Magic Quadrant for Data Science and Machine Learning Platforms

28 May 2025 - ID G00822217 - 41 min read

By Afraz Jaffri, Maryam Hassanlou, and 3 more

Data science and machine learning platforms provide software to build, customize and deploy AI models using various techniques. Data science, machine learning and AI leaders need to select platforms with awareness of AI agents enabling autonomous and interactive workflows and applications.

Strategic Planning Assumptions

The Magic Quadrant for Data Science and Machine Learning Platforms (G00822217) has been republished to correct the name of AWS Bedrock Guardrails to Amazon Bedrock Guardrails, and the Alteryx AI Platform for Enterprise Analytics to Alteryx One. In addition, Altair's listing has been updated to note that Gartner's analysis factors in Altair's offering following its March 2025 acquisition by Siemens. For more information, see the **Corrections** *page on gartner.com.*

By 2027, 50% of data analysts will be retrained as data scientists, and data scientists will shift to AI engineers.

By 2027, organizations will implement small, task-specific AI models, with usage volume at least three times more than those of general-purpose large language models.

Market Definition/Description

Gartner defines a data science and machine learning platform as an integrated set of code-based libraries and low-code tooling. These platforms support the independent use and collaboration among data scientists and their business and IT counterparts, with automation and AI assistance through all stages of the data science life cycle, including business understanding, data access and preparation, model creation and sharing of insights. They also support engineering workflows, including the creation of data, feature, deployment and testing pipelines. The platforms are provided via desktop client or browser with supporting compute instances or as a fully managed cloud offering.

Data science and machine learning (DSML) platforms are designed to enable a broad range of users to develop and apply a comprehensive suite of predictive and prescriptive analytical techniques. Leveraging data from distributed sources, code-first and low-code interfaces, and native machine learning and generative AI (GenAI) capabilities, these platforms enhance and automate decision making across an enterprise. They offer a range of proprietary and open-source tools to enable data scientists and domain experts to identify patterns in data that can be used to forecast financial metrics, understand customer behavior, predict supply and demand, and many other use cases. Models can be built on all types of data, including tabular data, images, video and text, for applications that require computer vision or natural language processing.

The supported machine learning techniques range from classic regression and decision trees to more complex deep learning, and reinforcement learning and GenAI. The models built using these techniques can be used for tasks within business processes, such as credit scoring, churn prediction, predictive maintenance, recommendation and image classification.

Organizations that use DSML platforms can significantly reduce the cycle time and barriers to entry for creating predictive and prescriptive models, generating insights, and distributing results. These platforms enable collaboration and reuse of assets across multiple teams and departments, as well as the orchestration of workloads to handle high volumes of data. They also provide consistent and reproducible training and development environments that establish lineage between data, code, and model assets, enhancing data scientists' productivity. Additionally, low-code, natural language interfaces and AI assistance enable domain experts and business users to create predictive models with simple interactions. DSML platforms support MLOps practices such as deploying models in production, orchestrating both batch and real-time workloads, and ongoing monitoring of model metrics and compliance.

Mandatory Features

- Import or connect to tabular data from data management systems, including databases, data warehouses, and content repositories located on-premises and in the cloud.
- Preparation of data using data transformation tools and packages.
- Code-based development environment.
- Building and evaluation of models using a library of core data science and machine learning techniques, methods, algorithms and processes.
- Collaboration and project management tools to allow multiple users and teams to use the platform.
- Deployment, hosting and serving models in the platform for usage in services and applications.
- Model life cycle management to promote, demote, retrain and retire models.
- Administration and configuration management for user roles, permissions and resource allocation.

Common Features

- Platform-generated recommendations for the best way to prepare, integrate and model data, as well as automated creation of machine learning models based on manually selected target prediction.
- Advanced interfaces that facilitate more complex modeling for simulation, optimization and deep-learning-based use cases.
- Custom software development kits (SDKs) that provide more control and flexibility for code-based model development and integration with services and applications.
- Utilize structured and unstructured data sources including text, images, video, audio and geospatial data.
- Low-code interface for model development and AutoML functions suitable for nonexpert data science roles, including business users and domain experts.
- Postdeployment model life cycle management to retrain, retire or adapt models based on detecting and analyzing data, feature and model drift.

- Support for MLOps-based processes and tools that enable machine learning (ML) models to be deployed at scale in different operational environments.
- Functionality for working with GenAI models, such as large language models, through tracking, selection and monitoring of prompts, models and outputs.
- Techniques and tools that increase the transparency, explainability and interpretability of models to understand how and why model outputs are generated.
- Metadata management and cataloging through structured repositories of key assets including data, code, features, models, logs and outputs.
- GPU support for training deep learning and generative AI models.
- Advanced data visualization to enable exploration and investigation of data for hypothesis testing, data validation and use case identification.
- Integrated data lineage and provenance tracking to ensure data integrity and traceability throughout the data science workflow.
- Advanced hyperparameter tuning and optimization tools to enhance model performance and efficiency.

