimpse Into the Power and Potential of Artificial Intelligence



Web: <https://www.omegavp.com>
Email: connect@omegavp.com
Tel: +1-650-446-9000
Social: [LinkedIn](#) | [Twitter](#) | [Forbes](#)

In 2020, Omega Venture Partners established the first US venture capital fund thematically focused on investing in AI / ML businesses at the early growth stage. Since then, we've seen a proliferation of AI utilization across software businesses. The availability of superior models and enhanced tools has solidified AI's presence across various industries. The release of ChatGPT in November 2022 helped make it evident to the world-at-large as to just how good AI has become.

Not coincidentally, in recent months we've been getting many questions about Generative AI from individuals and institutions seeking to better understand the category and opportunity.

After reading this report you will know things about Generative AI that you don't know now and something amongst the things you learn here will make a difference in your lives. Whether you are a business owner, investor, or an individual curious about the AI revolution, this report will equip you with a greater understanding of Generative AI, its implications, and ways to leverage it effectively in your professional and personal lives.

INTRODUCTION

ChatGPT's debut sparked an intense public and media discourse around AI's future, establishing the term 'Generative AI' in mainstream discussions. We believe that Generative AI holds the potential to transform content creation, enabling the rapid production of rich, tailored content. However, we also believe that its benefits extend beyond just 'generative' capabilities like generating marketing material or drafting essays.

Omega is working closely with portfolio companies in implementing real-world applications of Large Language Models, demonstrating the broad applicability of Generative AI beyond mimicking human composition. The shared feature across these applications is the employment of transformer-based language models — a staple in Generative AI endeavors. These models harbor significant potential, embodying an understanding of human language and knowledge in a manner that was unthinkable in the past.

We believe that Generative AI carries immense potential for a diverse array of applications, such as intelligent automation of business processes, boosting human learning and scientific advancements, knowledge management, and human-computer interaction. It enables individuals to generate custom content in a matter of seconds, signaling a shift in business paradigms. However, we maintain that the underlying technologies like Large Language Models (for text) and diffusion models (for images) enable a broader spectrum of applications beyond the 'generative' ones currently observed.

In this report, we will explore the implications and applications of Generative AI in a world that is continuously digitalizing and evolving. Our insights are drawn from our team's 20+ year immersion in Artificial Intelligence, frontline exposure to thousands of startups, and active engagement with our portfolio companies. We aim to illuminate the transformative power of this technology and how it's reshaping multiple industries. We will delve into the mechanisms of Generative AI, including the technological breakthroughs, like Large Language Models and Foundation Models, that have paved the way for its rise.

Read on to demystify Generative AI, understand its core mechanisms, and glimpse the horizon of its impact in your personal and professional realms.

AI AND MACHINE LEARNING

Generative AI is a class of Artificial Intelligence technology capable of creating diverse types of content, such as text, visuals, sound, and synthetic data. However, before we delve deeper into generative artificial intelligence, it's important to first establish a basic understanding of AI itself.

Often, we are asked about the meaning of Artificial Intelligence and the distinction between AI and machine learning (ML). Think of AI as a broad field, similar to biology. AI is a subset of computer science that focuses on creating intelligent systems that can reason, learn, and act independently. In essence, AI aims to build machines that can mimic human cognition and actions.

Machine learning is a subset of AI. It involves creating a system or program that trains a model using input data. This trained model can then generate meaningful predictions from new or unseen data that stems from the same distribution as the training data. Machine learning empowers computers to learn without explicit programming.

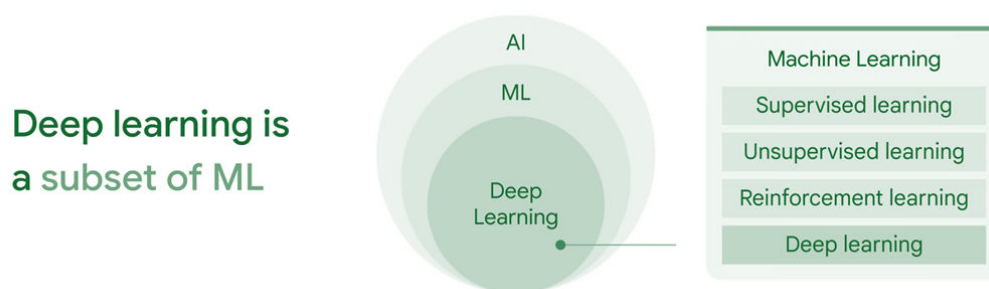
Two prevalent types of machine learning models are *unsupervised* and *supervised* ML models. The main differentiation between these two is that supervised models utilize labeled data. Labeled data refers to data tagged with a certain identifier, such as a name, a type, or a numerical value, whereas unlabeled data lacks such tags.

Consider a supervised model in the context of a bookstore owner. The owner possesses historical sales data concerning book genres and corresponding sales numbers. Using supervised learning, the model utilizes past data to forecast future values, such as predicting potential sales for different book genres.

In contrast, an example of an unsupervised model could involve a school principal aiming to understand the student population better. The principal might examine factors like students' grades and participation in extracurricular activities, grouping students to identify those potentially at risk of underperforming. Unsupervised models focus on discovering natural patterns or clusters within the data.

Two other types of machine learning that are worth being aware of are: semi-supervised learning and reinforcement learning. In summary, machine learning is AI that learns patterns from data. It includes:

- **Supervised Learning:** Where models learn from labeled data.
- **Unsupervised Learning:** Where models identify patterns and structures from unlabeled data.
- **Semi-Supervised Learning:** A combination of supervised and unsupervised learning (i.e., the model can be trained on a combination of labeled and unlabeled data).
- **Reinforcement Learning:** Where models learn to make decisions based on rewards.



