

Magic Quadrant for Cloud AI Developer Services

Published: 24 February 2020 **ID:** G00431332

Analyst(s): Van Baker, Bern Elliot, Svetlana Sicular, Anthony Mullen, Erick Brethenoux

The urgent need to deploy artificial intelligence puts businesses at risk of making bad choices. This report will help application leaders shortlist vendors of cloud AI services — in the key areas of language, vision and automated machine learning — that developers can use to enhance applications.

Strategic Planning Assumptions

By 2023, 40% of development teams will be using automated machine learning services to build models that add AI capabilities to their applications, up from less than 2% in 2019.

By 2025, 50% of data scientist activities will be automated by AI, easing the acute talent shortage.

Market Definition/Description

This document was revised on 2 March 2020. The document you are viewing is the corrected version. For more information, see the [Corrections](#) page on gartner.com.

This Magic Quadrant evaluates the important emerging market for cloud-hosted artificial intelligence (AI) services for application developers. It assesses vendors with offerings in three key service areas: language, vision and automated machine learning (AutoML).

Gartner defines cloud AI developer services as follows:

Cloud AI developer services are cloud-hosted services/models that allow development teams to leverage AI models via APIs without requiring deep data science expertise. These hosted models deliver services with capabilities in language, vision and automated machine learning. These services are often available

via API access and are typically priced based on the number of API calls. In some cases, services are usable via integrated configuration tools. Examples of these services include natural language understanding, sentiment analysis, image recognition and machine learning model creation.

Services in this market enable application leaders to have enhanced capabilities integrated into applications developed for their enterprise. These capabilities include, for example, those of chatbots, image recognition technology and machine learning (ML) models (see further “Machine Learning Alters the Role of the Developer”).

This Magic Quadrant is related to “Critical Capabilities for Cloud AI Developer Services,” which Gartner plans to publish later in 2020.

Magic Quadrant

Figure 1. Magic Quadrant for Cloud AI Developer Services



Source: Gartner (February 2020)

Vendor Strengths and Cautions

Aible

Aible, headquartered in Foster City, California, U.S., is a Visionary in this Magic Quadrant. Its unique value proposition is AutoML focused on business impact, rather than on “neutral” data science indicators like accuracy and recall. As such, Aible’s product is designed for an audience, including developers and business users, unskilled in data science. It is available in either a public cloud or a virtual private cloud on Amazon Web Services (AWS) and Microsoft (Azure) infrastructure. Aible offers a free 30-day product trial and provides excellent customer support. The vast majority of the reference customers for Aible whom we surveyed for this Magic Quadrant purchased from this vendor in order to drive innovation.

Strengths

- Aible’s focus on business impact and ROI makes the customer’s business an integral part of a model’s development and deployment. Aible automates the entire model creation and implementation process to ensure the resulting solution is deployed and used successfully. Aible explains what drives its predictions for different models and offers to compare the impact of those models. It incorporates predictions into customers’ enterprise applications, so that insights are available right where users work. It then monitors the business outcomes, compares these with what was predicted, and flags “model drift” when ROI diminishes.
- Aible’s approach to providing the best possible models for real-life business requirements and constraints is unique. Its product automatically asks business-relevant questions to understand business objectives, cost-benefit trade-offs and constraints in relation to, for example, domain, infrastructure, cost and governance realities. Aible lets customers compare the impact of a standard ML model with that of a model working under actual constraints. It also assesses the impact of true positives relative to false positives, according to the given limitations.
- Aible’s product addresses issues that could hinder AI adoption, such as difficulty of use. It has strong automation capabilities, such as for feature engineering, model building and model management. Its blueprints get customers off to a quick start by recommending variables that are typically good predictors for specific use cases. For security purposes, Aible extracts only metadata and, by design, does not see customer data or the trained model.

Cautions

- Development of Aible’s current product started in 2018, and some customers have yet to achieve full deployment. Although surveyed reference customers for Aible would recommend it to others, more than a half would do so with reservations, due to their growth expectations for this vendor. That said, Aible justifies these expectations by quickly adding new capabilities, and our survey indicates that most of its customers match their use cases to the kinds of AI problems that Aible can solve.
- Aible does not offer any of the language and vision services considered by this Magic Quadrant. It does, however, plan to deliver NLU, sentiment analysis (emotion AI), text analytics and image

recognition services by the end of 2020. In the meantime, many such services can be obtained via Aible blueprints and seamlessly supported third-party open-source frameworks such as MLflow, TensorFlow and scikit-learn. Chatbot virtual assistant frameworks and dialogue management, although not offered as stand-alone capabilities, are part of Aible's data flows.

- Some technical and infrastructure integrations require workarounds in Aible's platform. For ease of integration using standard APIs and tools, Aible received the lowest score of any vendor in this Magic Quadrant. A reduced focus on APIs and tools reflects Aible's vision to avoid the need for explicit integrations, which is illustrated by its already-available seamless, end-to-end integrations with Salesforce and Tableau. Integrations with external systems are possible via connectors, with the help of Aible's support team. Integration with AWS is straightforward, but integration is otherwise the area in which Aible is least mature — although it continues to introduce more capabilities in this area, with support for open-source third-party capabilities such as Apache Spark connectors and MLflow. Aible is also working on partnerships with integration vendors to alleviate integration issues.

Amazon Web Services

[Amazon Web Services \(AWS\)](#), headquartered in Seattle, Washington, U.S., is a Leader in this Magic Quadrant. It offers a comprehensive set of services and APIs, including ones for the language and vision service areas, which are further supported by Amazon's SageMaker ML services for developers and data scientists. Additionally, AWS announced SageMaker AutoPilot at its re:Invent conference in 2019, and offers ML frameworks and hardware for projects requiring more advanced capabilities.