Magic Quadrant

Figure 1: Magic Quadrant for Data Science and Machine Learning Platforms

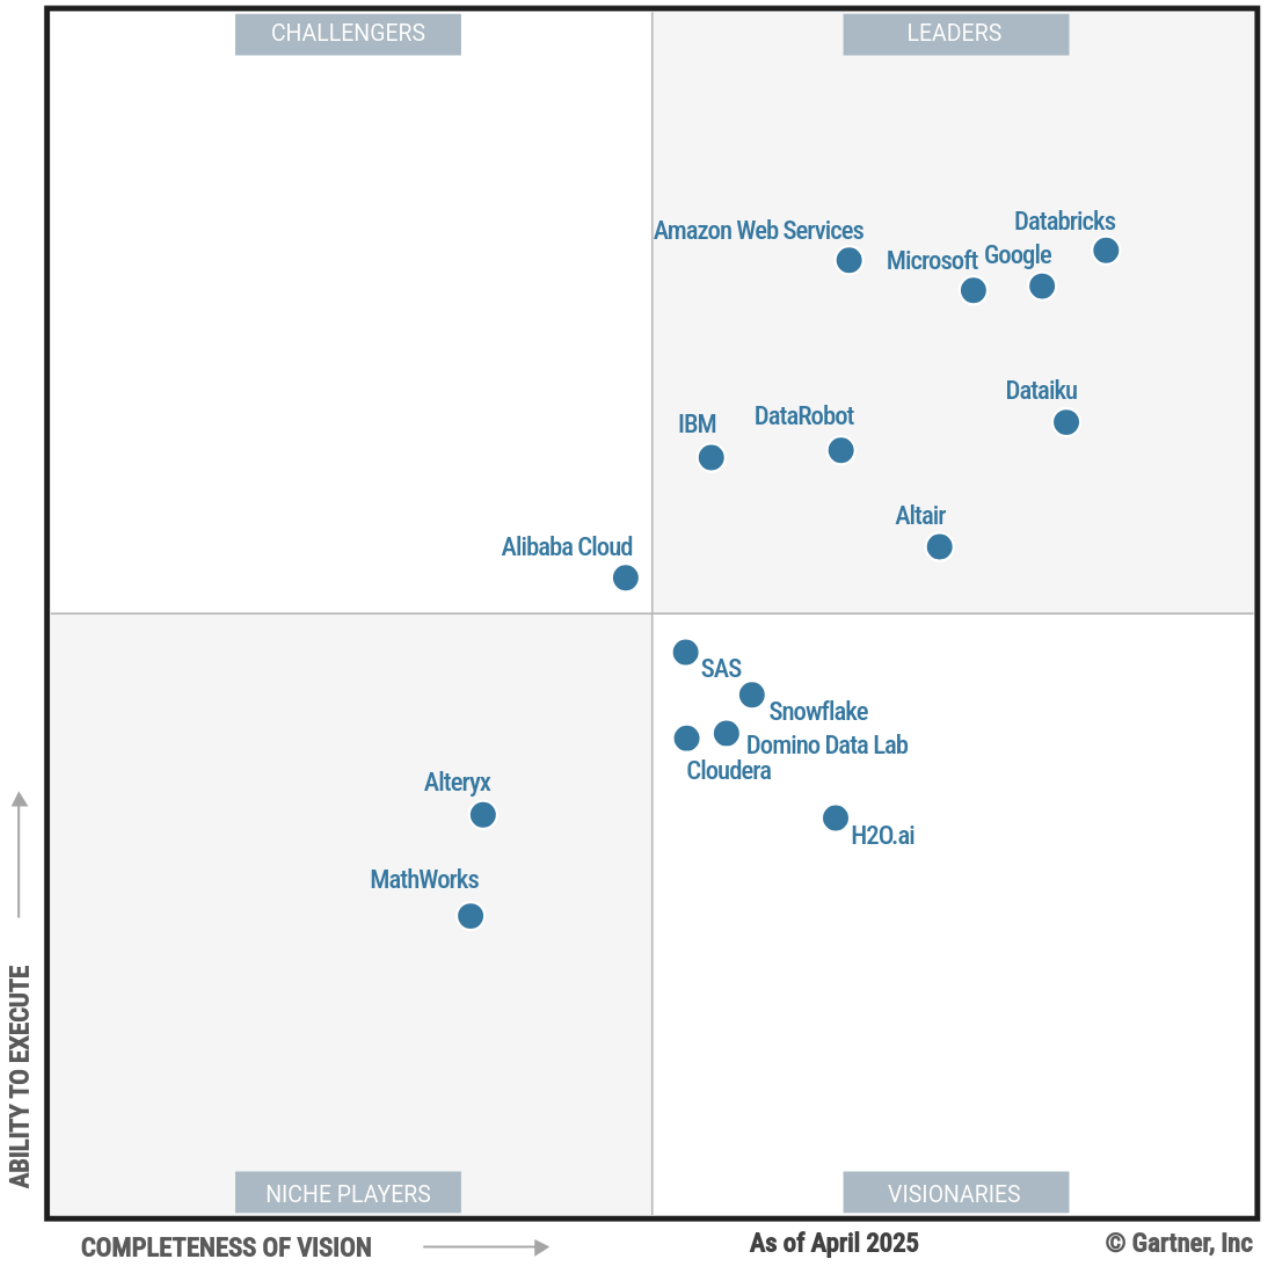

Gartner

Vendor Strengths and Cautions

Alibaba Cloud

Alibaba Cloud is a Challenger in this Magic Quadrant. The Alibaba Cloud Platform for AI (also known as PAI) provides high-performance computing resources and infrastructure that data science teams can use to build, deploy, utilize and manage ML and GenAI models. Its operations mainly focus on China, Southeast Asia, the Middle East and Australia, with a growing presence in Latin America and Europe. Its clients tend to be enterprises of various sizes operating across all sectors. In February 2025, Alibaba Cloud announced a $53 billion

investment in infrastructure over the next three years and a commitment to make artificial general intelligence the group's long-term objective.