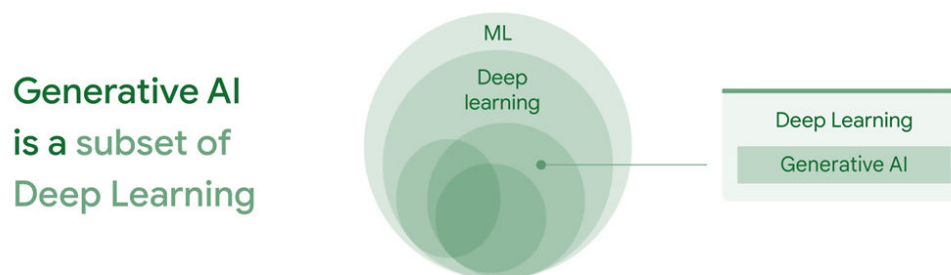
Deep Learning: Foundations for Generative AI

Now that we've covered the distinction between artificial intelligence and machine learning and shared with you an overview of the major flavors of machine learning, let's briefly explore where Deep Learning (DL) fits as a subset of machine learning methods.

While ML encompasses various methodologies, deep learning is a subset of ML that uses artificial *neural networks* with many layers ("deep" structures) to tackle more complex patterns than conventional ML approaches. These neural networks (loosely) draw inspiration from the human brain and consist of interconnected nodes or neurons capable of learning tasks by processing data and making predictions. Deep learning has been the key driver behind advances in image recognition, natural language processing, and other complex tasks.

Neural networks have the flexibility to utilize both labeled and unlabeled data (which is called 'semi-supervised learning,' as discussed above). In this setting, a neural network is trained on a limited amount of labeled data and a substantial volume of unlabeled data. The labeled data facilitates the acquisition of fundamental task concepts, while the unlabeled data enables the network to generalize its understanding to novel examples.

Generative AI is a subset of deep learning, employing artificial neural networks to process both labeled and unlabeled data through supervised, unsupervised, and semi-supervised methodologies. It is worth noting that Large Language Models also fall under the umbrella of deep learning techniques.



THE MAJOR TYPES OF AI

At this point let's examine the AI discipline from a bird's eye point of view. This should empower you with a good overview understanding of the major types of AI.

Artificial Intelligence is not a monolith; it is a collection of methodologies and techniques that aim to simulate human intelligence. The major types of AI include:

1. **Symbolic AI (Rule-Based AI):** The earliest form of AI, this involves creating explicit rules for the AI system to follow. It's most effective in areas with a limited problem scope and where rules are well-defined.
2. **Evolutionary Computation:** This is inspired by biological evolution, such as inheritance, mutation, and selection. Solutions evolve and improve over time, making this approach excellent for optimization problems.
3. **Swarm Intelligence:** This involves the collective behavior of decentralized systems, often inspired by natural systems such as ant colonies or bird flocking. It's useful for tasks requiring parallel problem solving or optimization.

4. **Fuzzy Systems:** This AI deals with reasoning that is approximate rather than precisely deduced. It's useful for handling real-world ambiguity in decisions.
5. **Causal AI:** An emerging area that focuses on understanding the cause-and-effect relationships in data, which can make AI models more transparent, trustworthy, and capable of generalizing from one situation to another.

Each type of AI has its strengths and ideal use cases. They are often combined in practice to create hybrid systems that can tackle complex, real-world problems.

As investors, understanding these different types of AI allows us to see the potential and limitations of various AI technologies. Some sectors may benefit more from rule-based AI systems, while others may need the data-driven insights provided by machine learning. Newer fields like causal AI may offer unique opportunities for early investment.

Generative vs. Discriminative Models

The terms "discriminative models" and "generative models" often come up in discussions about machine learning. However, they're not typically categorized as separate types of AI. Rather, they're considered different approaches or strategies within machine learning.


Deep learning models, including machine learning models in general, can be categorized into two main types: generative and discriminative. Both discriminative and generative models can be used within the context of various types of AI mentioned in the previous section.

A *discriminative model* is designed to classify or predict labels for given data points. It is trained using a labeled dataset, where it learns the relationship between the features of the data points and their corresponding labels. Once trained, a discriminative model can make predictions on new data points by assigning them appropriate labels.

On the other hand, a *generative model* is capable of generating new data instances based on a learned probability distribution derived from existing data. These models have the ability to create new content that resembles the patterns observed in the training data.

In essence, generative models generate fresh data instances, while discriminative models distinguish between different types of data instances. Traditional machine learning models focus on learning the relationship between the data and the label or the target variable to be predicted. In contrast, Generative AI models concentrate on learning patterns within the data in order to generate new content.

Deep Learning Model Types



Discriminative

- Used to classify or predict
- Typically trained on a dataset of labeled data
- Learns the relationship between the features of the data points and the labels



Generative

- Generates new data that is similar to data it was trained on
- Understands distribution of data and how likely a given example is
- Predict next word in a sequence

Will The Real GenAI Please Stand Up

A useful way to differentiate Generative AI from other types of AI is by examining the output. If the output consists of numbers, classes (e.g., 'approve loan' or 'deny loan'), or probabilities, then it is not considered Generative AI. However, if the output is in the form of content, such as speech, text, images, code, or audio, then it falls into the category of Generative AI. For instance, if a model generates a sentence like "define marketing" based on its training on large datasets, it is considered Generative AI as it responds with content. The response would be based on all the massive amounts of data the model was already trained on.

Simply put, in the traditional supervised and unsupervised learning processes, models are built using training code and labeled data. These models can provide predictions, classify data, or cluster data based on the specific use case. By contrast, Generative AI processes involve training models using training code, labeled data, and unlabeled data of various types. The resulting Foundation Model can then generate new content, such as text, code, images, audio, video, and more.

THE EVOLUTION TOWARDS GENERATIVE AI

We've made significant progress in moving away from *traditional programming* towards neural networks and generative models. Let's consider an example where we want to train a computer to recognize a specific flower. In the traditional programming approach, we would have to manually code the rules for distinguishing that particular flower, such as the shape of its petals, the color, the number of petals, and other defining characteristics.

However, with the advent of neural networks, we can feed the model or network a large dataset of images containing various types of flowers, including the specific flower we want to identify. By training the network on these images and asking it questions like "Is this the desired flower?" or "Which flower is this?", it can learn to predict and differentiate the specific flower from others. You may have already encountered this technology when using image recognition features in applications like image search on your smartphone.

