TPG INSIGHTS

The AI Revolution

> **For the first time in history, we're able to extract more shared human knowledge than any one person can physically retain … The result is AI systems that are very powerful cognitive co-pilots for us and operate as true "bicycles for the mind".**
>
> *- Lila Treitkov*

> **This platform shift is going to be something that has a profound impact on every one of our companies and not just those in software and enterprise technology.**
>
> *- Tim Millikin*

> **The rise of generative AI looks like the next big technology platform shift, though it's probably going to be a decade or more before we feel its full effects.**
>
> *- Mike Zappert*

The era of Artificial Intelligence has arrived. Despite advancing in fits and starts in recent decades, AI is experiencing nothing short of an explosion, the consequences of which will be with us for decades to come. Amid this pivotal moment of technological change, we bring together **Lila Treitkov**, Deputy CTO of Microsoft, and **Tim Millikin** and **Mike Zappert**, two of TPG's leading Software & Enterprise Tech investors, to provide their perspective on the promise and perils of AI, what's next for this astonishing technology, and where to expect the largest pockets of value creation. Treitkov, who's deeply involved in Microsoft's partnership with OpenAI, has little doubt: AI, she argues, will be as big if not bigger than the Industrial Revolution in terms of its economic impact. For their part, Millikin and Zappert generally agree, but they believe it might be some time before the full effects of AI trickle down to the more established segments of the tech stack. That said, all three see the biggest investment opportunities at the application layer, given companies' need for specialized AI tools and the massive size of the markets that AI could disrupt. As the pace of AI advancements continues to pick up, we also speak to **Art Heidrich**, an investor in TPG's Software & Enterprise Tech group, who answers several important FAQs on AI and the path ahead for its commercialization.

IN THIS ISSUE

GABE LIPTON GALBRAITH ggalbraith@advisorstpg.com

TPG INSIGHTS

Interview with Lila Treitkov

Lila Treitkov is the **Deputy Chief Technology Officer at Microsoft** where she leads the firm's efforts to develop the next generation of breakthrough technologies, including Artificial Intelligence. Treitkov has been deeply involved in the firm's partnership with OpenAI since Microsoft's initial investment in 2019. She's also a **board member** of TPG Growth portfolio company Onfido. Prior to joining Microsoft, Treitkov served as the **CEO of Wikipedia** from 2014-16. Below, she offers her perspective on the future of AI, concerns about AI safety, and where she sees the biggest opportunities for value creation looking ahead.

The views stated herein are those of the interviewee and do not necessarily reflect those of TPG.

Q: You've spent your life working in computer science and AI. How big of a breakthrough are the GBT models and ChatGBT?

Lila Tretikov: We tend to forget that AI is based on centuries of progressively sophisticated theoretical math. The latest breakthroughs are the product of accelerated algorithmic achievements applied in the context of modern computing and the Internet. There are three driving forces. First, we've built enormous supercomputers with hundreds of thousands of GPUs and connected them via high-performance InfiniBand networking, pushing hardware beyond its limits. Second, the availability of vast digital data has enabled the training of far larger models, and we've found that a lot more can be achieved with scale. Third, we've doubled down on the autoregressive transformer (the T in GPT) architecture, which has allowed us to discover the scaling properties of these models.

For the first time in history, we're able to extract more shared human knowledge than any one person can physically retain, deliver it in a system that can process knowledge efficiently, and surface it via a natural language interface. What we didn't expect is the speed of progress. We also did not anticipate some of the emergent behaviors of these systems that can feel human-like in intelligence but are actually quite distinct. The result is AI systems that are very powerful cognitive co-pilots for us and operate as true "bicycles for the mind".

Q: Are there any constraints on further advancements in AI?

Lila Treitkov: There are some structural bottlenecks around available compute, power, network capacity, memory, and the algorithms themselves. But many of those can be addressed with optimization techniques: while these LLMs are often quite expensive initially, we're able to bring down costs quickly over time. The far bigger–and self-imposed–constraint is ensuring the responsible development of AI, as should certainly be the case given the power of the technology. We go to great lengths to make sure that our models are aligned with our ethical principles, grounded in factual truths, de-biased, and used responsibly once they're released. This is incredibly difficult in practice given the extreme number of parameters that have to be considered. We often hold technology back until we're certain that all our ethical AI principles have been met, which certainly slows us down. But I'm glad we've built these safeguards.

Q: What's next in terms of the trajectory of AI?

Lila Tretikov: The LLM is the new central processing unit (CPU). They will be surrounded by other components such as memory, plug-ins, and application runtimes to create a form of an AI-based operating system (OS). On top of this infrastructure, myriad applications will be built, which we're already starting to see with the proliferation of what Microsoft calls co-pilots. I believe this application layer will drive the next phase of staggering growth. There will continue to be bigger foundational models or "central brains". But as with the Internet and mobile, the apps are going to create the most value and make almost everything in our lives easier. AI will enable cognitive workers to focus on high-value work, quite possibly increasing their productivity as much as the productivity boost to physical labor during the last Industrial Revolution (on the order of 50x).