Strengths

- **Foundation models:** Alibaba's Qwen series of models continues to be a strategic focus, with open-source and reasoning variants providing credible options outside of other major providers.
- **Security:** Alibaba has added provisions for confidential computing containers and model encryption to ensure safe hosting of models.
- **AI model inference:** Innovative techniques, such as LLM inference acceleration, are applied across all layers of the infrastructure stack to enable the development of large-scale distributed AI-driven apps.

Cautions

- **Customer adoption:** Growth in enterprises using Alibaba Cloud in Europe and the U.S. continues to be slow, with low market penetration in these regions.
- **Industry focus:** Enterprises that extensively use the platform tend to be in retail, consumer goods or technology domains, and targeted accelerators favor those industries over others.
- **Core data science:** Alibaba does not have a clear strategy for its data science capabilities, which are focused on insight generation, including data exploration, visualization and data science techniques.

Altair

Altair is a Leader in this Magic Quadrant. Its RapidMiner platform and other Altair products focus on providing enterprises with a platform for building AI-driven automation solutions in a wide range of settings. Its operations are geographically diversified, and its clients tend to be midsize to large automotive, aerospace, defense and healthcare enterprises. In the first half of 2024, Altair completed the acquisition of Cambridge Semantics, a knowledge graph platform and analytical graph database provider. The company has integrated this feature into its existing DSML platform as an "AI fabric" offering. This analysis factors in Altair's offering following its March 2025 acquisition by Siemens.

Strengths

- **Market understanding:** Altair's AI fabric vision for analytics and AI is differentiated and combines AI agents, data fabric and AI engineering at scale.
- **Acquisitions:** Altair has made targeted acquisitions in accordance with its overall vision, including Cambridge Semantics, to provide a knowledge graph platform and a regional services provider that expands the company's presence in Asia/Pacific.
- **Customer engagement:** Altair's AI Centre of Excellence provides a structured implementation program for project success without adding ongoing dependencies.

Cautions

- **Viability:** Altair's acquisition by Siemens leaves questions and uncertainty about the company's future direction and focus.
- **Customer awareness:** The AI fabric concept needs to be more effectively communicated to bring the platform into potential buyers' shortlists for DSML platforms.
- **Product portfolio:** The number of Altair offerings within and outside the analytics and AI space continues to cause integration challenges and confusion for existing and potential customers.

Alteryx

Alteryx is a Niche Player in this Magic Quadrant. In May 2025, Alteryx launched its Alteryx One Platform, a unified suite of AI-powered analytics capabilities encompassing low-code and code-centric DSML development. These products enable users to create data and analytics content, including visualizations and ML models. Its operations are geographically diversified, and its clients tend to be enterprises that range in size and operate across all sectors. Alteryx has hired a new CEO and executive team, following its acquisition in 2024 by Clearlake Capital Group and Insight Partners.

Strengths

- **Broad appeal:** Alteryx provides a unified platform for users of all skills and types, and is particularly strong at enabling data analysts and domain experts.
- **Data focus:** AI-ready data is a core driver of the platform's capabilities, across structured and unstructured data, that enable the creation of AI applications.

- **Flexibility:** Alteryx provides a composable platform that can orchestrate DSML workflows across multiple data and application platforms.

Cautions

- **Executive transitions:** Following the vendor's acquisition and new investment in 2024, its strategy and operational initiatives are recently defined and now being implemented under new leadership.
- **Persona focus:** Alteryx provides functionality for many different personas but has struggled to establish itself as a DSML platform for code-oriented users.
- **Innovation:** Alteryx adds GenAI-driven features to the platform but lags competitors for innovations specifically tailored to data scientists.

Amazon Web Services

Amazon Web Services (AWS) is a Leader in this Magic Quadrant. Both Amazon SageMaker AI and Amazon Bedrock are broadly focused on enabling all customers with the AI infrastructure, tooling, governance needed to build and deploy AI solutions (models and GenAI applications) responsibly. Its operations are geographically diversified, and its clients comprise organizations of all sizes and sectors. In December 2024, AWS announced the next generation of Amazon SageMaker, which brings together data storage and processing, analytics, DSML model building and GenAI development, all accessible through the new Amazon SageMaker Unified Studio. The company also announced new Amazon Nova models that can process text, images and video as prompts.

Strengths

- **Collaboration:** SageMaker Unified Studio combines data, analytics and AI as a fully integrated offering that simplifies collaboration and enables visibility across multiple projects.
- **AI app ecosystem:** SageMaker Partner AI apps allow third-party providers to offer their products within SageMaker Unified Studio for specific needs and use cases without the need to provide or operate infrastructure.
- **Responsible AI:** Amazon Bedrock Guardrails provides advanced features such as Automated Reasoning checks. AWS has also achieved ISO accreditation for Bedrock.

Cautions

- **Integration and flexibility:** As a new offering, SageMaker Unified Studio will need to address issues related to integration with other cloud platforms.
- **Foundation models:** AWS offers a choice of increasingly capable foundation models, but they have not yet become a more prominent consideration for customers when selecting a platform.
- **FinOps:** Monitoring spending and forecasting usage of AI infrastructure and services with AWS platforms is a challenge that can complicate cost forecasting.

Cloudera

Cloudera is a Visionary in this Magic Quadrant. Cloudera provides a unified open data lakehouse and fabric for data management, together with Cloudera AI, for collaborative AI systems and application development by data engineers, data scientists and AI developers. Its operations are geographically diversified, and its customers tend to be large enterprises across all sectors. Cloudera made two significant acquisitions in 2024: Verta's operational AI Platform, acquired in June, adding model management capabilities; and Octopai, a data lineage and catalog platform acquired in November, adding to the vendor's existing metadata management capabilities.