In the generative wave of AI, users now have the ability to generate their own content across different modalities, including text, images, audio, video, and more. For instance, models like OpenAI's ChatGPT, Google's PaLM (Pathways Language Model), Google's LAMDA (Language Model for Dialogue Applications), or Facebook's LLaMA (Large Language Model Meta AI) are trained on vast amounts of data from various sources on the internet. These models serve as foundation language models that can be accessed simply by posing a question or providing a prompt, whether typed or spoken. So, if you were to ask the model "Tell me about this flower," it would leverage its extensive knowledge to provide you with detailed information about that specific flower based on what it has learned from its training data.

GenAI is a type of Artificial Intelligence that creates new content based on what it has learned from existing content. The process of learning from existing content is called training and results in the creation of a statistical model when given a prompt.

AI uses the model to predict what an expected response might be and this generates new content. Essentially, it learns the underlying structure of the data and can then generate new samples that are similar to the data it was trained on. As previously mentioned, a generative language model can take what it has learned from the examples it's been shown and create something entirely new based on that information.

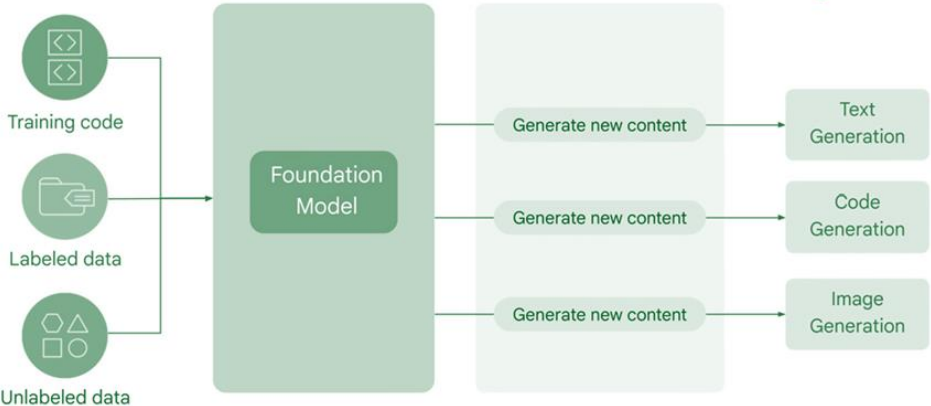
What is Generative AI?

- GenAI is a type of Artificial Intelligence that creates new content based on what it has learned from existing content.
- The process of learning from existing content is called training and results in the creation of a statistical model.
- When given a prompt, GenAI uses this statistical model to predict what an expected response might be—and this generates new content.

Generative Models vs. Foundation Models vs. LLMs

Generative AI model: This is a broad term that refers to any AI model capable of creating new content. This content can be text, but it can also be images, music, speech, or anything else that can be generated. For example, GPT-3 is a Generative AI model that generates text, while DALL-E is another Generative AI model from OpenAI that creates unique images from textual descriptions.

Foundation model: This term refers to large-scale machine learning models that are trained on a broad data foundation and can be fine-tuned for specific tasks. These models serve as a "foundation" because they capture a wide range of knowledge and capabilities, but usually require adaptation to perform well on specific tasks. A Foundation Model can be a Generative AI model if it's designed to generate content. An example of a Foundation Model is GPT-4, which is a Large Language Model trained on diverse internet text, but it can be fine-tuned to various language tasks.



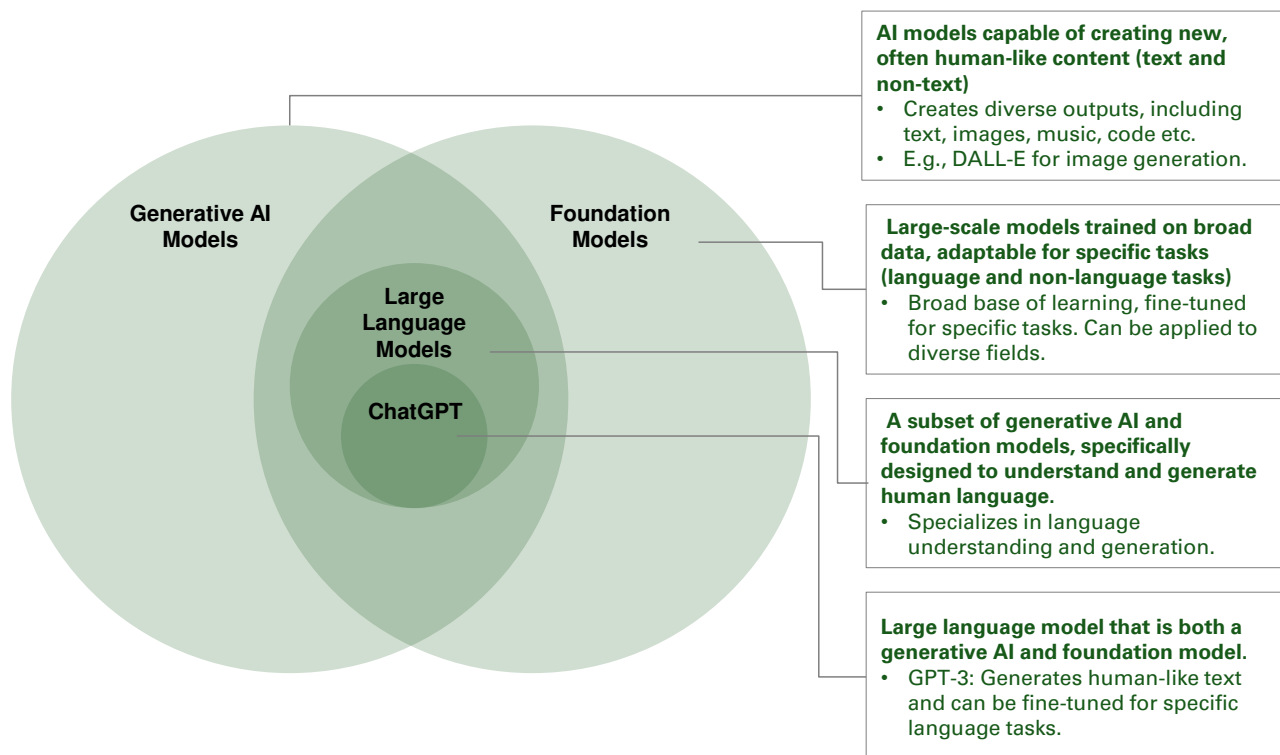
Foundation models are expensive to build, which is why you are seeing large companies take the lead on building them. The Oxford English Dictionary has about 171,000 words, although the vast majority of communication relies on about 50,000 words. Foundation Language Models are costly to build because they seek to ingest the entire corpus of publicly available text — to put that in perspective, today this corpus is equivalent to a human being reading for their entire lifetime across 100 lifetimes. To ingest all this data and train on it in less than 100 lifetimes, foundation models use massive farms of computers and GPUs to accelerate the process. Today, the state of the art is approximately 96 days to process all this data. But speed comes at a price. A recent report by OpenAI has found that

