

Blue Tower Asset Management, LLC
311 RANCH ROAD 620 SOUTH
SUITE 111
AUSTIN, TX 78734

October 21st, 2023

The Blue Tower Global Value had a -2.19% return net of fees (-1.94% gross) in Q3 2023 which was a smaller decline than most broad market indices. This quarter saw a sharp increase in the yields of US government treasuries as the Federal Reserve pushed back their forecasts of interest rate cuts. This decline in the prices of long-term bonds also led to a selloff in equity markets. Our biggest source of outperformance this quarter was the appreciation of our investment in Georgia Capital which we discussed in our Q1 2021 letter.

Please review your individual account statement to see the returns in your account as it will differ from the composite return. The sanctions on Russian investments, which prevent us from buying or selling Russian stocks, caused a dramatic increase in the internal dispersion of returns for our composite. Accounts that were opened after March 2022 will not own any Russian assets and will instead have a larger weight of all the remaining positions in the strategy. As there have been no major changes in the status of our Russian investments, I will limit discussion of them in this and future letters until major developments occur.

In 2023, the year's biggest technological development has been the rapid advances in generative AI. The release of GPT 4 by OpenAI and widespread use by the public has led to an AI innovation race between the tech giants and a large rally in their stocks. The S&P500's return of 13.07% for the first 3 quarters of 2023 was skewed by the returns of the largest 7 tech companies¹ (which collectively had an equal-weight average return of 87.83%!). If you looked instead at an equally-weighted average of the 500 companies in the S&P 500, its equal-weight index only had a return of 1.79% in the first 9 months of 2023. This huge rally in tech stocks is due to the excitement around recent developments in generative AI.

The Generative AI boom

AI has been around in various forms since the 1970s. Most of the research in the last two decades was focused on various types of neural networks and machine learning through statistics, decision trees, and linear regression. These methods are very powerful and have led to algorithmic tools with huge economic value including the ubiquitous recommendation algorithms used in video streaming and social media platforms.

The recent advances are an evolutionary jump and paradigm shift in a new model of AI. When people are talking about the recent AI boom in 2023, they are referring to transformer models. OpenAI's GPT acronym stands for "Generative Pre-trained Transformers." These transformer models are neural networks that track relationships in sequential data such as sentences. The networks use an attention mechanism to create context for subsequent tokens, words, based on

¹ The 7 largest US tech stocks are Apple, Microsoft, Alphabet, Amazon.com, Nvidia, Tesla, Meta

previous tokens. Tokens are the basic units of text that an LLM AI uses to process and generate language. Transformers were first described in a research paper “Attention is All You Need” from Google in 2017².

The major advantage that transformers have over previous techniques is that rather than trying to “remember” and process each token in turn, the network places special attention on more important tokens regardless of how early they may have appeared in an input token sequence. Previous models tended to “forget” early tokens before reading the end of long input sequences. This allows for more efficient learning requiring far less compute power to learn from much larger datasets. In general, the larger the training dataset of a neural network, the more powerful the resulting neural network will be. This current generation of large neural network models are the most powerful yet. Transformers are now replacing convolutional and recurrent neural networks, the previous state of the art neural network technologies, in many of the applications that used them.

