

Kennedy Capital Management®

Investment Perspectives

on Opportunities in
Artificial Intelligence





Contents

3

Introduction

4

What is the new wave of computational smarts called “AI” really about?

5

How does development now lead to deployment into the economy at scale?

7

How can we determine economic profitability?

9

How is AI similar and different from the dot.com bubble?

10

What are the limitations?

11

Conclusions

12

Disclaimers



Introduction

ChatGPT's public debut in late 2022 caused a seismic shift in not only the investment world, but also society at large — both from a business investment perspective (the winners and losers) or the perhaps more important technological and philosophical ones. In this paper, we attempt to provide context for:

1

Our view of the core nature of the current iteration of Artificial Intelligence (AI) systems

2

Outline many of the issues that we see these generative types of machine learning (ML) systems bringing to the forefront of the public's imagination

3

Suggest a framework of key points for considering investments in this area

What is the new wave of computational smarts called “AI” really about?



The most recent form of “artificial intelligence,” ushered into the public discussion by natural language tools such as OpenAI’s ChatGPT, is built on machine learning (ML) systems of the past but offers a different approach to their development and use. We have historically seen “smarts” built into programming in many applications – usually in the form of preprogrammed responses to potential inputs. In this way, a system could react to various sources or input conditions – but only those foreseen by the designers and only as explicitly determined by its programmers. In recent decades, related tools have been built on systems ‘learning’ patterns. Human programmers also designed that pattern recognition – but with the system itself developing some level of ‘learning’ based on its inputs. We might see this in the prediction of an autocomplete prompter or a search engine’s suggested finishes of top search queries.

What’s different about this version of AI are both the much more complete responses in natural language – known as “generative” answers – in the form of written text or even computer code, sounds and music, or images – and the way these systems are developed and trained. These models are the logical extensions of the computational inferencing tools made to interpret our environment and “fill in the gaps” and are built on a novel method for how machines could

“learn” more like humans – or at least seem to. This approach mimics human reasoning and the ability to manipulate sequences found in language and processes like computer code – hence the term “Large Language Model,” or LLM, the basis for how these generative systems work. The notion of taking these language inferences and predicting what comes next is the source for the “GPT” term, which stands for “Generative Pretrained Transformer.” The output is not just sourced but built piece by piece, which is why it is called “Generative AI” or “Gen AI.”

Gen AI works by interpreting the query based on a tremendous amount of previously completed inferencing to a network of connections, coupling that with additional processing to construct an answer. This process combines a set of generated terms into a natural language response (or other output), plausibly mimicking the sort of response a human could deliver. Some view this strategy as the same as (or at least analogous to) how humans think or achieve sentience, built upon prior knowledge and learning and constructing new understandings.

Regardless of definitions and views of true intelligence, there is clearly a separation between prior systems and this generation of AI or ML tools. We think the true difference could ultimately be a GPT-type system’s ability to be utilized by otherwise non-technically proficient users, and how it will likely permeate many use cases and lower-level human-dominated tasks. Those developing a proficiency in constructing queries and who can interpret, vet, or otherwise make creative use of the results will have advantages on a new and perhaps uneven playing field. Understanding the investment takeaways of how these tasks will be redeployed into the economy at scale can help us both as individuals and as investors.

How does development now lead to deployment into the economy at scale?

Prior generations of technological enablement can provide us with insight into how these systems work their way into our everyday lives and impact our work and human productivity. Gen AI and language-based models dramatically change how users work with data and could be considered as fundamental to data use as the web browser, with likely similar impacts to both users and use cases. We believe, in the long run, it's likely that the real benefit will accrue to those who leverage AI to drive efficiency gains and/or revenue growth over the next ten-year window.

So where are some promising areas? Some of the most cited areas for deployment might be generating background or basic text for articles, online reports, potentially dialogue for productions, or even computer code. One example of a function that could get commoditized unexpectedly early could include drafting legal documents, whose sometimes 'boilerplate' sections might lend themselves to repetitive AI tweaking. Beyond these textual cases, some observers grow alternatively excited and worried when contemplating the much-improved ability to mimic or produce what would have previously been considered possible only by human creatives – artwork, photographs, and music.