the cost of training large AI models is approximately \$100 million and could grow as much as five-fold by 2030.

The benefits of using Foundation Models are straightforward:

1. A single model can be used for different tasks. These foundation models are trained with petabytes of data and comprised of billions of parameters, and they are smart enough to solve different tasks.
2. Foundation models require minimal field training data when you tailor them to solve your specific problem. Foundation models obtain decent performance even with little domain training data. In other words, they can be used for “few-shot” or even “zero-shot” scenarios. In machine learning, few shot refers to training a model with minimal data, and zero shot implies that a model can recognize things that have not explicitly been taught in the training before.
3. The performance of foundation models is continuously growing when you add more data and parameters. Let’s take PaLM as an example. In April 2022, Google released PaLM, short for Pathways Language Model, a 540 billion-parameter model that achieves a state-of-the-art performance across multiple language tasks. PaLM is a dense decoder-only transformer model with 540 billion parameters.

Large language model: This is a type of AI model, specifically designed to understand and generate human language. These models are "large" because they have a high number of parameters (i.e., they learn a lot from the data they're trained on). GPT-4 is an example of a Large Language Model. These models are a subset of both Generative AI models (since they generate text) and Foundation Models (since they provide a broad foundation of language understanding that can be fine-tuned for specific tasks).



So, in essence, these three terms refer to overlapping categories. A Large Language Model like ChatGPT is also a Generative AI model (because it generates text) and a Foundation Model (because it provides a broad base of language understanding that can be adapted to specific tasks). However, not all Generative AI models are Large Language Models (since they might generate non-text content, like images or music), and not all Foundation Models are generative or language-focused (since they could be trained for other types of tasks).

ChatGPT

OpenAI's ChatGPT is one example of a generative language model that has gained widespread success because it is architected thoughtfully to generate meaningful responses to general-purpose text queries. It is important to know that ChatGPT is not the only game in town. Facebook, Google, and Microsoft have also announced sophisticated generative language models, and Amazon, Apple, and other large companies are believed to be building their own Foundation Models.

The creation of ChatGPT is rooted in extensive investigations into Generative AI systems, utilizing a machine learning technique known as Reinforcement Learning (RL). In an RL framework, an *intelligent agent* (considered an autonomous software in AI) engages with and gains insights from its surroundings. For ChatGPT, this 'surrounding' is the end-user.

GPT stands for Generative Pre-trained Transformer. We've already discussed what Generative means. Pretrained simply means that the model has already been fed a massive amount of publicly available information on the web and has pre-trained on it to learn about patterns and semantic relationships. The Transformer is the other recent breakthrough that makes this possible.

Transformers and Prompts

The power of Generative AI comes from the use of transformers. Transformers were invented in 2018 by researchers at Google and have led to a revolution in natural language processing. At a high level, a transformer model consists of an encoder and decoder. The encoder encodes the input sequence and passes it to the decoder, which learns how to decode the representation for a relevant task.

In transformers, *hallucinations* are words or phrases that are generated by the model that are often nonsensical or grammatically incorrect. Hallucinations can be caused by a number of factors, including because the model is not trained on enough data, the model is trained on noisy or dirty data, the model is not given enough context, or the model is not given enough constraints. Hallucinations can be a problem for transformers because they can make the output text difficult to understand. They can also make the model more likely to generate incorrect or misleading information.

A prompt is a short piece of text that is given to the Large Language Model as input. And it can be used to control the output of the model in a variety of ways. Prompt design is the process of creating a prompt that will generate the desired output from a Large Language Model.

Why Now?

The enhancements in the performance of Generative AI models can be attributed to several key factors:

Increased data. The unprecedented growth of digital data in recent years has provided a larger pool for training generative models. This abundance of data allows models to learn from a more extensive and varied data set, possibly leading to improved performance.

Improved algorithms. The past few years have seen numerous advancements in generative modeling algorithms, including Generative Adversarial Networks (GANs), Variational Autoencoders (VAEs), and transformers. These updated models are more adept at deciphering intricate data relationships and generating high-quality examples.

Superior hardware. The emergence of more robust hardware, such as GPUs and Google's Tensor Processing Units (TPUs), has empowered software engineers to develop larger and more intricate models. This capability to train expansive models has resulted in noticeable improvements in the quality of generative models.

Transfer learning. Generative models can now utilize transfer learning, where models pre-trained on extensive datasets can be refined on smaller, task-specific datasets or a set of cues designated for that task. By using the knowledge acquired during pre-training, the Generative AI can often deliver superior results with less data and training duration.

Despite the advancements in Large Language Models like GPT-3 (which forms the foundation for ChatGPT), BERT, and GPT-Neo over the years, they didn't receive widespread recognition from the public or media. It was only when OpenAI presented the abilities of GPT-3 (and now GPT-4) through a user-friendly interface through ChatGPT, that the public began to take note.

