

Viewpoint
David Meninger
 SVP &
 Research Director

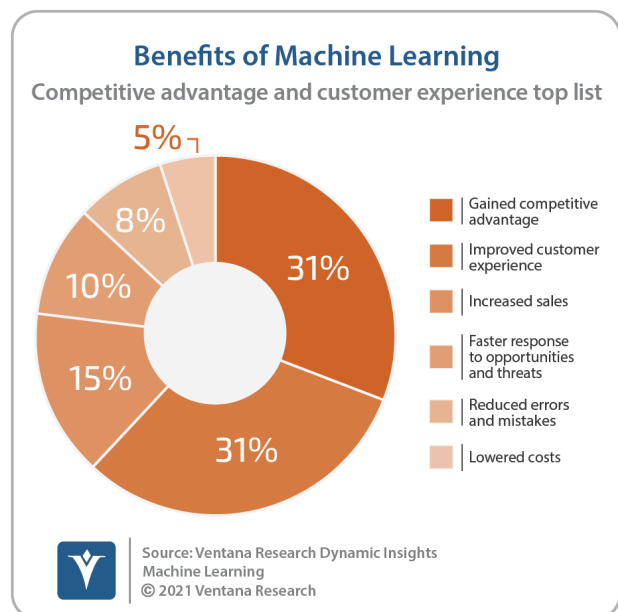
Can Your Data Keep Up with Your AI and ML Efforts?

As data volumes and data velocities increase, more and more decisions involving that data are being made using artificial intelligence and machine learning (AI/ML) technologies. However, operationalizing AI/ML models requires a dynamic, modern data architecture. Relying on traditional data management architectures that employ batch processing won't be sufficient to deal with the demands AI/ML place on an organization's data architecture.

AI/ML have become useful tools in the effort to improve an organization's operations. Our research shows that the two most common benefits organizations report from their investments in AI/ML is that they gained a competitive advantage and improved the experience of their customers. They also report increasing sales and being able to respond faster to opportunities in the market. The algorithms used in these techniques are sophisticated enough to automatically sift through and analyze millions or billions of records to find patterns and anomalies. Performing these operations manually would be impossible to complete in a timely or cost-effective manner.

But AI/ML processing can place significant demands on an organization's information technology infrastructure. First of all, large volumes of historical data are necessary to detect patterns accurately. Predicting fraud, for example, requires enough observations of fraudulent transactions to be able to identify the signs that indicate fraud. Since fraud occurs in a small percentage of all transactions, a far larger collection of all transactions is needed. Another common requirement is the ability to store and process not only structured data, but also semi-structured and unstructured data. Being able to deal with these data sources natively makes it easier to create and publish models based on data from these sources.

AI/ML processes require a bridge between analytics and operational systems. To continue the fraud example, the large historical database of transactions might be part of a data warehouse or data lake. Once the models are developed using this historical data, they should be deployed into operational systems to make it possible to detect and prevent fraud as it is happening rather than trying to catch up with the perpetrators after the fact. This process of applying the models is referred to as "scoring





records,” and the operational implications don’t end there. The models, once deployed, need to be monitored to determine if they are still accurate and replaced when a new, more accurate model is available. This process implies a constant connection between the analytical systems and the operational systems to ensure models are up to date as the records are scored.

To support these requirements, an organization’s data architecture should have three key attributes: scalability, elasticity and real-time processing capabilities. Scalability is necessary to be able to accommodate the volumes of data needed to develop AI/ML models. Elasticity—the ability to expand and contract compute resources—is necessary to handle the variations in processing needed during different phases of the AI/ML lifecycle. For instance, model-development phases require high amounts of resources when compared with model-deployment phases. And real-time processing, based on low-latency transactions is needed to score records as they occur, and then act on those scores in real time; otherwise the appropriate moment for taking action has passed.

Many organizations rely on cloud-based deployments to achieve both scalability and elasticity. In these deployments, additional computing resources can be added or removed relatively easily. While public cloud deployment is not a strict requirement, a cloud-style architecture certainly is. Teams developing AI/ML models need to be able to quickly create new instances of their databases for experimentation or they may need a burst of extra resources to try new modeling alternatives. If they must wait days, weeks or months to provision these resources, the organization will not be agile enough to take advantage of opportunities in the market as they occur.

Organizations must explore the value of AI/ML as these capabilities are practically a requirement to remain competitive in today’s markets. To deploy these capabilities, consider modern data architectures that unlock the value of AI/ML.



Dave Menninger - SVP and Research Director, Ventana Research

David Menninger is responsible for the overall direction of research on data and analytics technologies at Ventana Research. He covers major areas including artificial learning and machine learning, big data, business intelligence, collaboration, data science and information management along with the additional specific research categories including blockchain, data governance, data lakes, data preparation, embedded analytics, natural language processing (NLP) and IoT.