Strengths

- **Strategic acquisitions:** Cloudera's Verta acquisition adds key capabilities for operationalizing AI models and workloads. Octopai adds innovation to metadata management to keep ahead of market trends.
- **Private AI:** Cloudera enables the building and deploying of AI applications within enterprise boundaries, either on-premises or virtual private cloud, including hardware, with its AI-in-a-box offering.
- **Model availability:** Cloudera offers the option to choose between open and commercial models available on the day of release via Hugging Face and Nvidia NGC model registry.

Cautions

- **Awareness:** Cloudera AI is a newly branded platform and suffers from a lack of awareness of its innovative capabilities among potential buyers.

- **Complexity:** Cloudera's open-source strategy integrates several libraries and frameworks into the platform, which may increase the complexity of maintenance and upgrades.
- **AI governance:** Process workflows and risk controls are not supported as well as those from market leaders in this area. The Verta acquisition could bring enhancements in these areas.

Databricks

Databricks is a Leader in this Magic Quadrant. Its Databricks Data Intelligence Platform offers integrated lakehouse, data engineering, governance, analytics and model training, and serving capabilities, including GenAI models. Its operations are geographically diversified, and its clients tend to be midsize and large enterprises across all sectors. In 2024, Databricks announced a Series J round of funding for $10 billion of equity. It also made several acquisitions, including Tabular to improve data compatibility between platforms, Lilac for unstructured data evaluation for GenAI use cases, and Prodvana to deliver cloud-native infrastructure improvements. In February 2025, Databricks announced a partnership with SAP to integrate the Data Intelligence platform into SAP Business Data Cloud as a native offering.

Strengths

- **Vision:** Databricks encompasses data engineering, ML and GenAI in a comprehensive, agent-based ecosystem with assistive analytics and AI development features, suited for use cases in areas of financial and reputational risk.
- **Leadership:** The leadership team has remained essentially unchanged since its inception and makes decisive moves regarding acquisitions and the best way to utilize cutting-edge innovation.
- **Practitioner popularity:** Databricks has reached a desirable platform status among data scientists and ML engineers, increasing the skills available in the market.

Cautions

- **Learning curve:** Success in using the platform depends on the skills of the users, many of whom will need to build up a significant level of experience before becoming proficient.
- **Competition:** Databricks previously held the position of being the only integrated lakehouse and DSML platform, but offerings from other vendors are now available.

- **Infrastructure:** Some features available natively with cloud providers may not be available with Databricks, such as GPU serverless compute for ML and the use of custom processors.

Dataiku

Dataiku is a Leader in this Magic Quadrant. Its Dataiku product combines low-code development and AI-assisted tools with code-first interfaces, including notebooks and code editors, to enable collaboration across the enterprise and deliver AI applications. Its operations are geographically diversified, and its customers tend to be midsize and large enterprises across all sectors. In October 2024, Dataiku expanded its LLM Mesh feature with the LLM Guard Services suite to advance enterprise GenAI deployments from proof of concept to production, enabling quality assurance monitoring across cost, quality and safety factors.

Strengths

- **Core data science:** Dataiku has not lost sight of insight-generation-focused data science while delivering GenAI features and releasing its Stories feature to augment decision making.
- **Customer support:** Dataiku consistently receives high praise from customers who use its services team for setup and initial use cases.
- **Market understanding:** Dataiku's LLM Mesh targets the pain points associated with GenAI development, specifically governance and consistent use of models.

Cautions

- **Feature delivery:** Expansion of capabilities risks being unable to deliver best-of-class features for all parts of data science and AI delivery.
- **Buyer awareness:** Enterprises often utilize only parts of the platform to meet specific needs that cannot justify further investment.
- **Differentiation:** Dataiku operates on top of other vendors included in this Magic Quadrant, many of which provide overlapping capabilities within their offerings.

DataRobot

DataRobot is a Leader in this Magic Quadrant. Its Enterprise AI Suite provides a platform for developing predictive and GenAI-based solutions using models that can be trained with no-code interfaces for business users or coding environments for expert data scientists. Its operations are geographically diversified, with the exception of China, and its clients include enterprises that range in size and sectors. DataRobot refreshed its brand and positioning at the end of 2024, focusing on providing a platform that can deliver, manage and monitor AI apps through the use of AI models and agents across multiple horizontal and vertical domains.

Strengths

- **Vision:** DataRobot's strategic pivot to provide an AI-driven business app ecosystem with orchestration between agents resonates with the needs of enterprises looking to build agentic systems.
- **Acquisitions:** DataRobot's acquisition of Agnostiq, with its open-source distributed computing platform, Covalent, aligns with the vendor's vision for providing an agentic AI platform.
- **Market understanding:** DataRobot conducted the Unmet AI Needs Survey to fully understand the challenges associated with being the AI leader in an enterprise.

Cautions

- **Community:** The perception of DataRobot as being poorly suited for expert data science needs will persist after rebranding, reducing the ability to build a practitioner community.
- **User experience:** Frequent updates to the platform have made the low-code and pro-code interaction and interface difficult to navigate and manage.
- **Partnerships:** Building a true business ecosystem will depend on buy-in from application vendors to provide integration and support.

Domino Data Lab

Domino Data Lab is a Visionary in this Magic Quadrant. Its Domino enterprise AI platform emphasizes openness and adaptability, enabling data scientists to leverage on-premises and cloud infrastructures, alongside various development environments and programming languages, to develop, deploy and govern AI and GenAI models. Its operations are geographically diversified, and its clients tend to be global enterprises in the life sciences,

financial services, public sector and regulated industries. In October 2024, Domino added its Governance solution, which allows customers to orchestrate and automate collecting, reviewing and tracking all materials to enforce compliance with any internal or external policy.