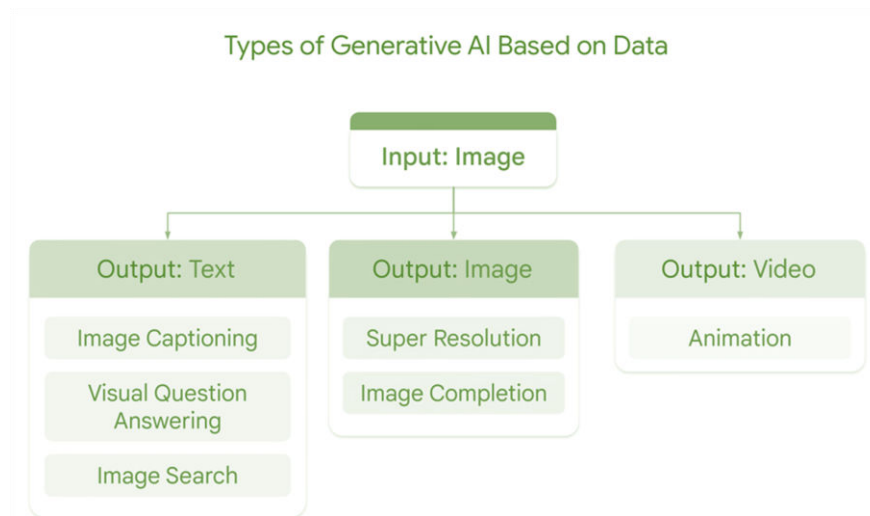
This enhanced accessibility to AI functionalities has introduced the broader public to the potential of Generative AI.

Types of Generative Models

As previously mentioned, GenAI depends a lot on the training data that you have fed into it. And it analyzes the patterns and structures of the input data and thus learns. But with access to a browser-based prompt, you, the user, can generate your own content. The associated model types include:

- **Text-to-text.** Text-to-text models take a natural language input and produce a text output. These models are trained to learn the mapping between a pair of texts — for example, translation from one language to another.
- **Text-to-image.** Text-to-image models are trained on a large set of images, each captioned with a short text description. Diffusion is one method used to achieve this.
- **Text-to-video and text-to-3D.** Text-to-video models aim to generate a video representation from text input. The input text can be anything from a single sentence to a full script. And the output is a video that corresponds to the input text. Similarly, text-to-3D models generate three-dimensional objects that correspond to a user's text description. For example, this can be used in games or other 3D digital environments, such as computer-aided engineering design.
- **Text-to-task.** Text-to-task models are trained to perform a defined task or action based on text input. This task can be a wide range of actions such as answering a question, performing a search, making a prediction, or taking some sort of action. For example, a text-to-task model could be trained to navigate a web UI or make changes to a doc through a word processor's graphical user interface.

The same logic extends to image inputs, as shown below. An image input can be translated into image-to-text, image-to-image, image-to-video, and so on.



Generative AI Applications

At present, the discourse surrounding Generative AI primarily emphasizes its ability to create textual outputs, such as academic essays or promotional material, via Large Language Models, and digital artwork through Latent Diffusion Models (LDMs) such as Stable Diffusion and DALL-E.

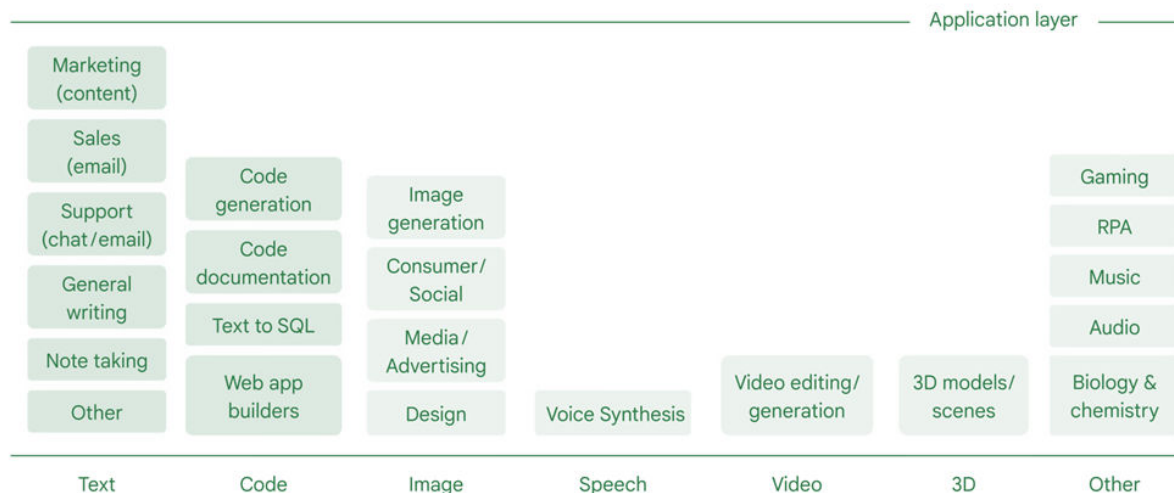
Nevertheless, we believe the capabilities of Generative AI extend beyond just content creation.

Consider, for instance, the potential role of Generative AI in promoting the uptake of AI and conversation-based interfaces. Organizations could utilize these base models in conjunction with their unique data sets. From our perspective, the wide-ranging accessibility of Generative AI models, offered by a plethora of providers and open-source initiatives, combined with the opportunity to refine these models using proprietary data, indicates that software enterprises with substantial data resources are poised to reap the biggest benefits. We predict an uptick in the number of LLMs that are trained on proprietary data and information.

Particularly for startup companies, constantly training LLMs might not always be practical due to high costs and environmental impact. However, adopting a sensible approach to the use of pre-trained LLMs could aid businesses in initiating their Generative AI endeavors. There are applications such as information condensation or elaboration based on a few key points, which can be implemented directly with these LLMs. Another instance might involve the use of other Natural Language Processing (NLP) methods, like search, that can enhance your prompt with your data to deliver customized dialogues for your clients.