Strengths

- AWS combines a very broad and capable AI portfolio with high visibility in the business and consumer sectors. Its visibility in the business sector is due partly to its storage and computing solutions; its visibility in the consumer sector is due to Amazon's online retail business and Alexa AI product. The international breadth and depth of AWS's market reach means that it is likely to be shortlisted when enterprise users and developers are considering their AI options.
- AWS's offer for developers caters to the needs both of those without ML skills and those seeking advanced functions. Those without ML skills can use pretrained AI services with continuously learning APIs; these services include a broad range of vision, speech, natural language and text-processing functions. Developers and data scientists looking for more advanced functions can use SageMaker's full ML suite.
- To help developers learn to apply ML, AI and deep learning, AWS offers over 30 digital training courses, as well as AWS DeepLens and AWS DeepRacer, which developers can use to learn the basics of deep learning and reinforcement learning. AWS also offers a significant number of data and compute infrastructure options. Users can build on these options to create high-quality AI solutions in the cloud, but not on-premises.

Cautions

- The complexity and breadth of AWS's portfolio poses some challenges, both for individual developers and enterprises' application leaders. Developers must determine whether AWS's functionality is the best available for their requirements. Application leaders must decide whether they want to consolidate AI projects on a single platform. AWS's breadth can also lead to confusion and governance issues — for instance, whereas Amazon Lex offers comprehensive translation services, the NLU portion of the virtual assistant framework supports only English language intents.
- Throughout the period during which we conducted research for this Magic Quadrant, Amazon SageMaker lacked some critical model management capabilities, such as active model monitoring and explainability for inferences provided by models. However, AWS has since released Amazon SageMaker Experiments and Amazon SageMaker Model Monitor to address this shortcoming.
- When users move from development to production environments, the cost of execution may be higher than they anticipated. AWS is aware of the cost and complexity of prediction and offers products like Amazon Elastic Inference to help address this issue. Even so, application developers should model the likely production costs at an early stage in their projects and, in some cases, plan to move deep neural network compute workloads on-premises.

Google

[Google](#), headquartered in Mountain View, California, U.S., is a Leader in this Magic Quadrant. It offers services spanning all three of the service areas evaluated: language, vision and AutoML. Google uses deep neural network models for most of its services, which it delivers from its public cloud, Google Cloud Platform.

Strengths

- Google has strong language service offerings that support more languages than those of other providers. Its NLU service can be used in both batch mode and real-time-streaming mode. Its chatbot tools support 22 languages, with more planned for the coming year. Its speech-to-text service supports 120 dialects across 64 languages — significantly more than other providers. Its translation service supports 104 languages. Google also has the most extensive optical character recognition (OCR) support for languages, with over 200 supported.
- Unlike its other services, Google's image recognition service can be deployed via a container in a virtual private cloud or on-premises, as well as on its public cloud platform. Google's What-If Tool offers a degree of explainability for its AutoML Vision and data-labeling services. Additionally, AutoML Vision offers the Vision API, which gives developers access to pretrained models via REST and remote procedure call (RPC) APIs. Developers can quickly classify images using automatic recognition of faces and objects. Google's AutoML Video Intelligence service enables developers with minimal ML experience to customize models that can classify and track objects.

- Google's AutoML Tables includes data support that provides information about missing data, cardinality and distribution for each feature in a dataset. This helps developers put together clean, effective training data. In training, AutoML Tables performs common feature-engineering tasks, such as normalization of numeric features, creation of one-hot encoding and embeddings for categorical features. Google's AutoML models use transfer learning and neural architecture search capabilities.

Cautions

- Most of Google's cloud AI developer services are available only via the public Google Cloud Platform. Except for the Vision services deployable in Edge and Mobile ML Kits, services may not be deployed in a virtual private cloud or on-premises. Google does, however, plan to include these deployment options for many of its services in the coming year. Its services are based on subsymbolic deep learning models only, though it plans to add symbolic learning to its models in the coming year.
- The Google Cloud Platform commands a lower market share than the offerings of AWS and Microsoft (Azure). Google appointed Thomas Kurian as head of the Google Cloud Platform in November 2018 — a change that attracted positive feedback — but the organization is still undergoing substantial change, the full impact of which will not be apparent for some time.
- Google has struggled to compete with other major cloud providers in the general cloud platform business, due to prior shortcomings in terms of sales execution and pricing. Although Google has demonstrated improvement, it needs to execute its new targeted-account sales strategy in order to gain traction with its cloud AI developer service offering, as this offering's success depends on that of the overall Google Cloud Platform.

H2O.ai

[H2O.ai](#), headquartered in Mountain View, California, U.S., is a Visionary in this Magic Quadrant. Driven by a mission to democratize AI for everyone, H2O.ai is known for its open-source ML and AI software. H2O Driverless AI, an AutoML platform, is the vendor's commercial offering and incorporates many open-source add-ons. In 2019, H2O.ai introduced a unique approach incorporating "recipes" of AI design patterns, best practices and techniques that can be uploaded into H2O Driverless AI in its AutoML process. Over 130 open-source recipes created and curated by the world's top data scientists (Kaggle Grandmasters) provide specific use cases and solutions. Customers can consume, extend, customize or create new recipes. H2O.ai aims to enable developers to build AI that can power applications and services in all major clouds, as well as on-premises. It plans to further simplify AI for business users, accelerate its AI innovation, and expand its sales and marketing globally.

Strengths

- H2O Driverless AI includes automatic capabilities such as visualization, feature engineering, model building, time series support, natural language processing (NLP) and ML interpretability, supported by collaboration capabilities and flexibility in data processing, model development,

deployment and operations. It automatically generates a portable, low-latency model deployment artifact called MOJO that can be deployed in Java, Python, R and C++ runtimes, and from the cloud to the edge.

