Unleashing the Power of Generative AI for Data and Analytics

By John Santaferraro, CEO and Head Research Analyst of Ferraro Consulting
and Amit Soni, Senior Product Manager at Google

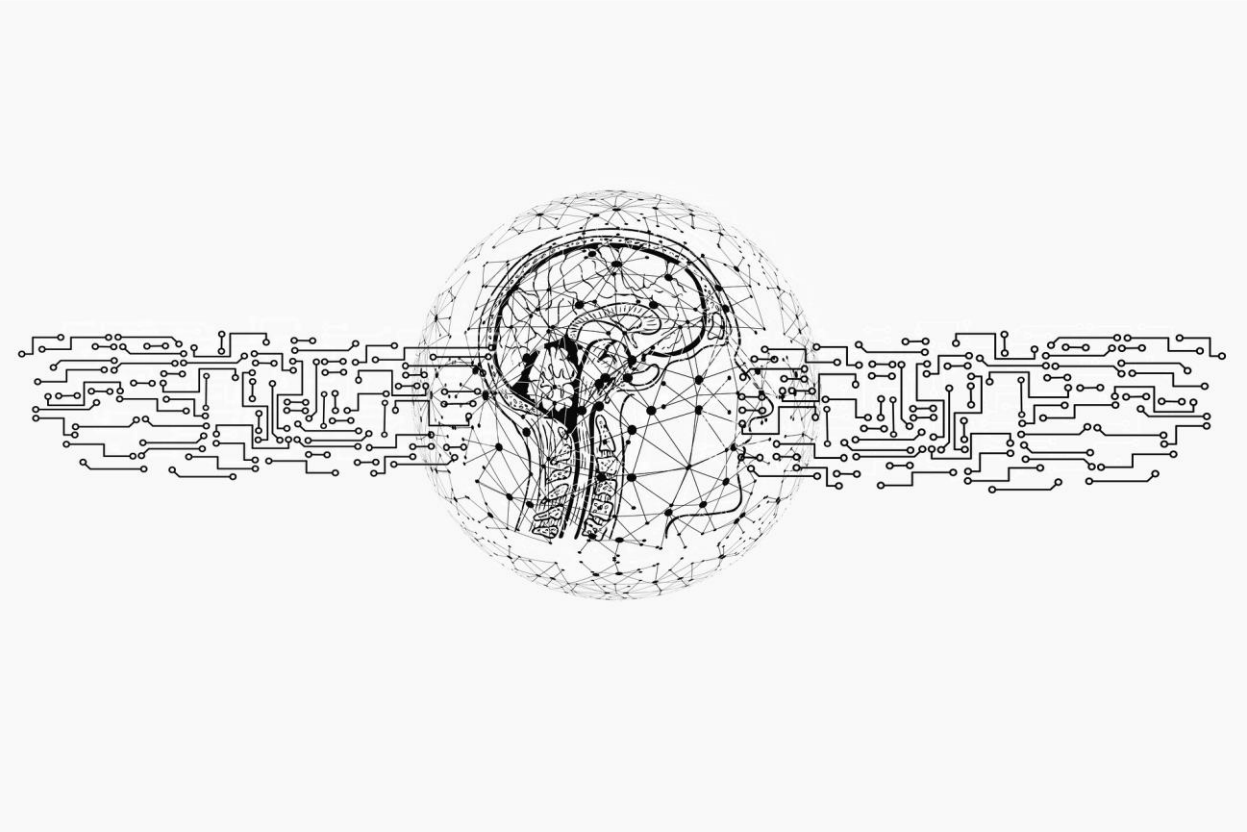

Published by **FERRARO** Consulting
Sponsored by ThoughtSpot.
www.ferraroconsulting.com

In the last year, Generative AI has created the wild, wild West for artificial intelligence. The tech world has gone from tacking ".ai" onto the end of their URL to the point where every vendor is now an AI company. The shift from a mobile-first to an AI-first approach was almost instantaneous, and whether we like it or not, we will all be changed by this latest invasion of AI.

In the words of Buckminster Fuller, the great architect, mathematician, and philosopher…

> "If you want to teach people a new way of thinking, don't bother trying to teach them. Instead, give them a tool, the use of which will lead to new ways of thinking."

With generative AI we have a tool that will lead to new ways of thinking about how we work, how we learn, how we create, how we interact, and how we succeed. Instead of fearing the new model, great thinkers and great leaders will ride the wave of innovation to new levels of success, new models for doing business, and new ways of living. Every leader is faced with a decision on how to implement general AI and generative AI as part of their digital strategy.