“There are likely some instances where Generative AI will replace human-performed work, however, it may also allow more human insight, which can move that human's work up the value chain.”

With announcements of initiatives and development plans happening at a dizzying pace, we expect that there will be many false starts and unfinished or abandoned projects as the developments are deployed into the economy at scale. We think what's most important is how integration can happen. We believe AI-based tools could accelerate a significant portion of business processes that haven't moved towards automation yet. There are likely some instances where Generative AI will replace human-performed work, however, it may also allow more human insight, which can move that human's work up the value chain.

Promising areas for AI deployment



Background text for articles
Boilerplate sections for
legal documents



Productions
dialogue



Computer
code

Assessment of these impacts suggests significant productivity boosts for lower-end, entry-level tasks (we'd estimate perhaps up to 40%). For higher value considerations, however, this productivity benefit is likely much less, and could even be neutral or negative. We'd expect that having these kinds of AI-leveraging capabilities will become a 'must-have' in many toolkits, much like analytics is currently for business process outsourcers, where the ability to employ business analytics is required to be invited to bid for projects. Consultants will need to be able to work with and integrate Generative AI tools as part of overall solutions in order to be competitive. ***We think that productive development and the long tail of sustainable, profitable opportunities likely lie in the ability to integrate these tools into an enterprise's specific domain expertise.***

We should also consider the human element inherent in developing these tools in the first place. Ironically, in order to make these AI tools effectively work (even if, at times, error-prone) seems to require human intervention – not just in coding the system, but in that actual training of the system's responses. Through a process called reinforcement learning with human feedback (RLHF), Gen AI's responses are scored by human evaluators and ones that are closer to the correct meaning (or "right") are valued higher by the system, reinforcing that set of inferences and steering it towards being more consistently accurate. This process is not itself without error and can have the impact of subjecting the human evaluators to precisely the kinds of wrong, biased and sometimes disturbing answers that cause concerns around AI systems.

The impact of AI's demand was clearly seen in Nvidia's May earnings report and the company's subsequent jump to over \$1 trillion in market value. It is clear that some of the early winners have been from the same mega-cap tech winners of the last decade – those with the resources to both drive the market for AI investing in its development, and harvest the initial benefits.

“Ironically, in order to make these AI tools effectively work (even if, at times, error-prone) seems to require human intervention.”

We think that the initial lift to build systems running these LLMs at the scale sufficient for them to be useful in a general purpose way has required entities with both the financial resources and the hyperscale system expertise needed to actually implement them, constraining that initial set to the level of Microsoft, Google, and Meta. There are undoubtedly others but not visibly at scale yet. While some, particularly OpenAI / Microsoft's ChatGPT, for now have a time advantage in the source material that they've trained on (potentially as much as 12 months), we don't think that will necessarily remain durable. ***As the use cases evolve into more specialized markets or tasks, perhaps implemented on top of this prior training, the door opens to more nimble implementations with more efficiently built systems at a smaller scale.***


What seems evident is that the early fundamental winners have been and will likely continue for some time to be the hardware providers for the computing infrastructure required (Nvidia, but also others like AMD or SuperMicro). Any AI startups would require significant capital to be on par with these industry leaders. In addition to their chip design expertise and lengthy history with parallel processing for hyperscale data centers, Nvidia has also ridden their mature suite of software design tools that make their products into an ecosystem. Others will follow, and being more open source may offer some long-run advantages. Outside of the computing and server providers, it's also evident that the uses and leverage for these novel approaches to data will provide a long-term benefit to networking and telecom equipment providers as AI adds another driver to long-term demand growth.

How can we determine economic profitability?

What is less clear to us, given the incredible cost to build these LLM AI systems and the massive computational power required to handle many complex queries (or even simpler ones), is the viability of their individual economic profitability. We may be seeing some loss-leader behavior or market-share land grabbing on the part of these initial deployments to get users accustomed to or integrated into a given ecosystem. While a logical strategy, that also doesn't insure against a newer version supplanting even successful ones. The landscape of these ecosystems to develop and deploy the code to run these models will likely become quite crowded, with again many false starts and new winners.