Strengths

- **AI governance:** Domino has evolved capabilities through its governance solution to support risk, compliance and governance requirements for regulated and other industries, ahead of competing vendors.
- **Industry focus:** Domino provides targeted solutions for specific industries, accelerating business value and adoption.
- **FinOps:** Domino provides visibility on infrastructure utilization and enables cost optimization across distributed cloud and hybrid environments.

Cautions

- **Awareness:** Domino's platform does not appear on as many shortlists as other vendors outside of their target verticals, and examples of successful implementations are often concentrated in specialized enterprise deployments.
- **Multipersona support:** Domino focuses on data scientists, ML engineers and AI engineers. Specific features for data engineers and line-of-business users are still evolving.
- **Integration:** Deploying and integrating the product into enterprise environments can be complex and may require further support.

Google

Google is a Leader in this Magic Quadrant. Its Vertex AI platform provides enterprise teams with data and AI governance, data engineering and analysis, data science and MLOps, and GenAI application and AI agent development tools using Google Cloud infrastructure. Its operations are geographically diversified, with the exception of China, and its customers are enterprises of all sizes across a wide variety of industries. Google is invested in advancing GenAI research and development with its Gemini family of foundation models produced by Google DeepMind. In March 2025, Google unveiled its latest GenAI model, Gemini 2.5 Pro, adding to the series of 2.0 models, including Pro, Flash, Flash-Lite and Flash Thinking.

Strengths

- **Unified governance:** Integrated with Vertex AI, Google Dataplex captures, stores and exposes all artifacts relating to datasets, models and features.
- **Retrieval augmented generation (RAG):** Vertex AI Search enables the creation of GenAI apps and agents across structured and unstructured data as well as grounding on Google Search, enterprise data, and data from trusted third-party providers.
- **Co-innovation:** Google provides an AI engineering team to apply cutting-edge research to industry challenges with a number of clients.

Cautions

- **Community:** The Vertex AI support offered by the community outside of Google is not as significant as other providers, which can lead to longer issue resolution times.
- **Complexity:** Vertex AI's core features can be difficult to use, depending on the level of familiarity with the platform and may require in-built abstraction layers to ease data scientist and developer experience.
- **Coherence:** Google Cloud offers various RAG-based solutions, including Vertex AI Search, Vertex AI RAG Engine and Vector Search, which could cause confusion for buyers.

H2O.ai

H2O.ai is a Visionary in this Magic Quadrant. Its Enterprise AI platform — featuring h2oGPTe for generative AI and H2O Driverless AI for predictive AI — enables the creation of domain-specific agents to automate workflows and deliver sovereign AI via on-premises and air-gapped deployments. H2O.ai serves a geographically diverse client base of large enterprises across financial services, telecommunications and the public sector. In March 2025, H2O.ai released the latest version of its h2oGPTe Agent, which achieved the top position on the General AI Assistants (GAIA) benchmark leaderboard, with an average score of 75%.

Strengths

- **GenAI techniques:** H2O.ai has specific expertise in creating tools for enterprise GenAI usage, including training small language models (SLMs), fine-tuning, distillation, vector search and AI agent creation.
- **Composite AI:** H2O.ai provides a platform to integrate predictive AI models into agentic processes, not only relying on LLMs.

- **Talent:** H2O.ai has 24 Kaggle Grandmasters who can help shape and deliver innovative customer solutions.

Cautions

- **Pricing:** AI Cloud pricing is based on "AI units" that can be difficult to understand and forecast accordingly.
- **Market awareness:** H2O.ai's capabilities as an AI platform and an innovator are not well-known by prospective customers.
- **Partners:** H2O.ai's partner network is expanding by adding new partnerships with Dell, EY and VAST Data, but it includes fewer consulting partners to provide deployment and development services than other providers.

IBM

IBM is a Leader in this Magic Quadrant. Its watsonx supports the development of business AI applications by integrating data management, governance, and predictive and generative AI tooling for use by data scientists and AI engineers. Its operations are geographically diversified, and its customers comprise enterprises of various sizes and sectors. IBM announced several innovations over the past 12 months for creating multiagent systems in its platform, including a new agent orchestration platform, watsonx Orchestrate, partnership with CrewAI to build and streamline agentic workflows, and the open-source BeeAI Framework. In February 2025, IBM announced its intention to acquire DataStax to add to its GenAI and data management capabilities.

Strengths

- **Flexibility:** IBM provides a wide range of tools, open-source and proprietary frameworks, foundation models and GPU options for all enterprise needs.
- **AI-ready data:** IBM's acquisition of DataStax will significantly add to its existing capabilities to deliver RAG-based enterprise solutions.
- **Innovation:** Research from IBM labs is being incorporated into products to provide differentiated capabilities, such as its Granite family of models, AutoRAG, and InstructLab offerings.

Cautions

- **Awareness:** The performance and capabilities of using IBM's foundation models and other AI innovations are not as well-known as other providers that also provide consumer-focused products.
- **Core data science:** SPSS is not highlighted as a component within watsonx for core data science, limiting its potential to be utilized for these use cases.
- **Multicloud:** watsonx does not have the same level of integration with GCP as AWS and Azure.

MathWorks

MathWorks is a Niche Player in this Magic Quadrant. Its MATLAB and Simulink products help data scientists, engineers, other scientists and business users apply AI to design, develop and maintain engineered systems. Its operations are geographically diversified, and its clients tend to be enterprises of all sizes in industrial sectors such as automotive, aerospace, industrial automation, telecommunications and medical devices. MathWorks has developed its MATLAB Copilot to bring intelligent assistance into the platform interface. The system helps with code editing, code explanation, debugging and offers guidance and answers sourced from documentation.