The data and analytics market is no exception to the great AI expansion of 2023. Every data management, business intelligence, and analytics platform already uses generative AI in some form. However, with a barrage of marketing madness, it is unclear who to believe, where to invest, where to use AI, how to use AI, and what percentage of resources should be devoted to AI. Fortunately, the maturity and measurability of data management make it easy to separate the leaders from the followers.

Generative AI is Already Transforming Data Management

The universal proliferation of generative AI in 2023 and the start of 2024 eliminates a "do nothing" strategy for every serious leader. To do nothing is to fall behind at such lengths that it may be impossible to catch back up. This urgency is not just because everyone is using generative AI, the imperative comes from the present proven value and the obvious future benefit of AI.

The real fear of falling behind is already motivating leaders to drive implementation and accurately measure the return on their investment. Even with many unanswered questions, the potential benefits and advantages of both general AI and generative AI far outweigh the risks. There are three strong areas of generative AI value creation: acceleration, innovation, and digital dominance.

Acceleration

Early adopters are already seeing a boost in data management productivity and the acceleration of time to insight based on current AI automations and recommendations. On top of AI acceleration, the creation of copilots using generative AI fuels developer and analyst productivity to even greater levels of efficiency and excellence.

Innovation

Innovation is almost always dependent on the speed at which organizations can iterate cycles of experimentation and fine-tuning new models. Increasing the speed of data management also accelerates both analytical and business innovation cycles, giving digital leaders what they need to establish data-driven differentiation.

Digital Dominance

Finally, AI adoption addresses a major shift in the digital economy. It is no longer enough to work toward digital transformation as the end state of modernization. Today's digital economy demands more. The combination of acceleration and innovation surpasses transformation and helps leaders establish and maintain digital dominance.

Identifying the Best of the Best

As Amit Soni explained in his paper, *Analytics alone can't win a battle!*, while generative AI affords great opportunity, it does not come without risks. The analytical prowess of a generative AI engine is not enough to propel it forward successfully. By now, everyone is well aware of the challenges of quality and accuracy when using generative AI. Unfortunately, platforms like ChatGPT have a good imagination and have no problem using it. There have already been a series of incidents where the ungoverned use of generative AI has generated as much cost as content. Soni explains the need for metadata, semantics, and feedback to guide AI to the right answers. Soni's early observations align perfectly with the Ferraro Consulting point of view on identifying the best of the best among data management and analytics vendors in the use of AI in their platforms.

The Crucial Role of Metadata for Generative AI

In the world of data management, metadata has already been established as the de facto way to govern the use of data and guide users to accurate insight for business decisions. Metadata is data about the data and often includes information like data definitions, ownership, lineage, quality, validation, and usage. More advanced uses include data certification, profiling, automation, and recommendations. The most advanced use cases include the automatic generation of metadata. Metadata can be applied to any type of data asset. Think

of metadata as a dictionary or glossary utilized to ensure that content being developed is true to the intended use of the data.

Just as enterprise metadata governance and management are fundamental to enabling safe, scalable, and compliant analytics solutions, they are equally foundational for the use of AI and automation. Given access to metadata, generative AI can govern the use of data and ultimately produce more accurate insights for users that are often unaware of the importance of governance. For example, instead of just giving insight to a user, it is possible to provide an explanation of the insight being returned or warnings about potential inaccuracies. Or, instead of waiting for users to ask where the insight came from, the explanation can offer information on the source data used to produce the insight including information data lineage, such as, where the data came from, how the data was transformed, and to what degree the data was validated or certified.

The Critical Role of Semantics for Generative AI

One of the biggest challenges in working with generative AI is the difficulty getting accurate results specific to the business. For this reason, data and analytics vendors that support a semantic layer, or a plain language description of the business, will greatly increase the output of reliable insight. The accuracy of the business model and the specificity of business processes descriptions both translate into competitive use of AI.

In order to produce accurate results, both general AI and generative AI require a logical representation of the business. A business model or semantic layer places guard rails on the use of AI and the resulting insight. While it is important to model the business at a high level, granularity widens the impact of AI across the entire organization. It is especially important to understand business process interactions and dependencies, as well as the potential value associated with each process. Defining roles and responsibilities within the organization and mapping them to the business processes extends AI impact even further. Role-based AI enables the fine tuning of automation in light of how people interact with data and technology.

The Critical Role of Feedback Loops for Generative AI

Because generative AI is young and immature at this point, it makes sense to build checkpoints to ensure that decisions are being made with accurate insight. Leading vendors are already establishing both technical and human feedback loops in their products to continually improve accuracy.

From a technical standpoint, some vendors have limited generative AI access to metadata and the semantic layer. Generative AI can deliver data as part of a response, but not alter the data. Since these vendors have already created a highly governed environment for the use of data, this separation reduces the risk of

hallucination. Other vendors have implemented a final validation step that feeds the results from the generative AI engine back into the analytics engine. The end user results come directly from the analytics engine.