For those looking to invest long term, it's worth considering both the strategically advantaged positioning of mega cap tech players and the risks involved with any upstarts. What we've witnessed over the last decade as the large hyperscale tech players have made themselves more entrenched and more profitable, is that it allows them to make massive long-term investments, and AI could continue to fortify their positioning. ***On the other hand, as we've seen multiple times, technology transitions give rare opportunities for doors to open and paradigm shifts to happen.***

Consequently, while much near-term economic value may reside at these bigger companies, we can make the case that as that economic value builds, it may only help some of them grow into their already lofty valuations. ***Over the longer term, we think it's likely that as AI diffuses into business practices and the workplace we'll see other areas of Tech, particularly information services, seeing significant benefits.*** This would be akin to where much of the post-2000 developments out of the internet bubble found its winners – down market, and in stocks that wound up compounding value over significant stretches – companies such as Salesforce or Broadcom / Avago, which started out as small and mid-cap names.



Where valuations have perhaps not decoupled from current economics, we think that there are several parameters investors should consider in evaluating true market positioning and business potential. These revolve primarily around a company's structural value and the resiliency of their business content and its adaptability. It can take the form of the specific uniqueness of their content (proprietary or not) and the importance of being right (accuracy). The value, perception and use cases of that data or content will significantly impact on any value of AI enhancements.

Finally, a company's willingness to proactively address implementing AI and adapting it to their specific business use cases will steadily translate into perceived value, especially if it can enhance the value of that critical content and data. We can see many potential opportunities across multiple sectors outside of Tech itself where this is possible, including Industrials, Healthcare, and Financials. We think that there are a great many human-driven tasks that will find themselves susceptible to automation by AI tools. The key, of course, will be to leverage the technology and the human activity it replaces into higher productive value tasks for those same people. We believe that in many of these cases, the places where these types of implementations will be accretive to corporate value will likely not be with the mega cap tech stocks but rather in smaller, more nimble firms.

Potential AI opportunities across multiple sectors



Industrials/Manufacturing
inventory tracking, quality control, preventive maintenance



Healthcare
disease detection, drug discovery, virtual assistants



Financials
fraud detection, algorithmic trading, risk management

How is AI similar and different from the dot.com bubble?

We have often been asked where we are analogous to the 2000s Tech bubble. This was a similar time of revolution, driven by the dawn of broad societal use of the internet, with dot.com winners and blowups, and the explosion of data traffic driving networking. The rise of a new generation of computing superstars emerged during this time, an environment that eventually led to some of the leaders of today (including Google and Meta). Our view is that we're still at quite early stages, perhaps before the analogous peak of the hype in early 2000. There may be some near term sorting out of winners and losers, but we fully anticipate that opportunities will present themselves down market for significant value creation.

We would, however, advise caution as high expectations that likely get priced into perceived beneficiaries of AI adoption must be considered when making investments over a multi-year timeframe. We can look at the leaders during the Tech bubble (companies like Cisco and Intel) and see that their expected sales and earnings growth in 2000 were meaningfully ahead of what transpired over the next couple of years, with a commensurate impact on stock returns. We should be watchful when companies see their EV/sales ratios hit 20x or more, no matter what the expected growth rate. We feel that there will be plenty of investment candidates with reasonable valuations and solid business cases, with the key, as always, being to identify potential ideas early and having the fortitude to harvest gains when appropriate.

“There may be some near term sorting out of winners and losers, but we fully anticipate that opportunities will present themselves down market for significant value creation.”



What are the limitations?

The cycle is fairly early and much of what will ultimately arrive as limitations is not fully apparent but primarily includes regulatory concerns and costs. We'd note that, given the incredible cost to build these LLM AI systems and massive computational power required to handle many complex queries, it's questionable whether on an individual basis that they're even currently economically profitable. We also foresee business use cases getting upended, and the downstream effects of AI-driven content that may overwhelm desired output. All of these are potential issues that could not only impact but substantially derail the development and deployment of AI.