An even more exciting promise, however, is that AI will enable us to decode aspects of the world that we do not understand today. Almost anything can be reduced to a form of language and represented by a string of numbers, making it accessible to AI. Human biology, for example, is an incredibly complex system that's difficult to express with a formula. Brute force math also will not work: the sheer volume of humanity's epigenomic information could not fit on all the computers in the world.

> **AI is a reflection of what we know and how we understand the world … Our task is to align it with what humanity thinks the world *should be*.**

But AI can be applied to biology, chemistry, and physics to isolate only the most important parameters based on the data it sees. This can help us generate cures for diseases or materials for sustainability. In many areas, AI will enable us to solve problems that are too complex for the human brain.

Q: Do you worry about the labor dislocations this might cause?

Lila Tretikov: We're entering a period of exponential change, and so we're going to have to help people retrain and reskill faster than ever before. This change is not new, but it is accelerating. At the same time, there's a massive labor deficit in many cognitive professions today. We need more resources, especially for reading and summarizing documents, processing bills, responding to inquiries, and transcription. This is exactly where AI shines. Professionals will generally benefit from the productivity enhancements that come from AI, though people earlier in their career will need to focus their training on leveraging AI and sharpening their analytical skills. There's going to be an entire generation of college graduates that will have to be taught how to manage AI ethically and responsibly. That's something only humans can do. It requires a degree of critical thinking that's only taught in universities today, but it will increasingly need to be a part of all students' education.

Q: Where will the greatest value be captured as this technology is commercialized?

Lila Tretikov: It depends on the layer. At the infrastructure level, startups have agility and an ability to take risks, which gives them an advantage over the large incumbent players. Plenty of startups, such as OpenAI, have had great success in building AI models because they were willing to push innovation in less explored areas. The startups will partner with the cloud providers for their compute. So, these providers, GPU manufacturers and networking infrastructure certainly stand to gain. But much of the value will appear at the application layer. Some applications will be "winner takes most", but the more common case will be another SaaS-like wave of competition, though this time with AI-driven applications. In this context, domain expertise will matter, and the winning solutions will match the specific regulations, processes, and requirements of each sector and sub-sector vertical. These are massive markets, and so these apps will create significant user value.

Q: Many policymakers and technologists have raised safety concerns about AI advancing beyond the point where we can fully understand and control it. Are these concerns warranted?

Lila Tretikov: Every powerful technology is dual use. It's our responsibility as individuals, companies, and society to ensure that the technologies we produce are used ethically. The good news is that Microsoft and many of the other tech companies take this very seriously. Microsoft laid out our principles for responsible AI all the way back in 2018, and there's an important group within the firm dedicated to these issues.

In fact, the first people who get access to our models are our so-called "red teams". They spend a couple of months studying any possible side effects before anyone else can build product. Of course, we can't guarantee that everyone will be responsible, which is why we've long supported well-designed tech rules and regulations. The key is making sure these regulations work in the best interest of society and don't end up backfiring.

Q: How do you create guardrails for AI models in practice?

Lila Tretikov: When I was the CEO of Wikipedia, we noticed a simple pattern: the systems we build hold a mirror to the world we create. They are a reflection of what we know and how we understand the world. People often justly complain about biases on Wikipedia. But Wikipedia is built by individuals and reflects the state of the world today–not the ideal state. AI is no different. It reflects the information that we have created. Our task is to align it with what humanity thinks the world *should be*. In the words of Mahatma Ghandhi, we're responsible to "be the change we want to see in the world". AI is giving us that chance.

In practice, we put the models through an extensive "alignment" and "grounding" process to make sure they maintain humanity's values and deliver factual information. It's an incredibly difficult problem because we also want to train them on as much information as possible, which often contains fantasy stories or accounts of war. Without this strong grounding and alignment, the models can hallucinate where they take something that they've read, say a piece of science fiction, and insert it into a conversation about a current event. This is by no means consciousness, but it can sound convincing because the model has learned from us.

The more difficult challenge, then, is that we, as a society, don't know what our values will be in the future. We're aligning our models to the world as it exists today, but the more important question is where we want to be tomorrow. Anytime we're inventing a new technology, we need to be thinking hundreds of years ahead. Take the combustion engine, for example. It was created more than 150 years ago, but it's only in the last 10-20 years that we've come to agree that its biproducts have environmental consequences. The promise and challenges of AI are of the same magnitude.

Q: What motivates you to be involved in this important project?

Lila Tretikov: I grew up studying science and art as a child, which for me were two sides of the same coin. Art reflects the mystery of the world. Science provides a way to understand it. But it's technology that offers a bridge. Technology is the connective tissue that enables us to realize our dreams and to make sense of the magic that is the world that surrounds us. As humans, we are eternally in search of meaning. Technology is one navigational instrument on our journey that is filled with a sense of wonder and the joy of discovery.