Along with technical feedback, leading data management vendors are using a combination of human and technical feedback loops to increase generative AI accuracy and reliability. Reinforcement Learning with Human Feedback (RLHF) is a methodology designed to optimize language models to eventually achieve the best possible results. In simple terms, the language model behind generative AI is already trained to a certain extent. Human experts then rank generative AI outputs and the rankings are used to create a reward model. The reward model is used as a guide for reinforcement learning, retraining the language model to produce better results.

The Critical Role of General AI for Generative AI

The greatest potential for generative AI in data management does not emanate from tacking generative AI onto existing technology. The phenomenon is as obvious as a glance at where the internet is today, compared to 30 years ago when people were taking their existing content and just posting it on the internet. Deep value is achieved when generative AI is combined with general AI to provide a common infrastructure for all enablement.

Having sat through dozens of generative AI demos, it is instantly obvious when someone has simply bolted ChatGPT or some other generative AI platform to their existing software. In almost all cases where this is happening, there is much hype and little value; most of the implementations are showing what other AI-enabled vendors have been doing since 2019, or earlier.

To realize the importance of a unified approach for general AI and generative AI, data management vendors must bring the two worlds together. For example, metadata services vendors like Unifi were already using general AI for automated usage analysis, sensitive data discovery, relationship discovery, data tagging, data validation, automated inventory, automated cataloging, search recommendations, and relevant dataset recommendations. When generative AI has access to the results from the general AI capabilities, it can do far more than simply access data or other non-AI enabled platform functions.

Recommendations for IT Leaders and C-Suite Executives

Identify opportunities for generative AI in the analytic lifecycle.

There are two natural starting places for exploration in the use of generative AI for data management: data integration and business intelligence. Since data integration and transformation often consumes up to 70% of every analytics and AI project, it is prime real estate for productivity gains. Business intelligence also rises to the top because of the way generative AI can speed time to insight, as well as extend the reach of insight to a broader audience including non-technical users.

Select data and analytic vendors based on their use of generative AI.

The most advanced data management vendors began using active metadata and machine learning back in 2018-2019 to automate formerly manual functions, and to provide users with workflow recommendations or automated insight (see Addendum). These vendors had a core AI engine already built into their platforms; and most had a rich set of active metadata. Some vendors also already had a business semantic layer to help users define their business models within the platform. As you consider the licensing of software for data management, look for vendors who meet the criteria defined in this paper: rich metadata, business semantics, feedback, and a central AI engine.

Foster a culture of innovation in the use of generative AI.

Opening generative AI too far can result in unforeseen risk. However, restricting access to generative AI can stifle potential innovation in your organization. A model that opens access in a protected environment and governs the process for moving into production provides the right balance. Consider a sandbox environment where access requires a business sponsor, this guarantees value creation. In the sandbox, provide minimum governance to ensure governability in production. In the move to production, select applications with measurable business impact. Finally, add complete governance in preparation for release.

Create an open and continuous learning environment around generative AI.

In order to get the most value out of generative AI in the use of data and analytics, assume that the general population only has a basic understanding of how the technology works and how it can be best used. Providing hands-on training with specific use cases is a great beginning. However, because AI technology is changing so fast, ongoing training programs with built in incentives for participants will keep employees and stakeholders engaged in continuous learning. Over time, you can develop role-based learning for both the business and technical teams.

Measure the business impact of generative AI.

While generative AI in the data and analytics space will surely reduce overhead and speed time to insight, it is still necessary to measure the business impact of AI. The most obvious starting place is in the analytics organization where you can look at headcount, response times for business requests, number of business users, and time spent finding, preparing, and analyzing data. If you can measure how analytical productivity impacts your business, that will ensure the continued funding of your AI capabilities.

Move Forward Well-Defined Buying Criteria

In conclusion, the keys to exceptional use of generative AI in data and analytics includes active use of metadata, the use of a business model or semantic layer, integration with general AI already used in the platform, and both technical and human feedback loops. Ferraro Consulting recommends that buyers consider these as checkpoints for all data and analytics purchases. AI augmentation from both general and generative AI will continue to expand, consuming more formerly manual or human activities in data management. Organizations investing in vendors leading the pack for AI augmentation can expect to streamline data management, speed the delivery of insight, and free valuable resources for use in innovation.

ThoughtSpot Aligns with Generative AI Best Practices

ThoughtSpot innovations over the last 10 years have led their way to a leadership position in the use of AI in business intelligence. Set apart from other business intelligence companies who are simply tacking generative AI onto their existing products, ThoughtSpot has excelled in all areas that align with the Ferraro Consulting POV on best practices for AI-enabled business intelligence: a core AI-engine with existing AI enablement, the use of business semantics and metadata as the interface between generative AI and data, and human in the loop feedback with reinforcement learning to improve relevance and accuracy of AI-enabled search.