- H2O Driverless AI is simple to deploy, thus enabling customers to reduce the need for mature data science organizations and providing a shortcut to starting new AI projects. It makes difficult ML and AI concepts simple to understand via an intuitive UI and documentation that shows how to get the most out of the product. H2O.ai's focus on ML interpretability and explainability facilitates AI adoption by helping to build trust and mitigate risk. It also helps developers learn the nuances of data science. Additionally, H2O Enterprise Puddle is a managed service that provides a secure way to run H2O Driverless AI instances on a customer's virtual private cloud.
- H2O.ai offers excellent customer support. It received top scores from its reference customers for the quality of its technical support, and has a dedicated support organization and customer success engineers and data scientists to support its users. H2O.ai employs 10% of the world's Kaggle Grandmasters, who are available to work with customers. It offers a free trial and comprehensive onboarding for H2O Driverless AI.

Cautions

- H2O Driverless AI is currently used more by data scientists than by developers — only 20% to 30% of H2O's users are developers. Increasingly, however, developers have been critically involved in the company's AI solutions, including the deployment, integration and monitoring of H2O Driverless AI models in final AI applications. Consequently, H2O Driverless AI requires developers to know basic data science techniques. It is built on top of open-source libraries and frameworks that enable model visibility in order for developers to better understand the product — a necessity for developers gradually embracing the use of ML models.
- Some natural language services are unavailable from H2O.ai, though others, such as sentiment analysis and text analytics, are available on H2O.ai's platform. Other services, such as chatbots, speech to text, automatic speech, text to speech and translation, are currently not available from H2O.ai. This deficiency can, however, be mitigated by using third-party tools that are supported by H2O.ai, such as TensorFlow and PyTorch, or custom recipes.
- H2O.ai does not currently offer any vision services, but image recognition and OCR are works in progress and will be incorporated into H2O Driverless AI in 2020. In the meantime, customers who need such services could develop them using third-party tools that are supported by H2O.ai, such as TensorFlow, or custom recipes. H2O Driverless AI uses an image classification and segmentation recipe to take image files and classify them on the basis of labels — all of which is done automatically.

IBM

IBM, headquartered in Armonk, New York, U.S., is a Leader in this Magic Quadrant. It offers a wide range of AI services across all three service areas covered in this report. Although IBM's Research entity is a powerful force for leading thinking about AI, it remains difficult to productize leading-edge innovations. IBM's recent shift to a hybrid cloud strategy, wholly embracing multiple public cloud

infrastructures, combined with its acquisition of Red Hat, is helping its cloud AI service endeavors. It is doing so by providing developers with a much broader playing field and the ability to deploy their systems in a wider range of domains. IBM also offers flexible and extensive support for a wide variety of SDKs, thus giving developers considerable freedom when it comes to picking their preferred integrated development environment.

Strengths

- According to its users, developing conversational agents on IBM's Watson Assistant platform is a relatively painless experience. The capabilities within IBM's NLP elements are well integrated, and they include innovative components that use deep-learning synthesis techniques, and tone and sentiment analysis. IBM also provides dialogue management, document scanning and text analytics functionality.
- IBM uses its deep foundations in ML and deep learning to consolidate the functionality behind its Watson Visual Recognition and Watson Video Enrichment services. Users can also use IBM Watson Studio's prebuilt models, which employ advanced ML algorithms and have NLP capabilities.
- IBM's Watson Studio and ML products benefit from the company's long history in the field of augmented ML. IBM offers a robust set of AI ML services, ranging from automated data preparation and selection of algorithms through a series of optimized metrics to an effective citizen data scientist experience.

Cautions

- IBM's breadth of offering makes for a large number of AI service elements. This complicates its task of bringing them all together in a single, coherent and well-integrated environment. This complexity, in turn, results in different products, from different divisions, being handled by various development teams and having various pricing schemes — an issue reflected in feedback from IBM's users. They often commented about the unequal strength, versions, integration levels and backward-compatibility levels of its services. Developers should test their assembled systems in real implementation conditions.
- Partly due to the different origins, levels of maturity and granularity of the components used, pricing levels and schemes vary across IBM's services. The company's recent cloud strategy shift, which allows services to run on external providers' infrastructure in addition to its own platform, will likely compound this issue; however, with the flexibility to deploy on a hybrid cloud infrastructure, comes the complexity of handling multiple price and API models. Developers should estimate the overall cost of their combined systems in production environments before starting to develop prototypes, to avoid any surprises at production time.
- Despite having strong components across the NLP spectrum, IBM's lack of robust natural language generation (NLG) elements is surprising, though, as with many other techniques, IBM Research is developing capabilities that might eventually appear in its service lineup. Among IBM's vision capabilities, developers find the autolabeling of images limited, in comparison with what other vendors offer. Also, with regard to AutoML functionality, there is a need for stronger

cost estimation for the training of ML models and the ability to track models' mathematical, technical and business-related key performance indicators (KPIs).

Microsoft

Microsoft, headquartered in Redmond, Washington, U.S., is a Leader in this Magic Quadrant. As with the other Leaders, Microsoft offers a wide range of services that developers can utilize to enhance their applications. Microsoft's familiarity to many enterprise development teams may give it an advantage in this market — enterprises and enterprise developers are likely to have existing relationships with Microsoft and are used to working with the company's tools. Microsoft's offering includes language, vision and AutoML capabilities. It has invested heavily in the creation of AI capabilities and embeds many into its commercial products that people use with everyday AI features. It also contributes to the development of standards and guidelines for best practices associated with AI, including the ethical use of AI and efforts to make AI explainable.

Strengths

- Microsoft offers a range of services based on both symbolic (parsed) and subsymbolic (vector-based ML) approaches to ML and AI. This offers Microsoft advantages for some of its cloud AI services. Its cloud AI developer services are among the most comprehensive on the market and are all highly competitive.
- Microsoft is among the more flexible providers of cloud AI developer services in terms of deployment options. Its services are deployable in the Azure cloud, a virtual private cloud or on-premises, based on enterprise customers' needs.
- Microsoft offers one of the largest selections of languages for its natural language services. This enables multinational organizations that have wide-ranging language needs to serve multiple markets with varying language requirements.