TPG INSIGHTS

The Explosion of Large Language Models

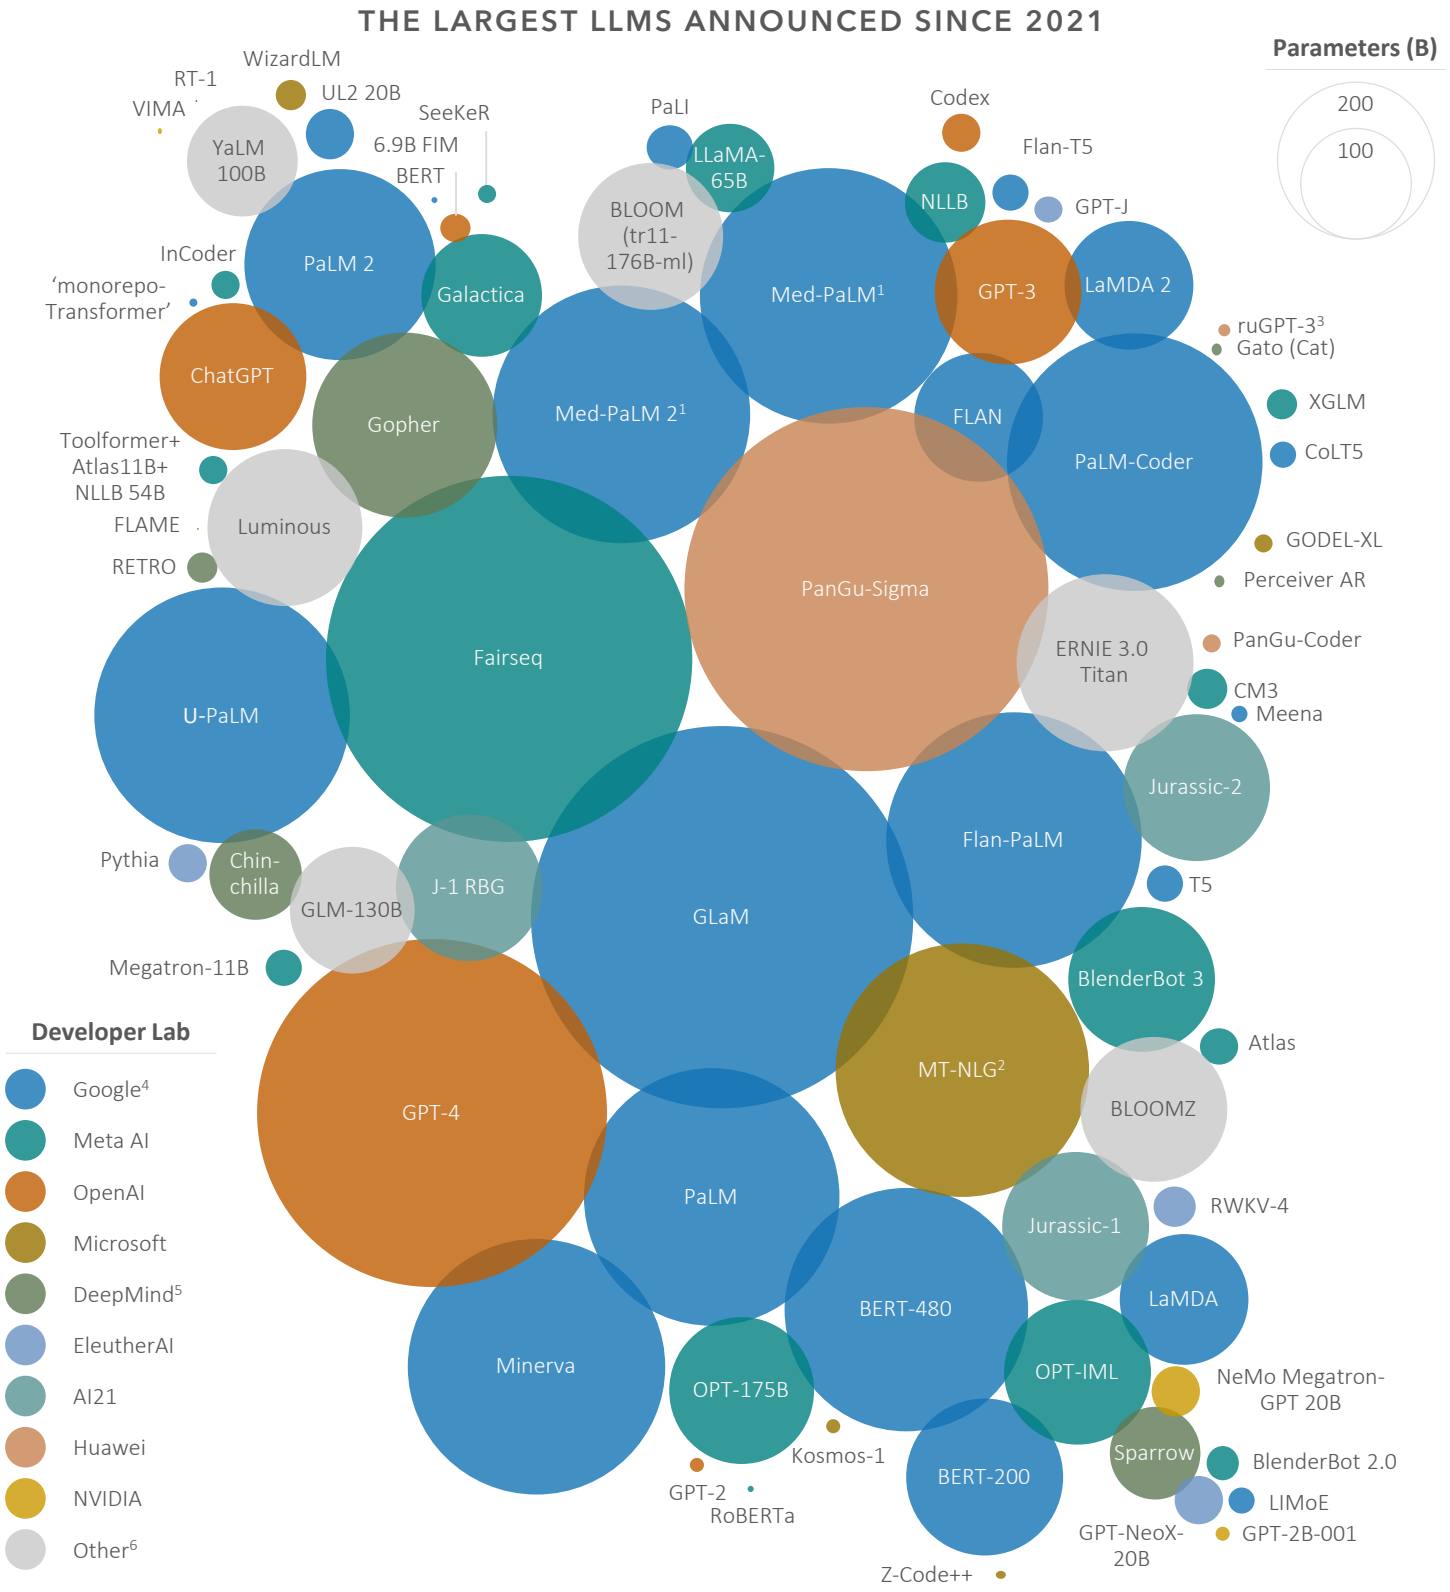

1. From Google & DeepMind labs. 2. From Microsoft & NVIDIA labs. 3. From Huawei & Sberbank labs. 4. Google includes Google Inc, Google Research & Google Brain. 5. Models categorized under DeepMind include those announced prior to its merger with Google Brain 6. Other includes Baidu Aleph Alpha, BigScience, Tsinghua & Zhipu, and Yandex labs.
Note: Data as of May 12, 2023.
Source: Dr. Alan D. Thompson, Arxiv, various online databases.

TPG INSIGHTS

Interview with Tim Millikin and Mike Zappert

Tim Millikin and Mike Zappert are Partners at TPG based in San Francisco. Tim **co-leads TPG's investment activities in software and enterprise tech**. Mike **leads software and enterprise tech investing across the TPG Growth** and **The Rise Fund platforms**, and he is actively involved with **TPG Tech Adjacencies (TTAD)**. Prior to joining TPG in 2007, Tim worked in the Technology M&A group at Morgan Stanley. Mike was previously a Partner at Adams Street Partners and also held various roles at 3i and Credit Suisse. Below, they provide their perspective on the latest developments in AI and their implications for private market tech investing.

Q: Given your extensive experience as tech investors, what are the biggest investment opportunities and risks surrounding the rapid advancement of AI?

Mike Zappert: The rise of generative AI looks like the next big technology platform shift, though it's probably going to be a decade or more before we feel its full effects. Like other shifts I've experienced during my career—whether from web to mobile or on-prem computing to the cloud—AI has the potential to radically reshape how businesses operate across the entire economy and create significant value in the process. The critical driver of the latest advancements in AI that have everyone talking is the move from models that require a lot of human training towards those that are effectively self-teaching, which has reduced the time it takes to build faster and more powerful AI systems. Progress in natural language processing, or the ability of AI systems to understand human language rather than code, has also enabled users to interact with the technology in a way that simply wasn't possible in the past.

In terms of risks, AI has now advanced to the point where it can perform complicated cognitive tasks at a level on par or even surpassing most humans, which naturally raises concerns that it could eventually disrupt significant swaths of the economy. The most oft-cited example of this is that ChatGPT went from scoring in the bottom 10% on the bar exam to the top 10% in just a matter of months between the releases of GPT 3.5 and 4.0. Additionally, the pace at which the latest AI models are progressing is well ahead of our ability to fully understand them, and so it's going to take many years before we're able to put in place the necessary rules and policies to govern how AI is used and regulated.

