Machine Learning is Driving an Innovation Wave in SaaS Software

By Jeff Houston, CFA
jhouston@Navidar.com | 512.364.2258

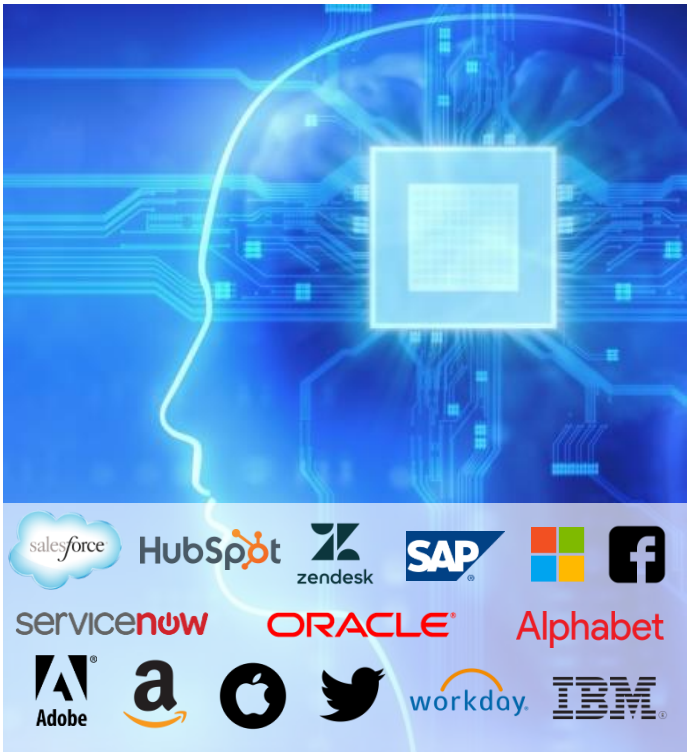

SaaS software vendors are enhancing their solutions with machine learning (ML) algorithms. ML is the ability of a computer to either automate or recommend appropriate actions by applying probability to data with a feedback loop that enables learning. We view ML as a subset of artificial intelligence (AI), which represents a broad collection of tools that include/leverage ML, such as natural language processing (NLP), self-driving cars, and robotics as well as some that are tangent to ML, such as logical rule-based algorithms. Although there have been multiple "AI Winters" since the 1950s where hype cycles were followed by a dearth of funding, ML is now an enduring innovation driver, in our opinion, due in part to the increasing ubiquity of affordable cloud-based processing power, data storage—we expect the pace of breakthroughs to accelerate.

These innovations will make it easier for companies to benefit beyond what is available with traditional business intelligence (BI) solutions, like IBM/ Cognos and Tableau, that simply describe what happened in the past. New ML solutions are amplifying intelligence while enabling consistent and better-informed reasoning, behaving as thought partners for business users. These endeavors will dramatically improve overall business operations, customer satisfaction, and workforces by eliminating automatable roles (e.g., telemarketer, dishwasher, court clerk) and place more emphasis on social intelligence, creativity, perception, and manipulation.[1] Enterprise SaaS providers (e.g., salesforce.com, Workday, and ServiceNow) have launched first-generation ML functionality to enhance the value of existing applications while legacy software titans (such as Oracle, Microsoft, and SAP), that have been acquiring SaaS companies, are following suit.

While still in the early innings of ML enriched SaaS solutions, we expect the aforementioned companies to dominate horizontal markets (e.g., salesforce automation, customer support, marketing, and human resources) through both organic and M&A initiatives. As for vertical-specific solutions, we think that a new crop of vendors will embrace ML and achieve billion-dollar valuations. VCs are making big bets that these hypotheses will come to fruition as demonstrated by the $5B they poured into 550 startups using AI as a core part of their solution in 2016[2]. Also, established vendors have been acquiring ML solutions, for example, salesforce.com alone has acquired BeyondCore, Metamind, PredictionIO, and RelateIQ.