Cautions

- Microsoft has yet to introduce NLG services, although they are on its roadmap. Companies with an immediate need for NLG must therefore look to another provider. Given that there are few such providers, this area could represent a future opportunity for Microsoft.
- Microsoft can be challenging to engage with, due to a confusing branding strategy that spans multiple business units and includes Azure cognitive services and Cortana services. This overlap often confuses customers and can frustrate them. Additionally, it can be difficult to know which part of Microsoft to contact, and often one part is unable to identify another business unit more appropriate for the customer.
- Microsoft often delivers AI services via its partner network and does not typically engage directly with enterprises to develop solutions. This approach is not unique to Microsoft, but it can nevertheless prove challenging for organizations new to developing cloud solutions themselves.

Prevision.io

[Prevision.io](#), headquartered in Paris, France, is a Visionary in this Magic Quadrant. Its set of cloud AI developer services is focused on AutoML, but there are also text and image analytics capabilities within its AutoML platform. Although this platform is primarily aimed at citizen data scientists (including developers), it also offers advanced features for experienced data scientists.

Strengths

- Behind the simplicity of Prevision.io's interface and controls lies a powerful, feature-rich ML platform. It has a good balance of ML model-building controls and more advanced capabilities accessible to experienced data scientists. Contextualized wizards guide developers and other citizen data scientists through a simple but powerful AutoML process. The platform helps temper the "black box" effect with a series of explanatory steps. It also offers cost estimation functions that help developers control the amount of resources needed to create models.
- Prevision.io's platform integrates advanced ML features. These include the abilities to build metamodels (an approach to ensemble modeling), automatically blend algorithms to produce more accurate models, and create embeddings to boost the performance of specific models.
- Prevision.io's interpretable AI functionality aims to "unpack" black-box models, so that users can understand how an ML model reaches its conclusions. The platform automatically generates simplified interpretable models and can visualize the behavior of these models. Users can explore the features and logic that contribute to decisions. The company also enables developers to interact with a model in order to understand its behavior better.

Cautions

- Although accessible to developers worldwide, Prevision.io essentially operates from only a few European countries. Even if, by definition, its platform should be self-sufficient, it is important to note that its support is concentrated in Europe. That said, Prevision.io, which is growing rapidly, aims to establish a presence in North America by 2021.
- A deployed model is not necessarily a managed model. Prevision.io's platform provides developers with schedulers and a repository mechanism (through its Store environment), but more functionality will be necessary to enable market-leading operationalization. The ability to monitor models' telemetry within applications and business processes (through technical, mathematical and business-related KPIs) will become crucial for developers.
- In environments where expert data scientists are not always available, thorough verification and validation procedures are crucial. These procedures are especially important at the front end of the process in order for data to be validated, to ensure not just its cleanliness and accessibility but also its relevance to the problem at hand. They are also important at the back end of the process to ensure proper instigation of model audit procedures. Prevision.io offers only basic capabilities in this regard, which it will need to expand if it is to remain a Visionary.

Salesforce

Salesforce, headquartered in San Francisco, California, U.S., is a Niche Player in this Magic Quadrant. For application developers and designers, it offers embedded AI capabilities in its cloud offerings. The company's central aim is to put AI in the hands of every CRM and front-office user by enabling administrators and developers in Salesforce's ecosystem who lack deep learning or ML knowledge to embed image recognition and NLP within their applications. Salesforce Einstein can be provided in the form of APIs or with a UI, which is more suitable for developers. Einstein Platform Services cover multiple core areas of functionality, with Einstein Prediction Builder, Einstein Language and Einstein Vision services, and Einstein Discovery.

Strengths

- A major strength of the Salesforce Einstein Platform services is the simplicity of the pipeline — predictions are built in the same environment where the training/test data resides, and results are written back into the same Salesforce objects (tables/fields). This results in fast model development, training and testing, and swift deployment into existing workflows and processes.
- Salesforce automates much of the data wrangling involved in developing predictions. It has also implemented features to ensure the quality of predictions and to reduce bias in relation to training data.
- Salesforce leads in terms of providing a simple, democratized approach to developing AI capabilities, one that requires no technical or development skills on the part of the user. Many skills are targeted at Salesforce administrators and business users, and the ability of Trailhead will make the AI service capabilities easy to acquire for both professional developers and citizen developers.

Cautions

- Salesforce Einstein provides two NLP capabilities — Einstein Sentiment and Einstein Intent — but both are rudimentary in comparison with the conversational APIs offered by most other vendors in this Magic Quadrant. Einstein Sentiment uses a global model that cannot be trained, nor can its behavior be modified. The Einstein Intent API simply takes a list of utterances and intent classifications as input to the training process. Buyers looking for a more comprehensive set of API services for developing conversational experiences may find more suitable options elsewhere, but this is a strategic area of investment for Salesforce with its Einstein Voice product line.
- Salesforce's Einstein Vision service is simple to use, and the approach taken is largely the same as that used for language services for which training data with labels is uploaded and the model automatically trained. However, specifying object detection bounding-box coordinates for training requires third-party tools and support. Furthermore, the Einstein Vision service lacks many of the features offered by competitors, such as data label outsourcing, detection of explicit content and edge-based model deployment. Einstein Vision APIs, though useful in limited sales, support and marketing contexts, are not suitable for broad or industrial use cases beyond CRM.

- Although Salesforce has rapidly increased its AI capabilities through acquisitions, the focus has been on empowering Salesforce Clouds with easy-to-use prediction, mining and segmentation tools. Enterprises with needs beyond Salesforce's core customer-centric use cases may find they need supplementary services from other vendors. Salesforce's AI services offering is much slimmer than those of other vendors in this Magic Quadrant, being confined to Einstein Platform Services, which provide only targeted developer services, not a full portfolio.

SAP

SAP, headquartered in Walldorf, Germany, is a Niche Player in this Magic Quadrant. It has two main offerings in this market. The first, Leonardo Machine Learning, provides developers and business users with both pretrained models and customizable models that can be used as a web service by calling simple REST APIs. They cover the service areas of vision, language and, to some degree, AutoML. The second offering is SAP Conversational AI, an end-to-end conversational bot building platform.