ThoughtSpot Existing AI-Enablement for Generative AI Leadership

ThoughtSpot began augmenting their business intelligence and analytics platform with AI and ML all the way back at their start in 2012. In 2019, Ferraro Consulting CEO, John Santaferraro, identified ThoughtSpot as a leader in AI-driven Business Intelligence. At that point, the BI company was already using AI to provide users with assisted insight, automated insight, next best insight and question, natural language explanation, natural language queries, voice recognition, visualization recommendations, image recognition, and autofill search

recommendations. For buyer reference, Ferraro Consulting recommends using long-term leaders, like ThoughtSpot, as a benchmark for other vendors under consideration.

ThoughtSpot Innovation for Generative AI Leadership

In addition to AI-enablement, ThoughtSpot runs on a business semantic layer powering their own natural language search. From the start, users were able to use natural language business terms to find the insight they needed, converse with the data, and get recommendations for the next best insight. By the time Chat GPT showed up on the scene in late 2022, ThoughtSpot had also already delivered human in the loop feedback to improve their search capabilities. In addition, they provided a business glossary where users could enter, define, and associate industry specific or company specific terms to further improve the search for intelligence.

When the company launched ThoughtSpot Sage in March of 2023, they were able to enhance basic generative AI with their fully functional natural language search, human in the loop feedback, and business semantic layer. As a result of the advanced state of their product, the company's use of generative AI was already far more intelligent than their competitors. In addition, the generative AI engine never touched the data, ensuring accuracy of the results.

ThoughtSpot Architecture for Generative AI Leadership

ThoughtSpot Sage augments the product's relational search technology with the power of LLMs to provide an even more intelligent natural language front end on top of a computational engine that can handle extreme complexity at scale. Where Chat GPT is weak in terms of accurate business representation, ThoughtSpot provides specific business context through search tokens and its semantic layer. Where Chat GPT is weak in its ability to handle complex analytics, ThoughtSpot provides a powerful compute engine. Where generative AI struggles with accuracy, ThoughtSpot ensures accuracy by learning from the human-in-the-loop training experience. As a result, their customers are already reporting up to 95% accuracy.

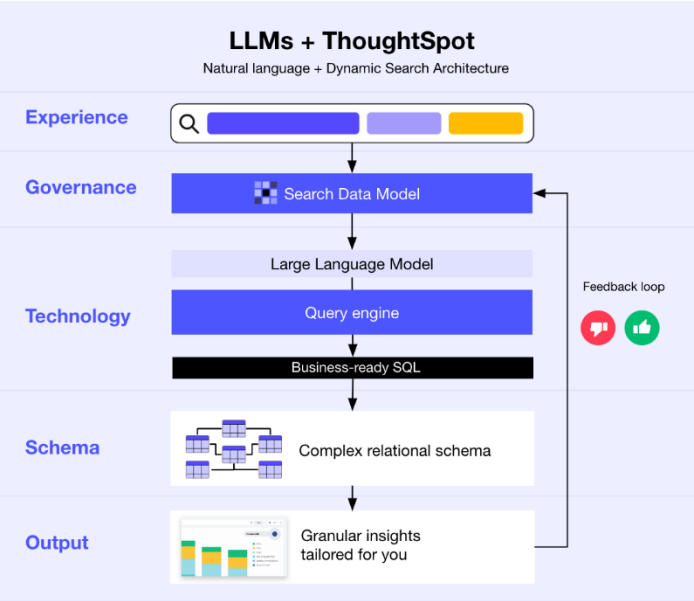

ThoughtSpot Advanced Human in the Loop Feedback for AI Leadership

To further enhance accuracy, ThoughtSpot Sage offers advanced capabilities in their human in the loop feedback for data professionals. Sage features five stages of feedback. First, there is positive-negative, thumbs up, thumbs down feedback which is common for generative AI offerings. Second, the user can make corrections to the business specific language used by the model, redirecting the query and search tokens as needed. These corrections are stored in the cache for future reference. Third, Sage breaks the question into phrases and allows for detailed association of business language with the data language in the model for a specific organization. Fourth, administrators and SMEs can see all corrections in a single control plane to make decisions on which corrections to prioritize in the fine tuning. Fifth, model stewards can update the business glossary to reflect the best possible terms, definitions, and associations in the model. This level of transparency and depth of feedback is unique among business intelligence service providers.

In conclusion, ThoughtSpot exceeds the Ferraro Consulting requirements for leadership in the use of general AI and generative AI in business intelligence and analytics.