In addition to conversational applications, Generative AI harbors the potential to revolutionize engineering designs by automating the process. This could lead to a quicker generation of design variations, optimization of materials, and predictions regarding the performance of a design under varying conditions, taking into account aspects like demographics, preferences, and physical characteristics. Moreover, Generative AI could increase the efficiency of testing iterations by digitally replicating the performance of a product under different conditions, yielding insights that physical testing might not be able to provide. Ultimately, design processes could be made more accessible to those without expert knowledge, as AI automates complex procedures and people can communicate with the AI through intuitive, natural language interfaces.

The generative AI Application Landscape



MAKING USE OF GENERATIVE AI

Business Opportunities

Opportunity Insight #1: Unique Data Moats Create Defensible Advantages

At Omega, we've identified a key strategic advantage for current and prospective portfolio companies: unique data reserves. We believe that companies with these robust data sets can build a long-lasting competitive edge, especially when they use the data wisely to sharpen their business models.

Take, for example, a business in the retail sector that's gathering and analyzing customer data. This exclusive data allows them to predict future consumer trends better than their competitors who rely on generic models. The hard-to-replicate nature of this data offers a sustainable edge that sets the startup apart from its peers.

In fact, existing companies with unique data reserves may have a leg up over startups in the same field. However, younger companies can also compete by using synthetic data generated by AI to create their initial data sets, a cost-effective way to establish a data foundation.

While maintaining large data sets can be costly, we encourage startups to work with just enough data to ensure a positive customer experience. Prioritizing quality over quantity can save resources without compromising customer satisfaction.

Here are a few ways to use unique data reserves:

- **Hyper-Personalized Experiences.** Using segmented user data, companies can generate personalized outputs for different customer groups. Think of a tutor chatbot assisting university students according to their individual learning styles, continuously improving its techniques based on user feedback and behavior.
- **Reinforcement Learning.** RL is a learning approach where an AI learns from its environment and adjusts its actions to maximize a reward. RL, when combined with human feedback is

called RLHF. Here are five ways companies can use RL and RLHF to compound defensible advantages from access to proprietary data:

- (i) **Customization and Adaptation:** RL algorithms can be trained to adapt to specific environments and problems. If a company possesses unique data, they can train RL models specifically tailored to that data, making the models more effective and difficult for competitors to replicate.
- (ii) **Continuous Learning:** One of the strengths of RL and RLHF is their ability to learn and improve over time based on the feedback received. This learning feature allows them to keep refining their models and creating better solutions or products over time, thus making it difficult for competitors to keep pace.
- (iii) **Developing Proprietary Technology:** If a company has access to unique data, they can use RL and RLHF to develop proprietary technology. For example, if a company has data on user behavior, they can create a recommendation system that accurately predicts and caters to user preferences. This system, driven by RL, would be unique to the company and serve as a significant barrier to entry for competitors.
- (iv) **Increasing Efficiency and Optimization:** RL algorithms can be used to optimize various operations within a company. For instance, they can be applied to logistics, scheduling, resource allocation, etc., to create efficiencies. If these efficiencies are based on unique data, it becomes difficult for competitors to achieve the same level of optimization.
- (v) **Creating Value from Data:** By applying RL and RLHF on unique data, a company can extract insights, generate predictions, and provide valuable services that might not be possible for competitors without access to similar data.

Opportunity Insight #2: Natural Language User Interfaces

While Conversational AI might pose hurdles to established firms more accustomed to classic UX than chatbots, we see chat interfaces as a game-changer that can be adopted with great success.

While traditional UX design will still hold its place, chat interfaces could potentially provide that edge over competitors. They enable UX designers to construct user workflows faster — the steps users follow to achieve a goal within an app. This speed comes from the simplicity of translating tasks into natural language prompts linked to key services providing actions, once a conversational AI interface is in place. This could be far simpler than designing traditional workflows and corresponding screens.

Most businesses would benefit from integrating conversational user interfaces, which can yield immediate dividends in the form of enhanced product or service usability and usefulness. Natural Language interfaces are becoming the new UX paradigm across technology, enabling users to input their objectives or queries in simple natural language phrases. Being able to use natural language to use software and other technologies naturally lowers technical barriers and facilitates broader adoption.

Here are four ways that companies can take advantage of NL interfaces:

- **Customer Support:** Utilizing Natural Language Processing (NLP), companies can deploy AI-powered chatbots to provide 24/7 support, answering customer queries quickly and efficiently. These chatbots can be programmed to understand and respond to complex requests, reducing response time and freeing human agents to deal with higher order issues. This can increase customer satisfaction, ultimately leading to higher customer retention as well as lower support costs.
- **Employee Self-Help Portals:** Implementing NL interfaces in internal tools can greatly enhance the employee experience. For example, HR systems can be designed with chat interfaces that answer queries about company policies, benefits, and more. This not only saves time for the HR team but also empowers employees to find answers to their questions without having to wait for assistance. This is particularly beneficial for remote workforces where live consultation can be more difficult.
- **Client Success:** Natural language user interfaces can help companies keep their clients engaged by offering personalized, efficient, and interactive solutions. For example, financial advisors can use AI-powered interfaces to help clients access their account details, discover investment options, make decisions, and understand complex financial terms. In software companies, the client success function is often about ensuring that clients are successfully adopting the software, demonstrating sustained engagement, and ideally increasing their usage over time. With natural language interfaces, the learning curve for clients can be greatly reduced, helping them get value out of a new product or service much faster while mitigating confusion and frustration.
- **Business Operations:** Natural language interfaces can be extremely valuable in data analysis and visualization tasks, such as business or sales operations. Rather than requiring complex coding or scripting to extract insights, users can simply ask questions in natural language to pull the required data or generate a specific analysis. This democratizes access to data analysis, enabling non-technical stakeholders to get the insights they need directly, without the bottleneck of data science teams. For example, a marketing manager could ask a system "Show me the trend of website visits for the past six months" and get the necessary graph without needing to understand SQL queries or data wrangling. This simplifies the interaction with data, making it more accessible and usable across the organization. And it makes managers more effective because they can now better measure what they are managing.