Strengths

- SAP provides many foundational AI services (both natively and via partners) across the areas of language and vision, but more important is its packaging of them together as “business services” or “scenarios” — an approach that could become a major differentiator. Scenarios can be used to represent workflows, such as invoice processing, or vision-workflow-related tasks. When combined with some of SAP's prebuilt models, this approach creates a quick development environment. (Note that, although the foundational services are available as APIs, the business service “macros” must be consumed inside an SAP environment.)
- Model management can be performed via API or the dedicated SAP Data Intelligence platform. The benefit of model management via API is that it enables third-party environments to provide master model management. Retraining models is also relatively straightforward, with simple retraining thresholds provided by default, but also amenable to configuration. Furthermore, developers can bring their own models (created with TensorFlow, R or scikit) into the Leonardo Machine Learning environment.
- Leonardo Machine Learning integrates easily with other SAP products and offers the clear benefit of making data from SAP repositories available. Furthermore, SAP has a comprehensive and progressive framework for bringing its native SAP employee-centric chatbots — integrated into its business unit platform offerings (such as those for ERP, HR and help desk) — together with custom consumer-centric bots. SAP Conversational AI provides language technology and a chatbot-building platform that enables development of digital assistants. This digital assistant acts as an aggregator for both SAP chatbots and non-SAP chatbots built using other chatbot-building frameworks.

Cautions

- SAP Conversational AI offers merely basic greetings and generic conversational elements, and SAP's vision services have only generic classifiers bundled with them, though SAP also offers

vision services from other providers. Prospective customers must be prepared to create their own training data and industry solutions — a task that SAP should attempt to ease by drawing on its huge network of customers, partners and system integrators.

- SAP's AutoML features are tied heavily to its language and vision services, rather than being a broader AutoML offering. The AutoML services offered focus on workflow and high-level configuration for retraining existing generic models as custom models. Feature engineering, again tied to vision and text/document analysis, creates output vectors for developers to compare. SAP has addressed the ability to develop more data pipeline and data science tools, as well as automated data preparation, with the release of SAP Data Intelligence, which was not, however, available in time for evaluation in this Magic Quadrant.
- Although SAP has rapidly increased its deployment of services, there were a few notable omissions at the time of evaluation. For example, vision for video content was lacking in its data preparation capabilities. SAP has, however, subsequently addressed the need for such capabilities with the release of SAP Data Intelligence.

Tencent

Tencent, headquartered in Shenzhen, China, is a Niche Player in this Magic Quadrant. A top 10 company in terms of market capitalization, it is perhaps best known as the world's largest gaming company, and for its billion-user WeChat and QQ messaging, social media and mobile payment applications. It has hundreds of subsidiaries in numerous industries and is one of the world's largest venture capital companies. Tencent launched its AI Lab in 2016 and in the same year opened an AI research center in Seattle, Washington, U.S.

Strengths

- Tencent uses its extensive AI resources for its gaming, vision, chat, and other products and services. Tencent AI Lab focuses on fundamental AI research; Tencent YouTu Lab focuses on face and image recognition. WeChat AI works on speech recognition, NLP, data mining and ML.
- Tencent has a full AI portfolio. Its language services include speech recognition, NLP, the Tencent Bot Platform and language generation. Its vision services include image recognition and OCR. Its AutoML service includes feature engineering, automated model building, model management, and multiple algorithms for data preprocessing, transformation and selection.
- Tencent has significant presence and visibility in the large Chinese market. Its ability to offer access to its user base makes its platform attractive to product and service providers.

Cautions

- Developers may find Tencent's AI marketing, messaging and offers fragmented and difficult to understand. There are multiple AI-related groups within Tencent, which conduct some overlapping AI research and development. Tencent also offers the AI Open Platform, which is a community where developers can get free trials of Tencent AI services. Their projects are redirected to the Tencent Cloud, once they need enterprise-level services.

- Although Tencent has a high profile in China, its AI functions lack the visibility and the language support needed to attract wide use by developers outside China.
- Developers may find that Tencent's AI services lack functionality in some areas. For instance, its AutoML services lack automated data preparation, and some of the feature-engineering functions are not available via API. The speech-to-text service supports only Mandarin, Cantonese, English, Tibetan and Uighur. The text analytics feature supports only English and Chinese.

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

Not applicable, as this is the first edition of this Magic Quadrant.

Dropped

Not applicable, as this is the first edition of this Magic Quadrant.

Inclusion and Exclusion Criteria

Inclusion Criteria

To qualify for inclusion, vendors must:

- Demonstrate a go-to-market strategy for cloud AI developer services.
- At a minimum, the services included must support two or more of the language services *and* one or more of the vision services *or* automated machine learning services listed below:
 - **Language services:**
 - **Speech to text or automatic speech recognition (ASR):** This service is a subset of computational linguistics that will take analog input and convert it to text output. This text output can be the final product or it can be entered into a natural language understanding model so that metadata can be extracted. Many computing devices have some ASR capability built in, such as PCs and smartphones.
 - **Natural language understanding (NLU):** This service is a subset of natural language processing (NLP) that deals with machine comprehension. It takes textual input and

extracts metadata from the text. Extracting metadata is relatively straightforward, but being able to understand the intent of the person entering the text is challenging and often requires supplemental models. These models can be built using semantic techniques, deep neural network models or a combination of both. NLU is often used as part of a chatbot or virtual assistant solution.