Q: There have been many cycles of tech exuberance that have ultimately disappointed. Do you think this time will be different?

Mike Zappert: Yes, we do. AI itself is a technology that's had its share of false dawns. The 70s and 80s, for example, saw an AI boom, but it was sandwiched between two AI winters. The biggest issue then was that computing power simply wasn't sufficient to support the scalable models that we have today. That said, with the caveat that we're by no means early-stage investors, it does seem likely that the capabilities that have been developed over the past couple of years will lead to a more durable period of expansive growth. The current moment feels a lot like when I was first starting out in the Cloud space in 2012. I remember walking the booths at AWS Reinvent, which at the time was only a few thousand people, and you could feel that a meaningful change was at hand. It's still early days for AI, but there are a number of important tailwinds that suggest to me that we've turned a corner.

> "The current moment feels a lot like the Cloud space in 2012 … It's still early days for AI, but there are a number of important tailwinds that suggest to me that we've turned a corner.

Tim Millikin: AI is as an enabling technology. There are a lot of underlying capabilities that were needed to get us to the current moment, including increased compute capacity, the advancement of programming and algorithmic tools, and the development of new applications. But now that the technology is on solid footing, the industry has turned its attention to how AI will actually be used by both consumers and businesses in the years ahead.

It's also important to distinguish between traditional AI, which has been in use for quite some time for things like document search, anomaly detection, and risk management, from generative AI. The latter has opened up an entirely new world in terms of content creation, our ability to operate on unstructured datasets for functions like customer service and sales and marketing, and the prospect for enabling businesses to make better decisions based on more structured data such as supply chains or sales pipelines. Many of these use cases are still in the idea phase, but when they do develop into commercially available tools, we believe these tools have the potential to create billions of dollars of market cap.

"

The rise of generative AI looks like the next big technology platform shift, though it's probably going to be a decade or more before we feel its full effects.

Q: How quickly might generative AI be commercialized?

Tim Millikin: Platform shifts like this one tend to occur much more slowly than people expect, but we feel strongly about the momentum behind generative AI. There are still a lot of large-scale impediments to commercial and enterprise adoption. First, it's going to be critical to figure out how to build applications with the underlying technology so that businesses can utilize it for their specific enterprise needs. Second, security and privacy are going to be important considerations for any enterprise that's looking to adopt this technology. In fact, you're already seeing some companies try to create their own instance of ChatGPT in order to protect their data, which is the same dynamic we saw in the Cloud and SaaS worlds 15-20 years ago and is still the case today.

> "We have a lot of confidence that AI is the wave of the future, but the question is just how near we are to that future in terms of a sweeping reconfiguration of software and enterprise tech. That probably will take quite some time.

Third, enterprises are concerned about IP and IP licensing when it comes to the ability of generative AI to create new content, as we saw with the deep fake of Drake and the Weeknd, which produced a number of legal questions. We have a lot of confidence that AI is the wave of the future, but the question is just how near we are to that future in terms of a sweeping reconfiguration of software and enterprise tech. That probably will take quite some time.

Mike Zappert: As was the case with the Cloud, it typically takes a decade or more for these technologies to go from an interesting idea to market dominance. The companies that are usually the winners in the commercialization process are typically those that can translate very advanced technology and make it available at a platform level or provide a best-of-breed solution that wouldn't otherwise be available.

We see a multi-year opportunity for companies to start to leverage generative AI on behalf of their customers, and what's particularly attractive is that there are a lot of network effects given the firms that have the earliest access to these models, the best data, and the longest track-records are going to be able to build more innovative and distinctive products as a result. That's what we've seen in the traditional AI space, and it will likely be the case for generative AI as the latest instantiation of it.

Q: As growth and buyout investors, how are you gaining exposure to the AI trend?

Mike Zappert: Automation is a long-standing theme of ours, and there's renewed focus on it this year as it encompasses many of the cutting-edge business applications of AI. While generative AI is perhaps the most complex automation tool that exists today, we're focused on many different automation technologies including AI/ML developer platforms, best-of-breed AI/ML apps, and low code/no code.

We also think the platform-as-a-service layer will become increasingly interesting over time. While there are a couple of players in the space at the moment, at some point there's likely going to be a preferred development environment for building AI applications. Many applications will eventually become DIY builds, and so there will be a provider of preference, as is the case with AWS, Azure, and Google Cloud Platform (GCP) for cloud computing. We will continue to watch this space, but given our investment orientation, we expect it will be 3-5 years before there's a purely generative AI company that's investable for us.

That said, all of our companies will certainly feel the impact of AI in terms of its ability to accelerate their roadmaps in certain areas or enable new products or features that wouldn't have otherwise been possible. We believe this is true for companies across all industries, not just those in software and enterprise technology.

> "This platform shift is going to be something that has a profound impact on every one of our companies and not just those in software and enterprise technology.