As the preeminent technology investment bank in the middle corridor of the U.S., Navidar closely tracks trends like these. We are excited about ML because most enterprises have access to a plethora of data and by prioritizing ML initiatives, they will be able to better solve complex, data-rich problems. Companies that are quick to make the arduous business-process and data-management changes, in our opinion, will establish durable competitive advantages. In this article, we examine ML technology, implementation steps, leading vendors' ML initiatives, and industries prime for disruption.

1. 80000hours.org 2. CB Insights

> When I look forward, what I see is an AI first world where every customer is going to be able to get a whole other generation of productivity out of artificial intelligence, ML, and deep learning."
>
> **Marc Benioff, CEO, salesforce.com**

> Machine learning, which I believe is not going to replace jobs, it's just going to make people more effective."
>
> **Aneel Bhusri, CEO, Workday**

> Building out what we refer to as these systems of intelligence. These are these digital feedback loops where you're building systems that help you engage with your customers, empower your employees, optimize your operations and reinvent products and business models."
>
> **Satya Nadella, CEO, Microsoft**

> Machine learning techniques shift the burden from the user asking the right questions to these applications finding the right context-sensitive answers that the user needs to know."
>
> **Amit Ganesh, SVP, Oracle Enterprise Manager**

ML Implementation Steps: Import, Build, Train, Deploy

Current ML approaches are effective at using a single data type to solve narrow problems, such as "which prospects have the greatest probability of closing", broader ones based on multiple data types, like "how to holistically optimize a go-to-market strategy", are not yet solvable because they require numerous micro-decisions—often, each one can use its own ML model. Still, we expect that within the next several years most applications developed will infuse ML. And, while ML will increasingly automate and complement human decisions, mankind will always be involved to an extent in each step. The most common roles will continue to be data scientists, programmers, subject-matter experts, and business/consumer users.

As described below, there are four steps necessary to implement ML as well as advanced approaches to consider.

1 Clean Data is Critical

The amount of digital data is doubling in size every two years and we estimate that over 70% of the data enterprises can rely on to make decisions is unstructured (e.g., mobile devices, social media, and wearables), residing outside of corporate silos and are not widely understood. Internet of Things (IoT), in particular, is increasingly contributing to corporate data growth and ML will play an important role in processing, analyzing, predicting, and responding to embedded sensors and machine-to-machine (M2M) connections. Furthermore, the importance of data is driving M&A, in our opinion, as demonstrated by Microsoft's acquisition of LinkedIn for $26.2B.

Despite access to a plethora of critical data, to implement a successful ML initiative, companies must first identify their objective and secure enough rich training data, otherwise failure is certain. While this step is easier if companies must simply adhere to specifications provided by pre-built models imbedded in SaaS software solutions, the data must still be clean and appropriately labeled. This cleaning process is much more stringent than the requirements of traditional BI because ML requires consistency among users and enterprise workflows as well as the flexibility to dynamically link real-time data. But, the process of labeling data can be a labor-intensive, expensive endeavor. To demonstrate the complexity, it could take a year to appropriately label 10M data records. Fortunately, vendors such as Alegion and CroudFlower provide micro-task, crowd-sourcing, technology-enabled services to help with this arduous task. Their solutions ensure quality work at an appropriate pace through the use of intelligent workflows and by grouping workers based on skillset. The workers they utilize are sourced from services such as Amazon's Mechanical Turk and business process outsourcers (BPOs).

> "The key to any effective AI/ML initiative is quality training data. Alegion uses crowd labor to create structured data that can be customized to meet any project at Enterprise scale."
>
> **Nathaniel Gates, CEO, Alegion**

///ALEGION

2 Build ML Model

Companies can skip the model building step if they are using ML models imbedded in SaaS software ML functionality. But for custom ML initiatives, they must select a framework and approach then actually build a model. The process of ML development is quite different from traditional SaaS application development because ML is created with models (not programs) that are trained (not debugged), retrained (not patched), and pass probability distributions (not absolutes) with dynamic release cycles occurring every millisecond (not quarterly).

ML model approaches are generally categorized into three areas, which are sometimes blended together.