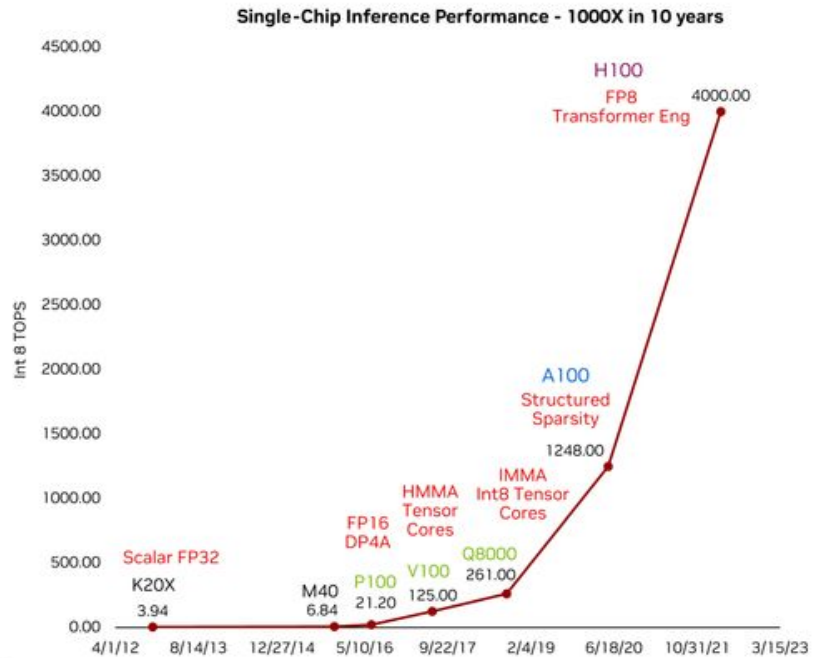
In addition to the use of larger datasets, the training speed of AI models has increased dramatically. Nvidia’s stock almost tripled in the first 3 quarters of this year with a 197% gain, and a large reason for this is the huge role they have played in recent AI improvements. Nvidia’s single GPU AI training speed performance has increased by a dramatic 1000x in 10 years with only 2.5x coming from Moore’s Law³ driven increases in chip density. Besides better chip manufacturing, there were three other improvement factors at play: simplifications in number representation for the weights of the neural networks, more complex mathematical instructions for reducing the computational overhead involved in mathematical calculations, and increased neuron sparsity (in neural networks, some neurons are useless and can be pruned from the network without reducing performance significantly). In addition to these single GPU improvements, Nvidia also made improvements in data center scale architecture that allows groups of GPUs to work more efficiently together.

² “Attention is all you need.” <https://arxiv.org/abs/1706.03762>

³ Moore’s Law is the phenomenon observed by Gordon E. Moore in 1965 that the number of transistors in a computer chip doubles every two years, as the number of transistors increases, so does processing power, and that this process is exponential. As transistor density increases, the cost per transistor decreases. This exponential growth in processing power has slowed in recent years.

Gains from

- Number Representation
 - FP32, FP16, Int8
 - (TF32, BF16)
 - ~16x
- Complex Instructions
 - DP4, HMMA, IMMA
 - ~12.5x
- Process
 - 28nm, 16nm, 7nm, 5nm
 - ~2.5x
- Sparsity
 - ~2x
- Model efficiency has also improved – overall gain > 1000x



NVIDIA

Source: IEEE Spectrum⁴

It is noteworthy that the vast majority of the improvement came from hardware architectural and software data improvements, rather than transition density. These improvements were likely the low-hanging fruit of training speed improvements as researchers will eventually converge on an ideal architecture. The simplification of going from 32-bit to 8-bit floating point numbers for measuring weights is a one-time gain that can't be repeated again. I expect the rate of improvement to slow down over the next ten years and eventually approach the levels of Moore's Law improvements in chip efficiency. The historical trend for computer hardware is for it to eventually be commoditized, and I believe this will eventually occur for Nvidia's GPUs as well.

On February 24, 2023, Meta released Llama, their LLM project, with the code being open-source but the training weights being kept proprietary. Nonetheless, within a week the training weights were leaked for the project. This release of the weights for 65 billion parameter model was a huge gift for open-source programmers to begin experimenting upon. Within weeks, other developers were creating innovations building on the model. There has been a race towards developing low-budget fine-tuning projects that can be trained for specific applications at the cost of a few hundred dollars in some cases. Some of these LLM projects can even be run on individual computers rather than data centers.

After the voluntary release of Llama's architecture as open-source, the leak of the model weights ended up being a blessing in disguise for Meta. The innovation and improvements that followed

⁴ IEEE Spectrum "The Secret to Nvidia's AI Success" <https://spectrum.ieee.org/nvidia-gpu>

were all built on Meta's AI codebase. They were able to get these innovations without needing to pay any of the open-source developers for their efforts. This also demonstrates that the business advantage of the large tech companies is not in their models or AI codebase. Any innovations that the companies develop are unlikely to stay within the walls of the companies due to the mobility of researchers being poached between companies by recruiters. However, as these new model improvement technologies are made, the big tech companies have a scale advantage. They will be able to employ it into models that have more data, more compute budget, more parameters and therefore more intelligence in the resulting neural net.

The moat of these tech companies then is not their code but their proprietary data built up over decades, the scale and infrastructure investment of their data centers, and their distribution ability for any resulting AI products or services. This dynamic means that the main winners of the new generative AI boom will be the incumbent tech giants.