Tim Millikin: AI is also relevant even in the large cap space. We have invested a lot of R&D into developing AI capabilities to improve the efficacy, for example, of a portfolio company's cyber detection tools, and there are plenty of examples of ways in which AI can be incorporated into our companies' operations to improve their products or operations. But I think the biggest opportunity of all is that as we're investing in companies that are serving specific business functions, whether in software, healthcare, financial services or elsewhere, we're actively thinking about how their business models can be powered by newer AI models. As a firm, we've always been more focused on the application side of new technological waves rather than taking direct technology risk in terms of what will be the winning AI algorithm or platform. This approach has served us incredibly well during past technology waves, and it will continue to be our posture moving forward.

TPG INSIGHTS

Data: The Oil Of The AI Era

AVERAGE AMOUNT OF DATA GENERATED PER MINUTE EACH DAY

Note: latest data as of 2022.
Source: Domo, various news sources.

TPG INSIGHTS

Interview with Art Heidrich

Art Heidrich is based in San Francisco, where he helps **lead TPG's investment activities** in **software and enterprise technology**. Prior to joining TPG in 2011, Art worked in the technology investment banking group at Morgan Stanley, where he advised on and executed M&A and corporate finance transactions. Below, he answers FAQs about generative AI and offers his perspective on the outlook for its commercialization.

Q: What's artificial intelligence and generative AI?

Art Heidrich: The goal of AI is to create intelligent machines. So, AI involves training computers and robots to mimic and learn human behavior. At its most fundamental level, AI is a combination of math and computer science, though obviously there's a lot more at work that goes into it. You can think of "traditional" AI as a catch-all for everything that came before generative AI, and the goal of these technologies in the simplest sense has been to recognize patterns in data and then use those patterns to make predictions.

With generative AI, we are applying machine learning methods to data to generate new content (hence the name "generative") in the form of images, text, audio, or video. What makes generative models distinct from other forms of AI is that they can also produce new instances of data based on an underlying set of inputs. It doesn't necessarily have to be new content, given the fact that the models are working from existing data, but you can think of it as content that's generated by the models themselves. The latest generative AI wave is also unique in that it's broken down the communication barrier between humans and computers. Humans can use "natural language" instead of "code" to interact with AI, which has made the underlying technology far more accessible.

Q: We often hear about different companies building their own AI models. What's a large language model (LLM) in AI?

Art Heidrich: Large language models (LLMs) are critical to generative AI. LLMs are unique machine learning models that are trained by working through a finite body of language or data, finding patterns and affiliations in that data, and then using an algorithm to generate responses to new information or inputs. These models are initially trained using a process called supervised learning, whereby a person repeatedly details the desired output to a particular question to the model. The models are then fine-tuned until they're able to provide outputs independently that are up to a pre-determined level of performance. It's not as though the AI has a mind of its own and is simply generating thoughts, but rather it's seen enough questions, affiliations, and answers that it can recognize patterns and create remarkably responsive outputs, even to entirely new questions. Importantly, the strength of the outputs is a product of the amount of data that the model is trained on, which makes data the fuel that powers LLMs. ChatGPT is the most popular and widely recognized LLM today, though there are a number of other models in development and use.

Q: One of the big breakthroughs in this latest AI wave is that anyone can now easily interact with this technology. What's natural language processing (NLP), and how has it contributed to this breakthrough?

Art Heidrich: Natural language processing (NLP) has been in development for quite some time, and there are a variety of businesses that have long been utilizing NLP in one way or another. A transcription service that takes in language and data and then runs it through a computer to generate a written version of a conversation is a rudimentary example of NLP. At its most basic level, NLP enables computers to read and interpret speech. It does this through two processes called syntactic and semantic analysis. Syntactic analysis identifies the structure and relationship of words in a sentence, while semantic analysis focuses on the meaning of the words and the broader context of the sentence. By parsing speech in this way, NLP techniques can be used to train machine learning and deep learning algorithms such that a model like ChatGPT can now have incredibly fluid conversations with a human that would've been unimaginable just a few years ago.

> **"There's a lot of disruption ahead, though investors and society tend to overestimate the amount of change in the short term and underestimate it in the long term.**

Q: Why is this rapid advancement in AI happening now?

Art Heidrich: It's really the confluence of a number of technologies, including those we've discussed but also advancements in cloud computing, computing power, and big data that have made the current moment possible. Without cloud computing, for example, there would probably only be one or two places in the world with sufficient computing capacity to run these LLMs, and you certainly wouldn't have a product that was commercially viable at scale.

We have far greater access to large datasets today because of the ubiquity of the Internet and the ability to leverage them via modern data science techniques, which have enabled the extensive training of these algorithms. There's also the fact that AI has become democratized over time in terms of what's freely available on the Internet for developers to access and work with. ChatGPT is still incredibly expensive to run computationally, which is one of the primary constraints on the scalability of generative AI. But that will change meaningfully in the next few years as the cost curve continues to come down.