Supervised entails predicting an outcome with a label based on a data training set that has labeled examples which are fed into a model. Supervised is the most common approach (we estimate that 95%-plus of current ML models are supervised) because it is easier to build a business case around as the outcome is known before the project starts.

Unsupervised excels at identifying contextual links that are difficult for humans to replicate. Since the outcome is not fully identified beforehand. It structures data based on inference without being provided correct results. The unsupervised approach is less expensive than supervised because data is not labeled.

Reinforcement discovers what works best through trial and error, often in dynamic environments. Common use cases include robotics, navigation, and gaming.

Deep Learning (DL) and End-to-End Neural Networks (NNs)

DL/NNs excels with unstructured data, such as images, video, audio, geo-location, and time-series data as well as unsupervised approaches. They are utilized by advanced ML models that require significantly more computing and data processing power than the traditional linear probabilistic approach because they simultaneously process multiple micro-models that feed off each other. This approach has just become viable in the past several years due to the improvement of graphics processing units (GPUs), which process data in parallel, much like central processing units (CPUs) perform tasks in parallel. Still, a DL/NN approach for ML can be overkill if a similar result can be achieved through less expensive, standard linear approaches.

Leading ML and DL Frameworks

There are multiple technical backbone options for creating ML models that programming languages (such as Python and C++) can tie into. The most popular framework is Google's TensorFlow. And Pytorch, which Facebook uses, is an up-and-coming alternative. Microsoft (Cognitive Toolkit), Amazon (Deep Scalable Sparse Tensor Network Engine), Intel (BigDL), and Baidu (Paddle) have also developed frameworks.

PYTORCH

TensorFlow

> The resulting level of quality output of any machine learning initiative starts with the input data. Qualia does this by using granular event level Intent data to inform and feed our Identity model. We combine billions of Intent signals with billions of Device signals to create a statistically relevant model that reinforces and validates each other."
>
> **Kathy Leake, CEO Qualia**
>
> *Qualia's adaptive digital identity platform creates undistorted consumer profiles, enables one-to-one consumer engagement and enhances marketing/advertising services for our customers.*

QUALIA

3 Train Model

After ML models are created, humans iteratively train them to achieve desired objectives until sufficient accuracy is achieved. To make this step easier, the aforementioned micro-task crowd-sourcing vendors, can also help with training and scoring of model results.

> The most productive machine learning models are attained through the collaboration of an intelligent system and knowledgeable humans. Conversable combines machine learning with adaptive systems to create experiences that have the power of automation, combined with the satisfaction of contextual relevance."
>
> **Ben Lamm, CEO Conversable**
>
> *Conversable provides an enterprise-class SaaS platform for designing, building, distributing AI-enhanced messaging and voice experiences across multiple platforms, including Facebook messenger, SMS and Amazon Echo*

CONVERSABLE

4 Deploy and Modify Workflows

The trained ML model is then deployed to live data to get predictions and the model's accuracy dynamically improves over time. When ML models are uncertain about particular scenarios, they notify humans that additional training is needed so that it can learn and potentially handle similar situations going forward.

Workflows

Business users must be empowered with specific workflows in this final stage for ML initiatives to be successful. Furthermore, individual business teams need to embrace the new workflows and understand how to test these new capabilities.

ML-Powered Cognitive Technologies

These technologies utilize ML and will increasingly be imbedded in SaaS software solutions, in our opinion.

Natural Language Processing (NLP)	Speech Recognition	Computer Vision
The ability of applications to process and generate text in the same way that humans do. An ML enabled NLP system, for example, is able to process misspelled words by employing probability to determine the intended word. See our **"The Rise of Chatbots! Identifying Winners in the Next Wave of Human-Technology Interaction"** article for more examples of NLP initiatives.	Automatically understands and transcribes human speech into text or another format.	Identifies and interprets objects, scenes, people, and activities in unstructured visual environments, such as images, videos, and virtual reality (VR). Use cases include self-driving cars, robotics, drones, digital check deposits, retinal scanning, and license-plate recognition.