Other potential winners will be companies with large propriety datasets that are unique within their industry. For example, I do not believe that there are any companies in the US with as large of a dataset of the borrowing behavior of short-term, subprime borrowers as Enova International (ENVA), a portfolio holding of our firm. Companies that fail to take advantage of newly developing AI tools will be at a great disadvantage to their competitors who do.

Best uses of generative AI today

A major problem with generative AI is the hallucination problem where the AI will produce a response that contains false information as though it was correct information. A famous example of this phenomenon is the recent experience of Steven A. Schwartz, an attorney who used an AI to generate a legal brief. In doing his research to support a response to his opposing counsel's motion to dismiss, he consulted with ChatGPT. The AI offered up several fake cases that supported his position along with citations and quotations from the legal opinions. When he asked about the veracity of the cases, ChatGPT assured him that the cases "can be found in reputable legal databases such as LexisNexis and Westlaw." After Schwartz filed the brief, the opposing counsel brought to the judge's attention that they could not locate the cited cases. The court was not amused by this error and fined the attorney submitting the fictitious citations.

Hallucination responses containing convincing yet false information will cause issues for other use cases, such as using it to generate software code. Many of these software segments that appear to be valid at first glance will fail to execute properly when used.

Due to these current issues, I believe the best use of generative AI today will be in submitting prompts that one already knows the correct answer to. There is a great deal of repetitive busy work in any profession, and generative AI can boost productivity in taking care of these tasks. A coder, for example, could use generative AI to generate a short code snippet rather than searching on a search engine or GitHub for the appropriate code. Graphic designers can generate quick image elements. Customer service representatives can use it to write common responses.

On November 1st, Microsoft will be launching their copilot AI assistant for Microsoft 365 customers on certain business and enterprise plans. Google recently launched a similar AI assistant as an add-on for certain corporate clients using Google workspace. It is too early to

judge these AI assistants, but I believe they have great potential to be productivity boosts for many office workers.

Blue Tower's philosophy and the use of AI in investing

AI-focused funds have had a history of poor performance. AI systems have advantages in rapid-decision investing such as high-frequency trading and market-making, but their performance has been poor at handling long-term investments. The use of AI for directing trading on securities markets has been a goal of asset managers since the most rudimentary AI systems in the 1970s. Human portfolio managers can get sick, go through personal difficulties, have off days, cognitive biases, or begin suffering mental decline with aging. An AI that does not suffer from these weaknesses would be at a great advantage.

There has been a boom recently in AI-utilizing investment strategies. However, the performance of many of these AI-directed strategies has been disappointing so far. Eurekahedge, a data provider specializing in hedge fund information, created an index tracking “a broad measure of the performance of underlying hedge fund managers who utilize artificial intelligence and machine learning theory in their trading processes.”⁵ The Eurekahedge AI Hedge Fund Index shows that these AI-utilizing funds have underperformed most broad indices of hedge fund performance over the past five years and also stock indices like the S&P 500. The conclusion I draw from this is that AI currently in use is not capable of outperforming expert portfolio managers at long-term fundamental analysis of securities. Long-term fundamental investing will continue to be managed by expert human portfolio managers, but software tools will continue to grow in importance.

My goal for Blue Tower remains the same as it has always been: to combine data analytics technology with the qualitative judgment of business models that has always been the core of classic value investing. The combination of human judgment with algorithmic screening tools is more powerful than either of them in isolation. AI will be able to summarize documents to highlight key points of interest, direct human attention to anomalous data points, identify liquidity holes as they are forming, and in general replace human “busy work” with the power of pattern replication. I can imagine a future of a human portfolio manager being advised by a dozen different AI agents all trained to look for a different type of financial market opportunity.

Understanding and using these AI technologies will not solely be the domain for software developers and tech workers. Regardless of what industry they are working in, all professionals should consider ways they can use AI to boost their productivity of work. The old ways of doing things will not be able to compete in the future. Eventually, AI will be as ubiquitous in business as the use of electronic spreadsheets.

⁵ <https://www.eurekahedge.com/Indices/IndexView/Eurekahedge/683/Eurekahedge-AI-Hedge-fund-Index>

I look forward to receiving your questions and comments, and appreciate your continued support of our firm.

Best regards,

Andrew Oskoui, CFA

Portfolio Manager

Disclaimer: This commentary does not represent a recommendation to trade any particular security, but is intended to illustrate Blue Tower's investment approach. These opinions are current as of the date of this commentary but are subject to change. The information contained herein has been obtained from sources believed to be reliable but the accuracy of the information cannot be guaranteed. Past performance is no guarantee of future results.