From a regulatory standpoint, even more than datacom / telecom or timeshifted streaming vs. linear media, AI is a technological wave that industry must get on the right side of, or risk having it regulated for them. As we've seen with privacy and copyright laws, trying to put the genie back in the bottle is quite often a losing battle. The primacy of AI, especially generative AI – looks to replace human actions and activities, even creative ones, and at least partially through appropriating their prior work product – calls into question the ownership of the content that is currently available in an accessible fashion. We could see a sudden suppression of significant portions of the internet, sealing content into walled gardens to prevent that content from serving competitors in a way not previously thought possible. We are already witnessing creative producers, including writers and actors, facing the reality that, without protections, they could in the very near future lose control of their professional images.

“From a regulatory standpoint, even more than datacom / telecom or timeshifted streaming vs. linear media, AI is a technological wave that industry must get on the right side of, or risk having it regulated for them.”

From a practical business standpoint, we see substantial concerns about how quickly generative AI is being thrown into seemingly every possible use case to see if it sticks, offering the real potential for the actual degradation of some business tasks either for the user or a potential customer. It's already bad enough to get stuck in phone tree limbo while contacting a customer service line; it's not too difficult to imagine having a rules-based generative AI as its replacement trying to solve a customer service issue and instead making it worse.

Finally, as a society we will need to ask the questions as to the true societal benefit when much of our day-to-day experience with Gen AI will come in the form of AI-produced content in our social media news feeds, spam or malicious emails in our inboxes, and potentially misleading copy posted as news articles. These are real, intractable problems that form some of the strongest basis for a robust regulatory framework to be defined before it becomes too little, too late.

Conclusions

The emergence of AI, in the form of natural language interfaces supported by these LLMs, has captured both the public's consciousness and investor excitement. Much like the groundswell of investment potential that came from the widespread adoption of the internet, we anticipate both short term and long-term opportunities to exist. We would caution against getting caught up with too much exuberance without consideration for the expectations required of some valuations. As the linearity of investment cases is in our opinion not clear, understanding their evolution will require vigilance as it unfolds. We believe that we're in the early part of this emerging cycle - the hardware buildout phase - with much of the revenue driving applications and genuine long term value creation ahead of us.



About Kennedy Capital Management LLC

Founded in 1980, St. Louis-based Kennedy Capital Management LLC (KCM) delivers investment strategies to corporate and public pension funds, endowments, foundations, multi-employer plans and high-net-worth individuals. As a registered investment adviser, KCM specializes in the management of micro, small and mid-cap strategies across the growth-value spectrum and as of June 30, 2023, managed \$4.49 billion in assets. KCM integrates environmental, social and governance (ESG) considerations into its research process and has over 20 years' experience managing socially-responsible client accounts. For more information, visit www.kennedycapital.com.

Advisory services offered through KCM, an SEC Registered Investment Adviser. All investing involves risk. The value of an investment may increase or decrease over time. Past performance is not a guarantee of future results. Future returns are not guaranteed, and a loss of principal may occur. The material and information contained within this presentation are not meant as investment advice and should not be construed as a recommendation; or an offer to buy or sell; or the solicitation of an offer to buy or sell any security, financial product, or instrument; or to participate in any particular trading strategy. All opinions expressed herein constitute KCM's judgement as of the date of the article and are subject to change without notice. Statements made are not facts, including statements regarding trends, market conditions and the experience or expertise of KCM. All statements are based on current expectations, estimates, opinions and/or beliefs of KCM. There can be no assurance that any KCM strategy or investment will achieve its objectives or avoid substantial losses.

Certain information contained in this document constitutes a "forward-looking statement," which can be identified by the use of forward-looking terminology, such as "may," "will," "seek," "should," "expect," "anticipate," "estimate," "intend," "continue," "believe," or the negatives thereof or other variations thereof or comparable terminology. Due to various risks and uncertainties, actual events or results or the actual performance may differ materially from those reflected or contemplated in such forward-looking statements, including a complete loss of investment.

2023090026