ML Initiatives at Leading SaaS Software Companies in Horizontal Markets

We expect ML will drive a new generation of productivity in horizontal SaaS software areas (e.g., salesforce automation, customer support, marketing, and human resources) where it is critical to understand consumer patterns and manage experiences. While we are in the early innings of such ML enrichment in SaaS software solutions, vendors are beginning to use ML to solve problems at a meta level, removing the complexity about how it is achieved for each customer, and enabling individual users to customize and predict anything they want. By applying ML to data surrounding successful deals, loyal customers, and standout employees, companies can identify winning attributes then encourage and automate ideal behavior. Following is a sampling of ML initiatives at leading SaaS software companies.

> "Right now, consumers are used to asking Siri about the weather tomorrow, but we want to enable people to ask natural questions about their own unique data."
>
> **Richard Socher, Chief Scientist, salesforce.com**

> "One of the things I think is going to be true in machine learning in the next 5 to 10 years is that it's going to be everywhere. It's going to be spread throughout lots of organizations, solving lots of problems."
>
> **Jeff Dean, Senior Fellow, Google**

Google

- **Einstein** – an AI layer infused into all of salesforce's applications. It provides recommendations and alerts.
- **Marketing Cloud Predictive Journeys** – compiles data from multiple sources (e.g., CRM, web browsing activity, email) to predict a customer's likelihood of taking certain actions. It then adjusts their customer journey according to these insights.
- **SalesforceIQ** – uses unstructured data found in emails and calendars to make actionable recommendations, such as encouraging salespeople to email customers and prospects that have not been contacted for a certain amount of time.
- **Field Service Lighting** – when data from IoT sensors indicate that there has been a malfunction (or the risk is high), ML assesses the status of related business processes (e.g., vehicle fleet, locations, and estimated time to arrival) then determines which technician is best suited for the job.
- The company has acquired several ML companies, including BeyondCore ($110M in September 2016), Metamind ($33M in August 2016), PredictionIO ($58M in February 2016), and RelateIQ (for $390M in August 2014).

- **Talent Insights** – analyzes employee data to predict which employees might leave their current role, estimates costs to replace them, and predicts which job functions or departments are at higher risk. When Workday back-tested the solution against internal employee data, it was 97% accurate about who had left.
- The company plans to release 30 to 40 more Insight Application over the next 2 to 3 years.

- **Automatic Answers** – leverages ML to automatically respond to customer-support tickets with relevant articles from Zendesk's knowledgebase.
- **Satisfaction Precision** – assigns satisfaction scores based on hundreds of signals, such as text description, number of replies, total wait time, and open customer tickets/conversations. This functionality enables customer-support personnel to identify tickets that need additional support, enter conversations with more context, and address problems before they escalate.

- **Task Routing** – ServiceNow automatically routes tasks to the customer-support people best equipped to address them. As a result, there is a reduction in their time spend on emails, phone calls, and other manual processes.
- **DXContinuum** – was acquired in January 2017 to build out ServiceNow's ML capabilities.

- **Kemvi** – HubSpot acquired Kemvi on 7/26/17. Kemvi provides AI/machine learning that automates sales peoples' manual processes of finding and using customer data to customize customer interactions (e.g., personalized emails). Its functionality will be integrated into HubSpot CRM. This deal follows other SaaS CRM providers AI initiatives, including salesforce.com's Einstein and Zendesk's Automatic Answers.

- **Human Capital Management Cloud** – recruiting feature uses ML to help HR departments sort through resumes.
- **Application Performance Monitoring** – continuously learns the behavior of each application component, such as a web server request, and detects anomalies in real time.
- **Adaptive Intelligent Apps** – for CX uses ML to enhance its customer experience solutions.
- **Oracle Analytics Cloud** includes domain-specific ML models.
- **Cross-Device ID** – Oracle acquired Crosswire for about $50M in April 2016, which provides ML-based cross-device data and serves to broaden the Oracle ID Graph for a more complete view of consumers' digital interactions across multiple devices.