Strengths

- **AI and simulation:** Simulation, which is often neglected as a core data science capability, is central to MathWorks' MATLAB and Simulink offerings.
- **Talent availability:** Many educational institutions incorporate MATLAB into their curricula, resulting in a workforce skilled in the product.
- **Agentic AI:** MathWorks is one of the few vendors that understands non-LLM-based AI agent use cases and has supported their development for years with, for example, its Reinforcement Learning Toolbox.

Cautions

- **Data scientist focus:** MATLAB would not be the primary platform of choice for data science teams outside of engineering, science and industrial applications, given the platform's specialist focus on engineering and industrial use cases that may or may not use AI.

- **GenAI models:** The platform does not provide features focused on training, fine-tuning or RAG use cases, focusing more on using GenAI to accelerate development.
- **Private cloud:** A private cloud offering that includes GenAI functionality is not yet available.

Microsoft

Microsoft is a Leader in this Magic Quadrant. Its Azure Machine Learning (Azure ML) product provides data scientists, AI engineers and pro-code developers with tools for model development and training, AutoML and MLOps. It also includes integrated governance and security as well as compute resources and custom AI infrastructure. Its operations are geographically diversified, and its customers comprise enterprises of various sizes and sectors. In November 2024, Microsoft announced Azure AI Foundry as a new offering, combining several existing and new services for developing and deploying AI- and agent-based solutions. Azure ML is one offering included in Azure AI Foundry.

Strengths

- **Innovation:** Azure AI Foundry Labs allows experimentation with leading AI frameworks in a safe environment.
- **Ecosystem:** Microsoft has partnerships with many third-party platforms and providers to enable best-of-breed and single-unified-platform usage.
- **Pricing:** The vendor provides flexible pricing options for GenAI usage and broad geographic coverage, enabling organizations of all sizes to build AI solutions.

Cautions

- **GenAI models:** Azure AI Foundry hosts models from many providers, but as the performance gap between OpenAI models and other providers diminishes, buyers can try alternative platforms without sacrificing performance.
- **Copilot:** Microsoft Copilot for data science tasks is initially available only within notebooks to help with exploratory data analysis and model building.
- **Product strategy:** Microsoft has had several iterations of AI offerings with different branding (e.g., current branding is Azure AI Foundry), causing confusion for buyers and users of AzureML.

SAS

SAS is a Visionary in this Magic Quadrant. Its SAS Viya product focuses on the entire data and AI life cycle. It has low-code and pro-code interfaces for managing data and ETL processes, developing models using multiple programming languages, and managing and governing models built using a wide range of DSML and GenAI techniques. Its operations are geographically diversified, and its customers tend to be midsize to large enterprises across various industries. In November 2024, SAS acquired Hazy synthetic data software. Integrated within SAS Viya, the software will generate high-quality synthetic data, accelerating the development cycle for AI projects to aid compliance with responsible AI and security practices, and deliver cost savings.

Strengths

- **Composable architecture:** SAS is enabling Viya platform components to be used as independent platforms, providing end users with flexibility and choice.
- **Model risk management:** SAS provides differentiated solutions for enterprises with low-risk thresholds, including regulated industries, that can manage multiple types of model risk.
- **Microsoft partnership:** Viya on Azure offers a compelling proposition for customers who want a complete package of integrated capabilities and compute resources.

Cautions

- **Vision:** AI agent capabilities are included in the platform, but the vendor's future vision for augmentation and automation lags its peers.
- **Viya adoption:** Gartner inquiries from SAS users that have not transitioned onto Viya more frequently consider alternative platforms as compared to other vendors.
- **Customer interest:** SAS has competition from smaller and newer vendors in DSML and is not as frequently assessed for GenAI capabilities as other vendors in this Magic Quadrant.

Snowflake

Snowflake is a Visionary in this Magic Quadrant. Snowflake provides a cloud-based infrastructure for DSML practitioners, data engineers and other personas. This infrastructure allows users to access, manage and govern structured and unstructured data within its data lake and data warehouse, run analytical workflows, develop agentic apps, and train and host

AI models. Its operations are geographically diversified, and its customers tend to be midsize to large enterprises across various industries. In early 2024, Snowflake promoted its SVP of AI to CEO to deliver accelerated growth in AI and ML. In November 2024, Snowflake acquired Datavolo, which developed a platform for building and managing multimodal data pipelines. In February 2025, Snowflake announced the availability of OpenAI models directly in Snowflake Cortex AI.

Strengths

- **Vision:** Snowflake has a strong vision for combining structured and unstructured data insights in a unified platform via data agents.
- **Marketplace:** Snowflake provides an AI marketplace for AI apps and products that simplifies integration and caters to a wide variety of use cases.
- **Training:** Snowflake has invested heavily in training and upskilling for various personas and roles through collaboration with online learning platforms.

Cautions

- **Pricing:** Users report that Snowflake's pricing model is complex, making it difficult to forecast costs accurately.
- **Availability:** Some LLM-based services may need to be assessed for regional and cloud provider availability.
- **MLOps:** Functionality for operationalizing and monitoring all types of ML models has been slower to arrive than on other platforms.

Vendors Added and Dropped

We review and adjust our inclusion criteria for Magic Quadrants as markets change. As a result of these adjustments, the mix of vendors in any Magic Quadrant may change over time. A vendor's appearance in a Magic Quadrant one year and not the next does not necessarily indicate that we have changed our opinion of that vendor. It may be a reflection of a change in the market and, therefore, changed evaluation criteria, or of a change of focus by that vendor.