- **Chatbot virtual assistant frameworks/dialogue management:** These frameworks allow enterprises and/or their partners to build models that supplement the hosted language services by allowing enterprises to add domain-specific learning to support the specific use cases needed by the business.
- **Natural language generation (NLG):** Services that create natural language from a machine representation such as concepts, datasets or minimal descriptions in a knowledge base or a logical form such as a return form that generates a letter to the customer. The body of text delivered in a natural language form can be thought of as translating data into language.
- **Text to speech:** This service converts textual input into analog output or speech. This can be done via normal conversion of text to speech or via symbolic linguistic representations such as phonemes.
- **Translation:** This service takes text input from the source language and converts it to a target language as output. This is a very challenging task as it is not just translating one word to the corresponding word in another language. Differences in language structure makes accurate services very challenging to achieve.
- **Sentiment analysis (emotion AI):** This service analyzes the words that are typically entered into a conversational or social model to analyze the language for positive, negative or neutral sentiment based on the words that are chosen by the user. These services may be supplemented by other services that analyze the tone of analog inputs.
- **Text analytics:** This service analyzes unstructured text using algorithms to extract elements such as concepts, topics and keyword attributes and add these as metadata.
- **Vision services:**
 - **Image recognition:** This service normally identifies what objects or people are contained in an image. Some implementations can also identify attributes of the elements in the image, such as color or patterns. A common use for this is to identify whether people or items of interest are in an image, and to add metadata to classify or tag images.
 - **Video content analysis:** This service normally combines image recognition and automated speech recognition to identify people and objects in the video, as well as to create a transcript for the audio portion of the video. Some services also identify objects or people in the video and track direction for the people across multiple frames of the video.

- **Optical character recognition:** This service converts electronic images of typed, handwritten, printed text or text in images or video into machine-encoded text and adds metadata to the content.
- **Machine learning services:**

Automated machine learning (AutoML): These services allow people without significant ML or data science skills to customize the services listed above or build specific-purpose models to predict actions or measures to meet their specific needs. Using AutoML services, developers can create custom models or supplemental models to be used in conjunction with the existing general services.

 - **Automated data preparation:** This service can prepare the datasets to be used for training of the models. It can cleanse and augment datasets from the raw data that is provided by the enterprise.
 - **Feature engineering:** This is a service that adds metadata to the dataset that is submitted to the provider. This augmentation can be done via the use of crowdsourcing capabilities or via ML models that evaluate the data and add metadata to the submitted dataset in an automated manner.
 - **Automated model building:** These models require the users to provide datasets that can be used to train the models. In addition to the raw data, users must provide datasets that include metadata tags that have the attributes they want the models to be trained to identify. They must also identify the variable that they want the model to predict. These AutoML services analyze the data and evaluate and recommend or select the potential algorithms that can be used to build and optimize the model based on the best results. They automatically optimize the performance and accuracy of the model by tuning the hyperparameters of the model.
 - **Model management:** These services provide analytics and, in some cases, allow for corpus management for data used to train models to ensure that models continue to perform optimally.

Additionally, at minimum, vendors had to fulfill one of the following criteria. *Either:*

- The vendor must have at least \$15 million annual revenue in 2018 from its cloud AI developer service offerings.

Or:

- The vendor must have added at least 20 new paying enterprise customers for its cloud AI developer services in 2018 across at least two of the following geographies:
 - North America
 - South America
 - Europe

- Middle East and Africa
- Asia/Pacific

Exclusion Criteria

We excluded vendors that:

- Offer only language-based services.
- Offer only vision-based services.
- Only offer these services as part of a professional services contract, where the services are used exclusively by the vendors' consultants.
- Services that are not native services created and delivered by the vendor.

Evaluation Criteria

Ability to Execute

Product or Service: Core goods and services offered by the vendor for the defined market. This includes current product/service capabilities, quality, feature sets, skills and so on, whether offered natively or through OEM agreements/partnerships as defined in the market definition and detailed in the subcriteria.

Overall Viability: Viability includes an assessment of the overall organization's financial health, the financial and practical success of the business unit, and the likelihood that the individual business unit will continue investing in the product, will continue offering the product and will advance the state of the art within the organization's portfolio of products.

Sales Execution/Pricing: The vendor's capabilities in all presales activities and the structure that supports them. This includes deal management, pricing and negotiation, presales support, and the overall effectiveness of the sales channel.

Market Responsiveness/Record: Ability to respond, change direction, be flexible and achieve competitive success as opportunities develop, competitors act, customer needs evolve and market dynamics change. This criterion also considers the vendor's history of responsiveness.

Marketing Execution: The clarity, quality, creativity and efficacy of programs designed to deliver the organization's message to influence the market, promote the brand and business, increase awareness of the products, and establish a positive identification with the product/brand and organization in the minds of buyers. This "mind share" can be driven by a combination of publicity, promotional initiatives, thought leadership, word of mouth and sales activities.

Customer Experience: Relationships, products and services/programs that enable clients to be successful with the products evaluated. Specifically, this includes the ways customers receive

technical support or account support. This can also include ancillary tools, customer support programs (and the quality thereof), availability of user groups, service-level agreements and so on.

Operations: The ability of the organization to meet its goals and commitments. Factors include the quality of the organizational structure, including skills, experiences, programs, systems and other vehicles that enable the organization to operate effectively and efficiently on an ongoing basis.

Table 1. Ability to Execute Evaluation Criteria

Evaluation Criteria	Weighting
Product or Service	High
Overall Viability	Medium
Sales Execution/Pricing	Medium
Market Responsiveness/Record	Medium
Marketing Execution	Medium
Customer Experience	Medium
Operations	Medium

Source: Gartner (February 2020)

Completeness of Vision

Market Understanding: Ability of the vendor to understand buyers' needs and translate these needs into products and services. Vendors that show the highest degree of vision listen to and understand buyers' wants and needs, and can shape or enhance those wants with their added vision.

Marketing Strategy: A clear, differentiated set of messages consistently communicated throughout the organization and externalized through the website, advertising, customer programs and positioning statements.

Sales Strategy: The strategy for selling products that uses the appropriate network of direct and indirect sales, marketing, service and communication affiliates that extend the scope and depth of market reach, skills, expertise, technologies, services and the customer base.

Offering (Product) Strategy: The vendor's approach to product development and delivery that emphasizes differentiation, functionality, methodology and feature sets as they map to current and future requirements.

Business Model: The soundness and logic of the vendor's underlying business proposition.