- **Automated sales forecasts** – predicts which opportunities are most likely to close.
- **Invoice matching** – automatically matches payments to invoices.
- **Resume matching** – identifies the best candidates for a given job description and/or the best job for a given candidate.
- **Additional initiatives include** – image based Ariba Commerce, contextual Concur Travel experience, vision enabled manufacturing, and drone/satellite based asset management for farming.

- **Computational Network and Machine Learning Developer Toolkits** – neural networks and ML for functions such as speech recognition, computer vision, NLP, knowledge management, and search.
- **Intelligent CRM** – enables users to act based on data from inside and outside the system, such as notifying when changes are happening with a customer on LinkedIn.
- **Skype Translator** – translates conversations in real-time across six different languages for audio and 50 for text.
- **Office 365** – understands content in users' inboxes (type of mail, people, context) to prioritize what matters most.
- **Delve** – labor augmentation and enhancement.

- **Adobe Sensei** – is the company's AI and ML framework that is being deployed across all of Adobe's solutions.
- **Intelligent assistant** – helps digital marketers process statistically relevant data and drive more personal and contextual customer relationships.
- **Adobe Target** – analyzes a customer's past interactions and models them against potential future interactions to recommend the most valuable offering.
- **Examples of functionality enabled through Sensei** – optimize geo-fence boundaries and automatically generate push offers, predict customer sentiment and orchestrate the customer journey, voice-centric brand experiences, and recommendations for layout and content personalization.

Alphabet

- **Google** – has an "AI first" policy. The company is applying ML in multiple products (such as Assistant, Photos, Home, and Lens) as well as APIs for developers to enable prediction, speech recognition, translation, and vision functionality.
- **TensorFlow** – open source ML platform which is also used internally in Google Photos and others.
- **Home** – powered by the ML in Google Assistant.
- **Shopping** – advisors and recommendations.
- **SmartReply** – makes automated email responses in Gmail more conversational.

- **Amazon DSSTNE** – open source machine learning platform, similar to Google's TensorFlow.
- **AWS Machine Learning** – ML as a service, enabling developers to build apps with NLP. It converts text into speech, enables conversations with voice or text, analyzes images, and recognizes faces, objects, and scenes.
- **Shopping Recommendation Engine** – has driven 30-50% of daily transaction value.
- **Amazon ECHO** – powered by Alexa, a digital assistant, has built out a voice chatbot ecosystem, called Skills.

IBM

- **IBM Machine Learning Initiatives** – include Watson Virtual Assistant, which is a chatbot; Watson Health features automated discovery; Blue Prism is robotic work/process automation; Watson Workspace is labor augmentation and enhancement offering.
- **Watson Conversation** – is priced at $0.0025 per API query per month.

- **Apple** – is incorporating ML into several applications, including Siri, Photos, Messages, and Maps to understand natural language and context as well as to make suggestions.
- **CoreML and ARKit Developer Frameworks** – CoreML enables ML to be built into apps while ARKit is for augmented reality.
- **Siri Watch Face** – incorporates the intelligent proactive assistant in iWatches to automatically display data that is most relevant to users.

- **Audience Insights** – a component in FB Analytics that uses ML to automate the output of campaign datasets and analysis for marketer-specific customization.
- **AR Glasses** – uses ML to feed users useful information directly into their eyes or ears.
- **Facebook Messenger Chatbot Platform** – more that 34k chatbots have been created since April 2016.

- **Twitter** – uses ML to identify, personalize, and notify users about content they want to see while delivering it to them more quickly.
- **Magic Pony Technology** – acquired in June 2016 for $150M due to its ML image and video processing functionality.

Vertical SaaS Software

Domain knowledge is critical in vertical ML applications and some industries (such as retail, healthcare, financial services, and manufacturing) tend to respond more quickly and aggressively than others to shifts in technology. There are elements of cyber-security imbedded in many vertical solutions. Following we provide examples of initiatives as well as opportunities in several industries that are ripe for disruption.