Opportunity Insight #3: Smart Model Management

The emerging field of Large Language Model Operations (LLM Ops) providers present a significant opportunity as businesses incorporate Generative AI into their existing systems. Processes and techniques that can streamline costs and simplify stages of the AI development cycle are invaluable allies for businesses seeking to scale.

For instance, a business might not require the largest model for a specific task. A model with a few million parameters could suffice instead of a few hundred million parameters, providing the needed accuracy for a particular use case. LLM Ops tools can assist businesses in identifying the ideal model size and configuration for their requirements. As we've experienced, a critical task when implementing Large Language Models is to make them smaller, despite the industry's growing tendency towards larger models.

LLM Ops can pinpoint strategies to optimize model size and training processes, a crucial aspect considering the high costs of LLM training. As a result, companies often compromise between cost and accuracy, leading to under-optimized models. For instance, while OpenAI's GPT-3 model was 10x the size of its predecessor, GPT-2, DeepMind's Chinchilla model proved more effective with only 70 billion parameters. This indicates that many current LLMs are likely oversized and undertrained.

Streamline Your Work and Personal Life

Individuals from various backgrounds, whether executives, investors, students, coders, artists, lawyers, or accountants, are actively exploring the potential of AI tools. To assist you in leveraging these tools to improve various aspects of your life, we will present two general approaches that can be applied to numerous situations.

The Power of Precision: Good Prompts

ChatGPT, Bing, and Bard are just a few examples of popular AI chatbots. Although these chatbots may appear straightforward, with a mere input-output mechanism, it is crucial to understand that asking questions in the wrong manner may yield generic, unhelpful, or even incorrect responses.

There is an art to crafting precise prompts that generate the most valuable answers. To maximize the potential of chatbots try the following techniques:

1. **"You are a...":** By commencing your prompt with these words, you can instruct the bot to emulate an expert in a particular field. For instance, typing "You are a financial advisor" or "You are a travel consultant" guides the chatbot to model itself around professionals in those respective domains. You can even say "You are an NLP engineer" to ask the system to generate the optimal prompt for you based on your initial description.

Such prompts provide additional context for the AI, enabling it to generate more informed responses. While the AI may not possess a true understanding of what it means to be a financial advisor or a travel consultant, the prompt helps the AI tap into specific statistical patterns from its training data.

A weak prompt devoid of guidance will result in less helpful outcomes. For instance, if you simply ask, "What should I do this week?" the chatbot might generate a generic list of activities, which may not be particularly interesting or relevant to you.

2. **"Ask me for any additional information you need to generate an excellent answer that will translate into tangible results":** To obtain more personalized results, such as tailored advice based on your specific goals or constraints, you can instruct the bot to request additional information.

For example, your prompt could be: "You are my golf coach. Create a plan for me to lower my handicap to eight or less. Ask me for any additional information you need to generate an excellent answer that will translate into tangible results." Subsequently, the bot might inquire about your current skill level, technical proficiency, strength and conditioning, practice routine, what courses you play, and mental game training routine to customize a coaching roadmap for you. The key here is to treat the bot as if it were a research assistant you just hired.

Break Down Your Query into Sequential Prompts

Try breaking down your query into a series of prompts. This approach will allow the chatbot to provide more detailed information for each subquery. For example, let's say you're looking for M&A comps for a particular investment target. Start with an initial prompt that asks for M&A data on companies across a sector. After receiving the results for the sector, you can continue with similar prompts for other sectors or sub-sectors and different time periods. Once you have gathered the data for each subquery, you can combine them into a comprehensive table.

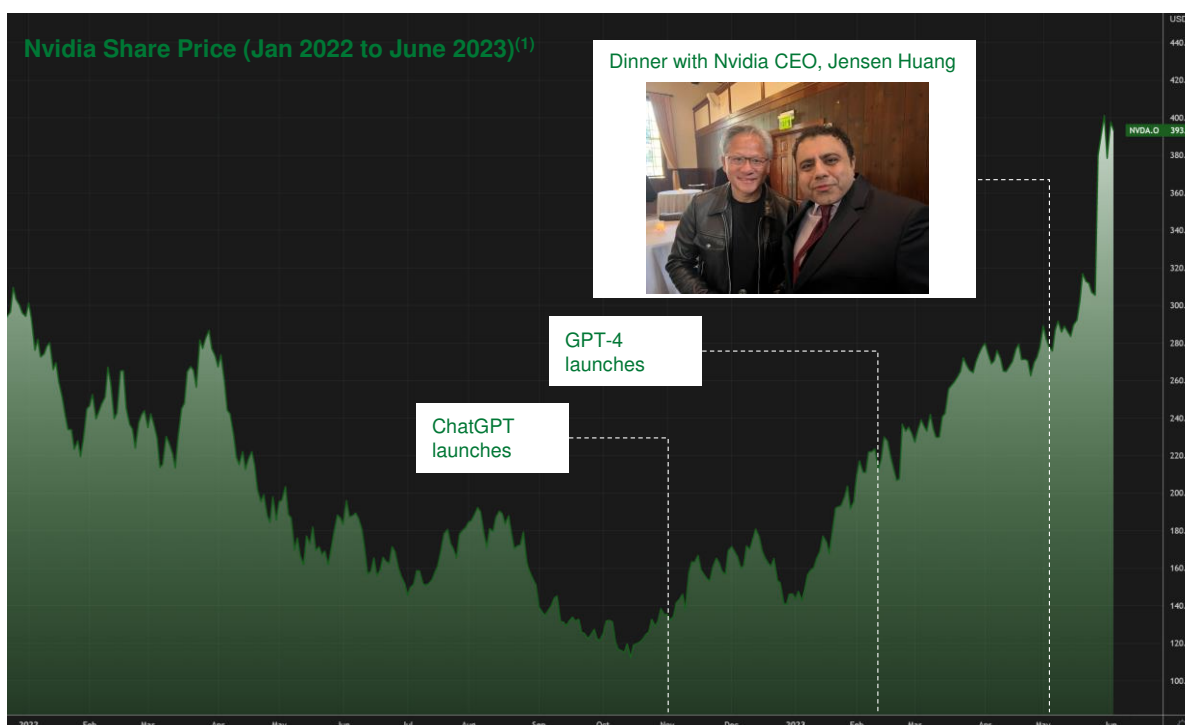
The key is to avoid treating your chatbot as a simple web search engine, starting afresh with each query. Remember, the chatbot retains a certain amount of contextual memory to understand how a new prompt fits within the context of prior prompts in the same conversation. This is something you should use to your advantage.