Vertical/Industry Strategy: The vendor’s strategy to direct resources, skills and offerings to meet the specific needs of individual market segments, including vertical markets.

Innovation: Direct, related, complementary and synergistic layouts of resources, expertise or capital for investment, consolidation, defensive or preemptive purposes.

Geographic Strategy: The vendor’s strategy to direct resources, skills and offerings to meet the specific needs of geographies outside the “home” or native geography, either directly or through partners, channels and subsidiaries, as appropriate for that geography and market.

Table 2. Completeness of Vision Evaluation Criteria

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

Source: Gartner (February 2020)

Quadrant Descriptions

Leaders

Leaders have robust offerings in all three key service areas: language, vision and AutoML. Their services are API-accessible and do not require developers to have data science expertise. Leaders also have ancillary services that support or enhance the capabilities of their core services.

Challengers

Challengers are often large businesses with substantial assets. They may have the resources to invest significantly in services, but limited portfolios because of other priorities that compete for their attention and resources.

Visionaries

Visionaries are likely to excel in AutoML, as this segment is viewed as the most important for application leaders and development organizations. Their AutoML services facilitate customization of ML models. These models can be developed to deliver predictions, classifications, next-best-action recommendations and numerous other uses that can enhance applications being delivered by enterprises.

Niche Players

Although one might expect Niche Players to offer only language or vision services, this Magic Quadrant requires providers to have, at minimum, both language and vision services or only AutoML services. Providers that offered only language or vision services were therefore excluded from evaluation. Niche Players may have limited penetration outside their home region. Although this Magic Quadrant requires providers to operate in at least two regions, these providers may be new to regions outside their home region. Additionally, Niche Players may be resource-constrained.

Context

This Magic Quadrant excludes vendors that offer only language services or only vision services. This is because we believe most developers want a range of services from one provider that enable them to create the ML models they need, which include both vision and language models. For the same reason, Gartner chose to include vendors that offer only AutoML services, as these services can be used for any use case a developer requires. At this relatively early stage in this market's development, we believe most development teams will use ML models for a wide variety of use cases. They are likely to want to move fluidly from enabling applications with language-based services, to doing the same with vision-based services and custom models, with no continual emphasis on one type of service.

For enterprises that need only language services or only vision services for specific purposes, several providers not evaluated in this Magic Quadrant could prove valuable partners. For enterprises that see a need for only language services or only vision services and no additional ML capabilities in the short term, it could be of value for the enterprise to broaden its consideration for these vendors.

Enterprises that want to use ML services or models in the Chinese market should consider Alibaba and Baidu, in addition to Tencent. (Alibaba and Baidu were excluded from this Magic Quadrant because they did not meet its requirement for offerings that span at least two major regions.)

The reason this Magic Quadrant focuses on three types of service — language, vision and AutoML — is that these services are the most mature. There are, however, many AI and ML services that focus on specific use cases, where models are built to address a predefined need or use case. Many of these models may be of interest to development teams. There are too many such offerings to discuss in this Magic Quadrant, however. Each should be evaluated on the merits of the particular capabilities of the models delivered.

Market Overview

Many different types of vendor are aiming AI and ML services at developers. They include cloud vendors, platform-as-a-service vendors, enterprise software vendors, and specialist AI and ML vendors. All realize that the short-term opportunities to utilize ML models in applications far exceed the capabilities of the limited number of data scientists and ML experts available. Furthermore, the potential use cases, which vary widely, have the potential to yield significant business benefits.

The three basic categories of language, vision and AutoML services have emerged to address these opportunities. Although these categories may seem narrow at first, there are many purposes to which these services can be put within enterprises. Some examples:

- **Language:** Virtual assistants and voice-initiated business processes
- **Vision:** Multifactor authentication workplace monitoring and asset cataloging
- **AutoML:** Sales lead scoring, risk assessment and next-best-action recommendation

Development teams are the natural audience for these services. These teams incorporate AI and ML models into existing enterprise applications to enhance their value. Developers know how to create applications that call APIs and incorporate the responses into applications. In short, development teams are the consumers of AI and ML models. They will operationalize them and thus create value for businesses.

Developers will use existing models and develop customized models to enhance their applications. They should be able to use an existing hosted service or an automated custom model-building service to enhance applications that already exist or that are being developed. Even if an enterprise has data scientists and ML experts, its needs will outstrip their availability and many of their capabilities. This is why vendors are aiming these services at developers.

There are many easy opportunities that, if seized, could deliver significant returns to an enterprise, based on simple classification of existing assets. Simply applying metadata to unstructured assets can yield a strong ROI. The ability to identify patterns in machine diagnostics could enable prediction of the need for repairs before failures occur. The identification of patterns in insurance claims could help predict fraud. The applications of simple ML models that are not hard to build and do not merit the attention of expensive data scientists are too numerous to list fully here.

Development teams that access the kind of services evaluated in this Magic Quadrant do not need to know how to use deep learning frameworks or have detailed knowledge of TensorFlow. They simply need to know how to incorporate an API call into an application to enhance its value. This is something that developers are already very familiar with. They can derive value from language and vision services in their current form — and, if they customize these services with AutoML, that value can be multiplied many times and tailored to the needs of a business. Customizing services' models may require data selection, cleansing and augmentation via feature engineering, but many of the service providers evaluated in this Magic Quadrant offer these services as part of their AutoML portfolios. In addition, such ancillary services are available from third parties. This means that any competent developer can build ML models and use them to benefit a business.

This market is far from mature, however, with new services and capabilities being introduced rapidly. This trend will continue, such that gaps in major vendors' service portfolios are likely to be filled in the next few years. In addition, we expect many vendors to apply AI and ML services to specific use cases and to offer these services to developers. All this will make for a confusing and volatile market, especially as AI continues to advance and mature. Already, even established models are being rebuilt from the ground up with some frequency. Additionally, services that become leading performers may be superseded not long after by fresh advances from other providers. It is therefore important to monitor these services' evolution.