> "Chat GPT is still incredibly expensive to run computationally, which is one of the constraints on the scalability of generative AI. But that will change meaningfully in the next few years as the cost curve continues to come down.

Q: How do you expect generative AI will be commercialized?

Art Heidrich: Commercialization is going to take on many different forms. One way to think about it is in terms of services that are internal and external to the firm. As we've seen with the robotic process automation (RPA) wave, there are a lot of manual tasks that happen repeatedly within a company that can be automated quite easily. AI will likely provide even more sophisticated workflow automation tools to simplify and expedite these routine tasks. This is certainly something that we're thinking about both in terms of investing in the companies that provide these solutions and how more mature companies in our portfolio can adopt these tools to their advantage. In terms of external services, the pace of product innovation will likely pick up significantly in coming years. The speed with which new AI applications can be developed and brought to market means that companies will be able to drive innovation and disruption at a far faster clip than in the past.

Q: But some observers are calling for a pause on the further development of AI. Could that slow down commercialization?

Art Heidrich: I don't think so. The pace of progress has understandably made people nervous. Even AI experts don't always know why a model comes up with what it does because it's weighing an extremely large dataset and then correlating across that data to generate a result. Given the limits to human cognition, it's hard to always know exactly what they're doing. This lack of control and forecastability naturally creates some trepidation in parts of the tech world and public. This will be something we need to work through as an industry, but we don't think the answer is to stop because innovation will continue globally. We believe it's important for the US and US-based companies to play an active part in generating the right solutions and guardrails that hopefully can become the industry norm. I don't see regulation stopping this technology outright, though the concerns that are being raised about the speed of progress are certainly important to ensure we are approaching growth thoughtfully and with a mind towards necessary "seatbelts" in order for this technology to truly succeed.

Q: Will generative AI benefit incumbent or upstart players more?

Art Heidrich: Right now, investment is flowing into this space at a startling rate, but it will be interesting to see what companies can grow into durable, standalone businesses. The reality is that sustainable differentiation in AI isn't about the underlying models. The models themselves aren't a secret. The main axis of differentiation is around data and who has unique datasets to train their models. We believe companies with unique data will be able to generate a flywheel effect whereby the more users and customers use their products, the more proprietary data is created, and the more their product offerings improve as a result, ultimately attracting additional customers. The key for investors will be to find those unique data pools, which will enable sustainable differentiation.

> "The sustainable differentiation in AI isn't about the underlying models … The main axis of differentiation is around data and who has unique datasets to train their models.

Q: To what extent are companies going to have to play defense as this technology threatens their core business?

Art Heidrich: The risk is significant. A lot of companies will need to start focusing on this soon. Generative AI has opened the door for analytics systems that not only parse data but also provide detailed interpretations of that data based on specific business questions. For companies that don't have direct end user or customer relationships, and aren't capturing unique data, generative AI is risky because it can cut out the middleman. On the other hand, companies that have unique data and are a system of record will have a more durable moat. Big picture: there's a lot of disruption ahead, though investors and society tend to overestimate the amount of change in the short term and underestimate it in the long term, which is something we're keenly aware of at TPG.

TPG INSIGHTS

Glossary of AI Terms

- **Artificial Intelligence (AI):** The development of computer systems that can perform tasks that normally require human intelligence, such as learning, reasoning, problem-solving, perception, and natural language comprehension.
- **Natural Language Processing (NLP):** A subfield of AI focused on enabling machines to understand human language. NLP techniques involve breaking down language into smaller components and then using statistical or machine learning methods to analyze and interpret these components.
- **Machine Learning (ML):** A field of artificial intelligence that involves training computer systems to learn from data and make predictions or decisions without explicitly being programmed for each task. Machine learning algorithms are designed to automatically improve in performance over time through a process of model training and tuning.