Ecommerce/ Retail

- This sector has been an early adopter of ML for inventory planning and recommendation engines: "customers who purchased this item also bought." It combines an item being viewed with order history and the available demographic attributes of the user like age, sex, and location to create a set of targeted recommendations they are most likely to purchase. It then recommends the right products to customers at the right moment.
- **IoT is driving a swell in retail data** – driven by RFID tags, shelf beacons, smart hangers, and location-sensing WiFi. In fact, the amount of data available to retailers doubles every two years, and 90% of retail data is less than two years old.
- **Amazon,** for example, generates more than half of sales from its recommendation engine based on what's in users' baskets, their past searches, and what other shoppers have purchased in similar situations.

Healthcare

- Only 20% of the knowledge that clinicians use today is evidence-based and keeping up with new medical knowledge could require as much as 160 hours of study per week, according to Sloan-Kettering.
- **Electronic Health Record (EHR) Systems** – are accelerating data collection in healthcare. There is potential for ML to detect diseases based on symptoms, lab results, and X-rays/CT-scans. We also expect it to drive proactive health management by correctly identifying patients' symptoms in order to recommend the appropriate medications.
- **IBM Watson for Oncology,** for example, has a 90% successful rate for detecting lung cancer, compared with 50% for human doctors. Also, its Health Cloud aggregates personal, clinical, research and social health data from clients, partners, and researchers to surface new connections.
- **ML is driving M&A activity in healthcare.** Speaking of its June 2017 acquisition of Praxity for $63M, Athenahealth's CTO Prakash Khot said, "leveraging advancements in machine learning and natural language processing, Praxify has invested several years in developing highly intuitive technology to enhance the delivery of high quality patient care."

Financial Services

- The volume of data captured in financial services is becoming too much for traditional relational database systems.
- **Risk Calculations** typically occur overnight in batches, which limits quick responses and results in lost revenue.
- **There are myriad of applications in the financial services industry,** including predicting fraud and evaluating credit quality in retail finance as well as automating investment processes and algorithmic trading in investment management.
- **Q2 Ebanking's SMART SaaS solution** utilizes ML to unlock actionable customer insights to drive improved targeting, products, growth and revenue across channels. And its *Patrol* offering utilizes ML to identify potentially fraudulent digital banking sessions.

Manufacturing

- **Manufacturing** – Internet of Things (IoT) is driving demand for ML within the broad manufacturing sector, including defense, aerospace, automotive, and energy. Use cases include estimating warranty reserves, forecasting demand, optimizing process and predicting maintenance.
- **Energy** – oil, gas, and utilities are embracing ML. An example is optimizing energy supply/demand and intelligent power grid management.

"Industrial Internet is responsible for approximately one-third of the world's data today. The sensors and operations of the industrial world are producing volumes of data that humans cannot make sense of. The only way to extract efficiencies and improve safety is through smart software.

The side effect of Industrial IoT is susceptibility to cyber attacks. The notion of cyber risk is not limited to credit card or customer data, but the risk of turning industrial assets into vehicles for physical destruction. The complexity of industrial assets and increase in sophistication of cyber malware make this a complex challenge—one well suited for Artificial Intelligence and something SparkCognition is taking head-on."

Usman K. Shuja, SparkCognition, General Manager, IIoT

SparkCognition's AI products are improving throughput, efficiency, and safety for many industrial clients across manufacturing, aerospace, and energy. For instance, our products are not only great at identifying patterns to predict failures but also performing root-cause analysis and explaining reasons for predicting to reduce time to remediate. We are spending significant research into building systems that explain what they did, which we call "Explainable AI". The idea is to make AI systems better communicate with human experts and gain their trust so they can take necessary actions.

sparkcognition

ML Startups

To compete with the scale of companies such as salesforce.com, Google, and Microsoft, some ML startups target specific workflow areas within horizontal solutions and others focus on specific industries. The hard costs related to a ML startup are very low given the proliferation and commoditization of cloud computing power and storage. And proven entrepreneurs are able to raise capital and highly-knowledgeable technicians are able to hire themselves inexpensively. But, the cost of sourcing and cleaning data as well as securing PhD level data scientists can be prohibitive. Two ML-related startups described the process as follows.