* * *

FINAL THOUGHTS

Investing in Artificial Intelligence requires a strategic approach that goes beyond following trends. While the AI hype has been soaring lately, it's essential to exercise real intelligence when allocating investments.

The success of ChatGPT and the excitement surrounding Generative AI have led to a surge in interest and investments. Chip maker Nvidia briefly reached a market value of over \$1 trillion, driven by the demand for its chips in the AI space. While Nvidia stands out as an obvious candidate due to its chips' importance in the AI realm — we know that Nvidia's GPUs are powering ~90% of AI model training and inference (execution) workloads today — the stock has already experienced substantial gains.



¹ Callouts should not be construed to imply causality with share price.

This is where Omega's approach to investing in early-growth stage private companies proves advantageous. We have been prescient in recognizing that AI and ML are shaping the Fourth Industrial Revolution, driving a tsunami of innovation, and creating once-in-a-generation opportunities. We were amongst the first to act on this insight, with Omega I launched in 2020. And this is backed by our team's 20+ year immersion in AI. Our expertise extends to identifying businesses with the potential for substantial shareholder value growth through the judicious integration of AI, which our team is uniquely equipped to help catalyze.

We seek to consistently and prudently back category leaders capitalizing on the transformative power of AI, the vast majority of which are private companies. In this environment, our deep domain expertise is more important than ever in order to correctly and efficiently separate the wheat from the chaff. Over the last few years, we have sought to establish Omega as the preferred investment partner for leading innovators and entrepreneurs who are building consequential companies at the forefront of the Artificial Intelligence revolution. Our results support that these efforts are bearing fruit.

The fact that AI has achieved an inflection point in technological viability and commercial feasibility is beyond clear. We view AI as the biggest driver of technological innovation and investor value creation for the foreseeable future. Our purpose-built investment platform equips us well to capitalize upon this generational value-creation tsunami. We are more enthusiastic than ever to leverage our momentum in the journey ahead.

* * *

IMPORTANT DISCLOSURES

THESE MATERIALS AND THE VIEWS EXPRESSED HEREIN ARE INTENDED FOR INFORMATIONAL PURPOSES ONLY, AND SHOULD NOT BE USED AS OR AS A SUBSTITUTE FOR INVESTMENT, FINANCIAL, TAX OR OTHER FORMS OF ADVICE OR FOR ANY OTHER PURPOSE. THESE MATERIALS ARE NOT INTENDED TO ACT AS AN OFFER OR SOLICITATION FOR THE PURCHASE OR SALE OF ANY FINANCIAL INSTRUMENT. INFORMATION REFERENCED OR USED HEREIN HAS BEEN OBTAINED FROM SOURCES BELIEVED TO BE RELIABLE BUT OMEGA VENTURE PARTNERS (OMEGA) DOES NOT WARRANT THE COMPLETENESS OR ACCURACY OF ANY THIRD-PARTY INFORMATION. PAST PERFORMANCE IS NOT INDICATIVE OF FUTURE RESULTS. GRAPHS AND CHARTS CANNOT BE USED IN AND OF THEMSELVES TO DETERMINE THE ADVISABILITY OF INVESTING IN SECURITIES OR THE TIMING OF THE PURCHASE OR SALE OF ANY SECURITY.

THESE MATERIALS CONTAIN OPINIONS AND/OR ESTIMATES OF OMEGA AND ITS PERSONNEL THAT CONSTITUTE ITS AND THEIR PERSPECTIVES AS OF THE DATE OF THESE MATERIALS AND ARE SUBJECT TO CHANGE WITHOUT NOTICE. THE OPINIONS CONTAINED IN THESE MATERIALS DO NOT TAKE INTO ACCOUNT INDIVIDUAL CIRCUMSTANCES, OBJECTIVES, OR NEEDS AND ARE NOT INTENDED AS RECOMMENDATIONS OF PARTICULAR SECURITIES, FINANCIAL INSTRUMENTS, OR STRATEGIES. THE RECIPIENT OF THESE MATERIALS MUST INDEPENDENTLY UNDERTAKE ITS OWN DILIGENCE AND THEREAFTER MAKE ITS OWN DECISIONS REGARDING THE ADVISABILITY OF ANY SECURITIES OR FINANCIAL INSTRUMENTS REFERENCED HEREIN.

OMEGA WILL NOT BE RESPONSIBLE OR HAVE ANY LIABILITY FOR ANY INJURIES OR DAMAGES OF THE RECIPIENT OF THESE MATERIALS OR ANY THIRD PARTY, INCLUDING INJURIES OR DAMAGES THAT MAY RESULT FROM (I) THE RELIANCE OF ANY PERSON UPON ANY INFORMATION OR OPINION PROVIDED HEREIN OR (II) ANY ERRORS, INACCURACIES, OMISSIONS IN, OR ANY OTHER FAILURE OF, THE DATA HEREIN, FROM WHATEVER CAUSE. OMEGA SHALL NOT BE LIABLE FOR DIRECT, INDIRECT OR INCIDENTAL SPECIAL OR CONSEQUENTIAL DAMAGES RESULTING FROM THE INFORMATION CONTAINED HEREIN AND/OR FROM ANY DECISION TAKEN ON THE BASIS OF SUCH INFORMATION, REGARDLESS OF WHETHER SUCH DAMAGES WERE FORESEEABLE. SOME OF THE COMPANIES SELECTED FOR INCLUSION IN THIS REPORT ARE HELD BY INVESTMENT VEHICLES OR CLIENTS MANAGED BY OMEGA.