

REPORT REPRINT

Kogentix shines a light on AMP, looks to build an automated machine-learning business

KRISHNA ROY

21 AUG 2017

Having entered the machine-intelligence fray with an automated machine-learning platform in early 2017, the startup has illuminated early uses cases, in addition to detailing the technical bedrock and target market for its offering.

THIS REPORT, LICENSED TO KOAGENTIX, DEVELOPED AND AS PROVIDED BY 451 RESEARCH, LLC, WAS PUBLISHED AS PART OF OUR SYNDICATED MARKET INSIGHT SUBSCRIPTION SERVICE. IT SHALL BE OWNED IN ITS ENTIRETY BY 451 RESEARCH, LLC. THIS REPORT IS SOLELY INTENDED FOR USE BY THE RECIPIENT AND MAY NOT BE REPRODUCED OR RE-POSTED, IN WHOLE OR IN PART, BY THE RECIPIENT WITHOUT EXPRESS PERMISSION FROM 451 RESEARCH.



©2017 451 Research, LLC | WWW.451RESEARCH.COM

Kogentix was privately beta-testing its machine-learning offering when we caught up with it in January. The two-year-old startup's automated machine-learning platform (AMP) is now in production with customers, enabling management to illuminate use cases. Kogentix has also fully elucidated the technology behind AMP, along with its go-to-market strategy.

THE 451 TAKE

As we noted in our previous report, some folks may be confused by what Kogentix has to offer, given that the vendor started out providing professional services for Hadoop, announced an intention to deliver Hadoop-based business apps in 2015 and subsequently shifted focus to enabling artificial-intelligence-driven applications using AMP. However, we feel any potential confusion has lifted now that AMP is out in the market and customer, as well as pilot, use cases have emerged. While Kogentix is operating in a market full of potential rivals, AMP has some solid differentiators, which the startup should build on by raising series A funding later in 2017.

CONTEXT

Kogentix is looking to build a business in the vibrant but increasingly crowded machine-learning sector. As we noted in a previous report, Kogentix is no ordinary startup. The firm is fairly large and growing quickly. It cites a headcount of 250 employees, up from 200 in January. Kogentix has various offices around the world, including worldwide headquarters in Schaumburg, Illinois; a recently opened office in Jakarta; and an operation in Hyderabad, India.

Furthermore, Kogentix already has a revenue stream from consulting engagements, which the firm has used to build AMP. That said, these days the startup is focused on product sales, noting that it doesn't engage with a customer or prospect unless there is an opportunity to sell its AMP machine-learning stack. AMP became generally available in February, after a fairly long private beta period.

Kogentix intends to raise series A funding later in 2017. In the meantime, the startup is focused on building an initial set of customers and communicating use cases for it.

CUSTOMERS

Kogentix reports that it has the first four paying customers now in production for AMP, and is working with a further 27 accounts, 12 of which are in the US and 15 of which are in Asia-Pacific. Initial customer traction – along with the firm's involvement in a fair number of pilots – has enabled Kogentix to discern some adoption and use-case patterns.

The startup cites customer-360 use cases as a key usage scenario. AMP is being employed to help with customer churn, marketing campaigns, and other marketing use cases in the telco, banking, pharmaceutical and consumer packaged-goods industries. AMP is also in use for Internet of Things (IoT) scenarios involving the prediction of machine failure, and for outlier and anomaly detection in the healthcare and finance sectors, according to management.

Kogentix takes the view that data shouldn't leave a customer's platform, so AMP is designed to deploy where a company's data resides. Current AMP deployments are either on-premises or in a customer's private cloud, although the startup supports public clouds, including AWS and Microsoft Azure.

TECHNOLOGY

AMP is pitched as an end-to-end platform for building machine-learning-driven applications. It is designed to handle all aspects of the creation of these apps – from the collection, storing, and integration of data in multiple formats and states to the deployment of the apps into some sort of operational environment. AMP is being primed to make data scientists more efficient. It is also targeted at non-data scientists with subject matter expertise who want to create machine-learning-driven apps without data science knowhow, using the automation capabilities built into it.

The offering's features and capabilities can be broken down into four categories: data ingestion and discovery, feature management, model lifecycle management, and operations and tracking. AMP's data ingestion functions include the ability to ingest data in batch and real time, and within a company's internal and external data sources.

AMP also provides automated data correlation and feature extraction, data profiling to aid in the data-discovery process, and surfacing of data distribution and trends to provide further insight.