Added

Snowflake

Dropped

Anaconda: Anaconda provides data science tooling, including libraries, repositories and package managers, to enable enterprises with secure, open-source Python development. It was dropped from this Magic Quadrant due to the withdrawal of its Data Science & AI Workbench product and a change in strategy to focus on advancing open-source AI and software supply chain security.

KNIME: KNIME provides a low-code-centric platform for end-to-end DSML development, including KNIME Analytics Platform and KNIME Business Hub. The vendor did not meet the inclusion criteria for the DSML platform to be recommended for use in predominantly code-centric data science teams for this Magic Quadrant.

Posit: Posit offers Posit Workbench, Posit Connect, Posit Package Manager, and Posit Cloud products. These products provide an open platform for data scientists to build predictive and prescriptive models and deliver insights to business users and decision makers. Posit was dropped from this Magic Quadrant because it did not meet the Customer Interest Index threshold for inclusion.

Inclusion and Exclusion Criteria

To qualify for inclusion, providers need to provide:

- A go-to-market strategy for their DSML platform that focuses on the persona of a professional data scientist, machine learning engineer or AI engineer.
- The DSML platform to be recommended for use in predominantly code-centric data science teams.
- At least $100 million calendar year 2024 (estimated) platform revenue, excluding hardware and services, or
- At least $50 million calendar year 2024 (estimated) combined platform revenue, excluding hardware and services, and at least 20% year-over-year growth when compared to calendar year 2023, or

- At least $25 million calendar year 2024 (estimated) combined platform revenue, excluding hardware and services, and at least 40% year-over-year growth when compared to calendar year 2023.
- At least 10 paying customers (logos), as of 1 January 2025 for their DSML platform in each of two of the following regions:
  - North America
  - South America
  - Europe, the Middle East and Africa
  - Asia/Pacific
- Achieve a score above 55 points for this Magic Quadrant according to Gartner's Customer Interest Index (CII). Data inputs used to calculate relevancy include interest from Gartner clients, customer reviews and key vendor growth metrics.

Honorable Mentions

Teradata: Teradata offers its VantageCloud product platform for analytics, data warehousing and lakehouse requirements. It offers in-database analytics through its ClearScape Analytics engine, enabling model development and deployment within the data warehouse to reduce data movement. Teradata supports integration with common DSML frameworks and provides capabilities for using LLMs and GPU resources for GenAI applications on enterprise data. Teradata did not meet the inclusion criteria for this Magic Quadrant.

Aible: The Aible platform provides a full GenAI orchestration and value extraction stack primarily for business personas with the ability to extract insights from data and develop AI and GenAI applications. It features "chat with Aible" functionality for advanced natural language querying, recommendations and explanations based on AI simulation, guardrails and hallucination detection, and the ability to use many different GenAI models for development. Aible did not meet the inclusion criteria for providing a go-to-market strategy that focuses on the persona of a professional data scientist for this Magic Quadrant.

Evaluation Criteria

Gartner evaluates a vendor's ability to execute in the data science and machine learning platforms market by applying criteria that assess its capability to address current market trends. This involves demonstrating a high-quality track record of maintaining visibility and brand awareness in the AI space and operational practices that can attract and sustain high levels of customer demand and engagement.

Ability to Execute

Product or Service: This criterion covers the assessment of the vendor's capabilities to deliver features and functionality within its platform for core data science and ML-engineering-based use cases.

Overall Viability: This criterion covers the assessment of a vendor's key financial growth metrics and diversity of its customer base.

Sales Execution/Pricing: This criterion covers the assessment of customers' willingness to make strategic investments in a vendor and the vendor's ability to support client interest or demand.

Market Responsiveness and Track Record: This criterion covers the assessment of the quality of a vendor's release cycles for its product, how it responds to customer needs and the ability to prioritize feature requests.

Marketing Execution: This criterion covers the assessment of a vendor's market visibility, brand awareness and community development.

Customer Experience: This criterion covers the assessment of a vendor's quality of technical support, level of customer retention and customer advocacy.

Operations: This criterion covers the assessment of a vendor's employee retention rates and commitment to corporate social responsibility (CSR) initiatives, including data for good.

Ability to Execute Evaluation Criteria

Evaluation Criteria	*Weighting*
Product or Service	High

Evaluation Criteria	*Weighting*
Overall Viability	High
Sales Execution/Pricing	Low
Market Responsiveness/Record	Medium
Marketing Execution	Medium
Customer Experience	Medium
Operations	Low

Source: Gartner (May 2025)

Gartner evaluates a vendor's Completeness of Vision in the DSML platforms market by applying criteria that assess its ability to understand current market trends across diverse user groups, including data scientists, AI engineers, software developers and business use This involves recognizing how emerging technologies can advance the goals of data-drive decision making and the development and deployment of AI systems.

Completeness of Vision

Market Understanding: This criterion covers the assessment of a vendor's vision and strategic positioning with respect to GenAI and the DSML market as a whole, the competitive differentiators of the vendor and the awareness of customer needs, buying personas and end-user roles.

Marketing Strategy: This criterion covers the assessment of a vendor's go-to-market approach, considering targeted buyer and user personas and market visibility in the context of dedicated resources.

Sales Strategy: This criterion covers the assessment of a vendor's ability to produce different sales strategies for target customer personas, strategic partnerships and effective use of sales resources.

Offering (Product) Strategy: This criterion covers the assessment of a vendor's depth and breadth of product portfolio and its open-source strategy and commitment.