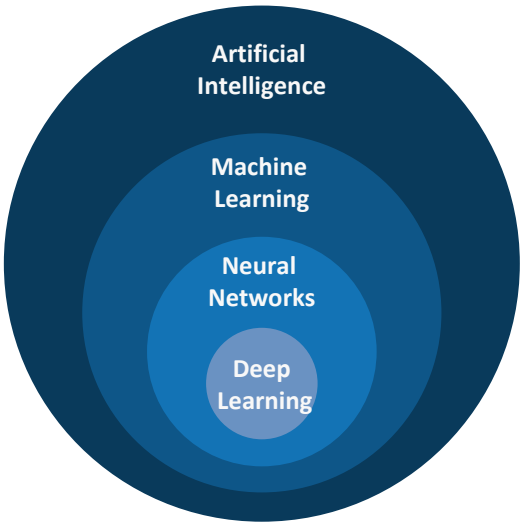

AI is a broad concept encompassing ML, Neural Networks & Deep Learning

- **Supervised Machine Learning:** A type of machine learning where the algorithm is trained on a labeled dataset such that each data point has an associated label or target value. The algorithm is provided with a set of inputs and outputs, and it tries to learn the underlying relationship between them.
- **Unsupervised Machine Learning:** A type of machine learning where the algorithm is trained on an unlabeled dataset, and it tries to find the underlying structure and relationships. The goal of unsupervised learning is to find patterns and structures in data without any prior knowledge of what these patterns might be.
- **Neural Network:** A machine learning algorithm inspired by the structure and function of the human brain that consists of a network of interconnected nodes or neurons that are organized into layers. By leveraging a series of continuous feedback loops, artificial neural networks can improve their predictive analysis over time.
- **Large Language Models (LLMs):** A type of machine learning model that's trained on a large parameter of inputs while using NLP to process queries. The computational power of transformer models to process data sequencing in parallel on massive datasets is the driving force behind LLMs.
- **Transformer Technology:** A type of neural network architecture critical to the field of natural language processing. The transformer architecture is based on a mechanism called "self-attention", which enables the model to weigh the importance of different words in a sentence in order to make predictions.
- **Deep Learning:** A subfield of machine learning that uses neural networks with multiple layers to extract increasingly complex features from data. Deep learning algorithms learn independently to identify features and prioritize data attributes when they are fed with input datasets. The key difference between traditional ML and deep learning is the use of neural networks to train the algorithm rather than human intervention.

Source: Pitchbook, various news sources.

IMPORTANT NOTICE

General

This presentation (the "Presentation") is provided to you solely for informational and reference purposes only and is not intended to be, and must not be, taken as the basis for an investment decision. Unless otherwise noted, the information herein has been compiled as of May 25, 2023 and there is no obligation to update such information. The delivery of this Presentation will not under any circumstances create any implication that the information contained herein has been updated.

Statements in this presentation represent the subjective views of TPG and cannot be independently verified.

This Presentation does not constitute an offer to sell or a solicitation of an offer to buy any securities and may not be used or relied upon in evaluating the merit of investing in any TPG Fund.

Information throughout the Presentation derived from sources other than TPG has not been independently verified and TPG does not assume any responsibility for the accuracy or validity thereof.

No representation or warranty, express or implied, is made as to the accuracy or completeness of the information contained herein, and nothing shall be relied upon as a promise or representation as to the future performance of any investment. Past performance is not necessarily indicative of future results. Differences between past performance and actual results may be material and adverse. There can be no assurance any unrealized values presented herein will ultimately be realized as indicated. These materials may include comparisons of TPG Fund performance against an index (e.g., the S&P 500) or benchmark. The performance of any index or benchmark is provided for reference purposes only. TPG Funds differ from an index or benchmark in that, among other things, TPG Funds are actively invested and is subject to various fees and other expenses. Investors cannot invest directly in the index or benchmark. In addition, there are significant differences between a TPG Fund's investments and any index or benchmark presented, and they will therefore have different risk and reward profiles. Investors should only rely on annual audited financial statements when evaluating the performance of their investment.

Insofar as this Presentation contains summaries of existing agreements and documents, such summaries are qualified in their entirety by reference to the agreements and documents being summarized.

If you believe any content, branding, information or other material incorporated into this presentation has been included in violation of applicable law, agreement, or other restriction, or that any other portion of these materials is otherwise improper, please notify us at compliance@TPG.com.

Forward-Looking Statements

All statements in this Presentation (and oral statements made regarding the subjects of this Presentation) other than historical facts are forward-looking statements, which rely on a number of estimates, projections and assumptions concerning future events. Such statements are also subject to a number of uncertainties and factors outside TPG's control. Such factors include, but are not limited to, uncertainty regarding and changes in global economic or market conditions, including those affecting the industries discussed herein, and changes in US or foreign government policies, laws, regulations and practices. The market analysis, estimates and similar information, including all statements of opinion and/or belief, contained herein are subject to inherent uncertainties and qualifications and are based on a number of assumptions. Opinions expressed are current opinions as of the date of this presentation. Should TPG's estimates, projections and assumptions or these other uncertainties and factors materialize in ways that TPG did not expect, actual results could differ materially from the forward-looking statements in this presentation, including the possibility that investors may lose all or a material portion of the amounts invested. While TPG believes the assumptions underlying these forward-looking statements are reasonable under current circumstances, investors should bear in mind that such assumptions are inherently uncertain and subjective and that past or projected performance is not necessarily indicative of future results. Investors are cautioned not to place undue reliance on such forward-looking statements and should rely on their own assessment of an investment.

TPG INSIGHTS

Issue 3
May 2023

www.tpg.com

TPG

TPG is a leading global alternative asset management firm, founded in San Francisco in 1992, with $135 billion of assets under management and investment and operational teams around the world. TPG invests across five multi-strategy platforms: Capital, Growth, Impact, Real Estate, and Market Solutions and our unique strategy is driven by collaboration, innovation, and inclusion. Our teams combine deep product and sector experience with broad capabilities and expertise to develop differentiated insights and add value for our fund investors, portfolio companies, management teams, and communities.

www.tpg.com

GABE LIPTON GALBRAITH
ggalbraith@advisorstpg.com