AMP's feature management capabilities are for data wrangling and data preparation. Kogentix is looking to provide a full set of feature management facilities inside AMP – not just basic data-joining functions. AMP has a drag-and-drop interface for data wrangling, and will augment data, if it is sparse, by surfacing other relevant data that the user may want to employ. It also includes automated event correlation and facilities to operationalize a feature engineering pipeline by essentially giving users a data flow.

The offering's model-building capabilities are designed to support the whole process – from building the model to training and validating it. Kogentix has also built recommendations into the model-building process to suggest the best model from an ensemble of machine-learning algorithms. AMP supports clustering, classification, regression, time series and graph-based modeling. Kogentix is planning to add support for deep-learning algorithms within AMP by year-end. The offering suggests the most appropriate model, and surfaces features and models that are not relevant. It also houses a rules engine, so the user can augment machine-learning models with rules.

A data scientist could also use R or Python languages to configure models, as an alternative to using the offering's automated modeling functions, which are also designed to support model validation and training via a one-click process. Users can also build models using Spark's MLlib machine-learning environment. Model-building is also supported via a drag-and-drop interface for ease of use.

AMP's operations and tracking features aim to provide an integration point into a system of record in order to create business results out of insight – otherwise known as operationalization. AMP also provides a window into how the model is performing. It provides key performance indicators, so a user can ascertain, for example, if customer churn is reducing and how it is reducing. AMP also automatically retrains and recalibrates in order to handle situations where a prediction is no longer aligned with the actual data. The offering is designed to enable the user to deploy the model as a streaming application or REST API in one click.

It is also worth noting that AMP requires Spark as a prerequisite because it uses Spark as a processing engine. AMP is designed to run natively on Amazon EMR in Spark, as well as in other Hadoop distributions. Cloudera is an initial go-to-market partner, although management notes that AMP can run on any Hadoop distribution.

COMPETITION

IBM Watson, DataRobot and Dataiku are the cited primary competitors AMP has squared up to so far. We would expect Kogentix to encounter all three because they are high-profile machine-learning platforms, addressing many of the same functional areas and use cases as Kogentix's machine-learning stack.

Kogentix is looking to differentiate AMP from DataRobot, which is also focused on automating machine-learning-driven analytics, by providing a broader set of capabilities than push-button predictive modeling, which is DataRobot's primary purview. The startup is seeking to single itself out from Dataiku, which is also pitching a machine-learning platform play, through features such as AMP's rules engine that enables rules to be included in the model-building process.

Additionally, Kogentix cites SAS Institute as a competitor, noting that SAS is often incumbent in many of the accounts the startup is targeting. SAS is also looking to up its ante in machine-learning-driven analytics – an area the firm has served for many years, but is now targeting with a packaged offering on its new Viya cloud-friendly architecture.

As we have previously noted, the machine-learning-driven advanced analytics sector is a crowded one, with many fellow young guns, including Nara Logics, Continuum Analytics, DMWay Analytics, BigML, Yottamine Analytics, H2O.ai, Domino Data Lab, Alpine Data, FORMCEPT and Ayasdi, competing for a piece of this business, albeit with different technology and approaches to each other, and to Kogentix. We think Formcept's offering, with its focus on enabling data-driven apps, is probably the closest in nature to AMP of this group.

It is worth noting that we have seen the first wave of M&A in 2017. DataRobot reaching for Nutonian to add time-series analytical modeling to its offering, and Progress Software picking up DataRPM in order to deliver a platform for building and delivering AI-driven apps are prime examples of acquisition activity.

SWOT ANALYSIS

STRENGTHS

Automating machine learning while providing visibility into it is a genuine market requirement that AMP is able to fulfill. The offering also houses some differentiated capabilities, such as a rules engine built into its modeler layer, and fairly extensive data discovery and preparation capabilities.

WEAKNESSES

The startup suffers from a low market profile. AMP may take some evangelism since it is a first-generation offering from a relatively unknown player.

OPPORTUNITIES

Raising series A funding will provide the startup with fuel for product and business development, as well as giving it investor validation, but the startup is right to land marquee accounts and an initial customer base for AMP before seeking an infusion of venture capital.

THREATS

AMP will elicit comparisons with a number of established and high-profile machine-learning-driven analytics offerings. Communicating the difference between the likes of SAP, Oracle and Salesforce's AI-driven apps approach and AMP is crucial - they have already built machine intelligence into their respective apps, while AMP is a platform for crafting apps of this ilk.