Cloud AI developer services are available from a variety of providers, and some interesting vendor groupings are apparent. These are described below, for each of the quadrants that have vendors in them in this Magic Quadrant:

- **Leaders quadrant:** This quadrant contains large cloud vendors that offer a wide range of hosted services, including multiple language, vision and AutoML services. These providers give developers the greatest flexibility to use a wide range of services to enhance their application suites with ML models and services. The same vendors have shown continued leadership in the AI and ML markets, with heavy investment in this rapidly evolving field. The scope and variety of the services they offer to developer teams distinguishes them from other competitors.
- **Visionaries quadrant:** This quadrant is populated by emerging vendors in the AutoML segment. These vendors are important because they offer development teams innovative solutions that facilitate the preparation and building of models of all types without requiring expertise in deep learning or ML frameworks. These automated services enable developers to build and deploy ML models easily in order to augment applications' capabilities, without having to learn Python or R or use other tools.
- **Niche Players quadrant:** This quadrant hosts two groupings of vendors — or at least has the potential to. The first comprises two enterprise software platform vendors that are increasingly embedding AI and ML capabilities into their platforms to enhance the applications developers build using them. Although not as flexible as pure API-callable services, these capabilities are easy to use as part of each vendor's platform. The second grouping currently contains only one large Chinese cloud AI service provider, but, were it not for this Magic Quadrant's requirement for customers in at least two major regions of the world, Tencent would be joined by other vendors of a similar nature.

Gartner Recommended Reading

Some documents may not be available as part of your current Gartner subscription.

“How Markets and Vendors Are Evaluated in Gartner Magic Quadrants”

“Machine Learning Alters the Role of the Developer”

“Critical Capabilities for Cloud AI Developer Services”

Evaluation Criteria Definitions

Ability to Execute

Product/Service: Core goods and services offered by the vendor for the defined market. This includes current product/service capabilities, quality, feature sets, skills and so on, whether offered natively or through OEM agreements/partnerships as defined in the market definition and detailed in the subcriteria.

Overall Viability: Viability includes an assessment of the overall organization's financial health, the financial and practical success of the business unit, and the likelihood that the individual business unit will continue investing in the product, will continue offering the product and will advance the state of the art within the organization's portfolio of products.

Sales Execution/Pricing: The vendor's capabilities in all presales activities and the structure that supports them. This includes deal management, pricing and negotiation, presales support, and the overall effectiveness of the sales channel.

Market Responsiveness/Record: Ability to respond, change direction, be flexible and achieve competitive success as opportunities develop, competitors act, customer needs evolve and market dynamics change. This criterion also considers the vendor's history of responsiveness.

Marketing Execution: The clarity, quality, creativity and efficacy of programs designed to deliver the organization's message to influence the market, promote the brand and business, increase awareness of the products, and establish a positive identification with the product/brand and organization in the minds of buyers. This "mind share" can be driven by a combination of publicity, promotional initiatives, thought leadership, word of mouth and sales activities.

Customer Experience: Relationships, products and services/programs that enable clients to be successful with the products evaluated. Specifically, this includes the ways customers receive technical support or account support. This can also include ancillary tools, customer support programs (and the quality thereof), availability of user groups, service-level agreements and so on.

Operations: The ability of the organization to meet its goals and commitments. Factors include the quality of the organizational structure, including skills, experiences, programs, systems and other vehicles that enable the organization to operate effectively and efficiently on an ongoing basis.

Completeness of Vision

Market Understanding: Ability of the vendor to understand buyers' wants and needs and to translate those into products and services. Vendors that show the highest degree of vision listen to and understand buyers' wants and needs, and can shape or enhance those with their added vision.

Marketing Strategy: A clear, differentiated set of messages consistently communicated throughout the organization and externalized through the website, advertising, customer programs and positioning statements.

Sales Strategy: The strategy for selling products that uses the appropriate network of direct and indirect sales, marketing, service, and communication affiliates that extend the scope and depth of market reach, skills, expertise, technologies, services and the customer base.

Offering (Product) Strategy: The vendor's approach to product development and delivery that emphasizes differentiation, functionality, methodology and feature sets as they map to current and future requirements.

Business Model: The soundness and logic of the vendor's underlying business proposition.

Vertical/Industry Strategy: The vendor's strategy to direct resources, skills and offerings to meet the specific needs of individual market segments, including vertical markets.

Innovation: Direct, related, complementary and synergistic layouts of resources, expertise or capital for investment, consolidation, defensive or pre-emptive purposes.

Geographic Strategy: The vendor's strategy to direct resources, skills and offerings to meet the specific needs of geographies outside the "home" or native geography, either directly or through partners, channels and subsidiaries as appropriate for that geography and market.

More on This Topic

This is part of an in-depth collection of research. See the collection:

- [The AI Crunch Demands Practical Responses: A Gartner Trend Insight Report](#)

GARTNER HEADQUARTERS**Corporate Headquarters**

56 Top Gallant Road
Stamford, CT 06902-7700
USA
+1 203 964 0096

Regional Headquarters

AUSTRALIA
BRAZIL
JAPAN
UNITED KINGDOM

For a complete list of worldwide locations,
visit <http://www.gartner.com/technology/about.jsp>

© 2020 Gartner, Inc. and/or its affiliates. All rights reserved. Gartner is a registered trademark of Gartner, Inc. and its affiliates. This publication may not be reproduced or distributed in any form without Gartner's prior written permission. It consists of the opinions of Gartner's research organization, which should not be construed as statements of fact. While the information contained in this publication has been obtained from sources believed to be reliable, Gartner disclaims all warranties as to the accuracy, completeness or adequacy of such information. Although Gartner research may address legal and financial issues, Gartner does not provide legal or investment advice and its research should not be construed or used as such. Your access and use of this publication are governed by [Gartner Usage Policy](#). Gartner prides itself on its reputation for independence and objectivity. Its research is produced independently by its research organization without input or influence from any third party. For further information, see "[Guiding Principles on Independence and Objectivity](#)."