Business Model: This criterion covers the assessment of a vendor's evolution of its business model in line with enterprise needs, the significance of DSML to its overall business and the viability of the vendor in the context of current market dynamics.

Vertical/Industry Strategy: This criterion covers the assessment of a vendor's clear and consistent vertical focus and understanding of industry needs.

Innovation: This criterion covers the assessment of a vendor's GenAI feature quality and roadmap, and commitment to ongoing research that results in significant product enhancements.

Geographic Strategy: This criterion covers the assessment of a vendor's clear and consistent regional focus and understanding of regional market dynamics.

Completeness of Vision Evaluation Criteria

Evaluation Criteria	*Weighting*
Market Understanding	High
Marketing Strategy	Medium
Sales Strategy	Medium
Offering (Product) Strategy	High
Business Model	Medium
Vertical/Industry Strategy	Low
Innovation	High
Geographic Strategy	Low

Evaluation Criteria	*Weighting*

Source: Gartner (May 2025)

Quadrant Descriptions

Leaders

Leaders in this market have a mature, refined and targeted company and platform strategy that incorporates and leverages GenAI and AI agents to drive their customers' business value. They see opportunities for leveraging agents that other providers may not see or hav made significant investments above and beyond standard offerings. They have the capabili to innovate at a speed that outperforms other vendors. In addition, they can clearly articulate how they provide value to the multiple types of personas involved in the process of building data science and machine learning models.

Challengers

Challengers in this market have the operational capacity to serve a wide variety of enterpri needs in the data science and machine learning space through brand recognition and complementary product offerings. Their current limitations are centered on the appeal of t platform among their target users and limited geographic and/or industry focus. They shov potential for adding innovative and differentiating features from their product roadmap tha could gain traction in specific industries or use cases that are in high demand.

Visionaries

Visionaries understand the DSML market and its future direction, and offer a differentiated view of solutions that need to be provided to meet enterprise AI needs. They offer industry-specific functionality and clearly demonstrate value to their customers on an individual and enterprise level. They are limited by not having the necessary recognition of their product for complete end-to-end DSML capabilities due to historical brand association or limited marketing initiatives and community influence.

Niche Players

Niche Players are characterized by their specialized focus on specific industries or a limited set of user groups, such as low-code or code-first practitioners. They are known to deliver solutions that meet the needs of their target demographics but fail to demonstrate a broader

appreciation of market trends and enterprise needs, especially in regard to GenAI. Their appeal is limited beyond a core audience, and they struggle to grow in line with the average market rate. They struggle with user adoption of the platform beyond a basic set of use cases, such as data preparation and exploration, that do not drive as much business value as competitor offerings.

Context

The DSML platform market remains at the forefront of buyers' minds when selecting AI delivery platforms for the enterprise. The market has grown beyond enabling classical data science and AI techniques to provide a holistic AI platform for innovative generative AI solutions, including retrieval augmented generation and AI agents built by AI engineering teams. Use of GenAI in the enterprise will also require other types of models built with various techniques to be added in a composite fashion, which is a key differentiator for many of the platforms.

DSML activities within enterprises have grown outside of centralized, core DSML teams. This results in many enterprises needing to deal with multiple platforms while maintaining best practices for development, monitoring and responsible AI use. Many of the vendors in this Magic Quadrant offer partnerships and integration with each other, meaning that platform selection sometimes is not simply a straight choice between platforms. AI and analytics leaders should understand how these platforms complement each other, whether for low-code, code-first, data management or operationalization-based functionality, to provide the best means of delivering value.

Market Overview

The data science and machine learning (DSML) platform market is increasingly focused on empowering the creation of agentic systems with integrated large language models (LLMs) and GenAI-driven assistants to enhance the data science workflow. According to Gartner's 2024 Analytics and AI Engineering survey, over 50% of respondents report that AI tools for automated insights and natural language queries are being used in AI development.

The strategic value of DSML platforms within enterprises continues to rise, fueled by the growing demand for advanced AI solutions. GenAI is a significant catalyst, but the challenge

of integrating data, models, code and infrastructure into reliable, scalable products remains. In addition, composite AI systems that combine predictive and generative models will become a standard methodology for AI development, further cementing the importance of a DSML platform.

DSML platforms also need to be capable of handling data and become operational across multiple environments, including multicloud, hybrid and on-premises infrastructure. As data engineering, pipeline construction, model training, experiment tracking, deployment and user-facing applications mature, data scientists need to evolve into AI engineers.

The movement toward making data science accessible to a broader range of business roles continues, with platforms increasingly supporting individuals aligned with business units and functional areas. AutoML remains a foundational feature, now supplemented by GenAI-driven assistants that provide coding support, natural language interaction and workflow creation. A significant challenge for data science and AI leaders is overseeing and governing the activities of dispersed DSML teams while optimizing collaboration with centralized resources. Enhanced AI governance and management capabilities, linked across data sources and other assets, are now must-have capabilities.

Major cloud providers continue to expand their influence due to their comprehensive compute, data and infrastructure offerings needed for DSML development. However, other vendors still have ample opportunity to succeed, particularly by fostering collaboration between teams — a critical factor for DSML and GenAI advancement. Expanding DSML techniques throughout enterprises and various business areas presents a significant opportunity for vendors and users alike.

Acronym Key and Glossary Terms

AutoML	Automated machine learning
DSML	Data science and machine learning
GenAI	Generative AI
LLM	Large language model

ML	Machine learning
MLOps	Machine learning operations
RAG	Retrieval-augmented generation

⊕ Evidence

⊕ Evaluation Criteria Definitions