"The challenges of building a machine-learning startup are finding the right problem and uniting a group of experts around that common purpose. Computing power is cheap, but well-trained brains are expensive and rare. Fortunately, the advantage of startups over big companies is that they have more freedom to choose the problem because they have less to lose."

Jonathan Mugan, CEO, Deep Grammar

Deep Grammar uses ML to power natural language processing as a grammar checker.

"As a ML startup, we use cloud storage as an agile commodity to iteratively refine large, innovative ML knowledge structures and models of care. Our advantage as a startup is that we are focusing on the cutting-edge of ML-based care delivery whereas larger players are more focused on diagnostics."

Andy Liwen, co-CEO, Aplacare

Aplacare meets the challenge of Value-Based Care by using AI/ML to deliver intelligent patient data at the point of care, empowering healthcare providers to innovate their clinical operations resulting in an improvement in quality, inefficiencies, and optimize financial outcomes.

APLACARE

AI Winter is Over

Hype cycles followed by a dearth in funding, also known as “winters”, are common in emerging technology, exemplified by the dot-com crash. AI, in particular, experienced two meaningful hype cycles in the 1970s and the 1980s. These crises were caused by unrealistic predictions and insufficient data as well as inept algorithms, processing, and storage power. Modern circumstances also different because of the increasing ubiquity of affordable cloud-based processing power, data, and storage as well as improving algorithms and hardware (GPUs). While there will always be the risk of unrealistic predictions, there is little chance for another AI Winter. ML is established as an enduring innovation driver and expect the pace of breakthroughs to accelerate.

In Summary

- Machine learning is driving the next wave of SaaS software solutions. The underlying technology has been around for years, but has made significant advances in the past decade due to the increasing ubiquity of affordable cloud-based processing power and data storage.
- ML-based innovations will make it easier for companies to benefit beyond what is available with traditional BI solutions that simply describe what happened in the past.
- The steps to developing ML solutions are simplified when using those imbedded in SaaS software solutions, but even for more complicated custom solutions, it all begins with clean data. The next steps are designing the ML model, training, and deploying.
- While we are in the early innings of ML being imbedded in SaaS software solutions, we expect enterprise SaaS leaders (salesforce, Workday, and ServiceNow) and legacy titans (Oracle, Microsoft, and SAP), that have been acquiring SaaS companies, to dominate horizontal markets (e.g., salesforce automation, customer support, marketing, and human resources) through organic and M&A ML initiatives.
- VC dollars will continue to flow into AI and ML oriented solutions, in our opinion, not only in horizontal markets but also in specific verticals applications.
- Several industries (such as ecommerce/retail, healthcare, financial services, and manufacturing) are ripe to be disrupted by ML-based startups that will eventually garner billion-dollar valuations, in our opinion.

About Navidar

Navidar LLC is a Registered Investment Advisor and a subsidiary of Navidar Holdco LLC (which was referred to as “Navidar” in this report). For more information, please visit **www.Navidar.com**.

Exhibit 1 – Notable Machine Learning and Artificial Intelligence Acquisitions

Buyer	Target	Price	Revenue Multiple	Date
Amazon	Harvest.ai	$20M		1/9/2017
Amazon	Angel.ai			9/20/2016
Amazon	Orbeus			10/1/2015
Amazon	Amiato			5/31/2014
AOL	Sociocast Networks			6/9/2015
AOL	Convertro	$101M	19.0x	5/6/2014
AOL	Gravity	$90.7M		1/23/2014
Apple	Lattice	$200M		5/15/2017
Apple	RealFace	$2M		2/20/2017
Apple	Tuplejump Software			9/22/2016
Apple	Turi	$200M		8/5/2016
Apple	Emotient			1/7/2016
Apple	Perceptio			10/5/2015
Apple	Vocal IQ			10/2/2015
Apple	Acunu Ltd.			3/25/2015
Apple	Novauris Technologies			4/3/2014
Baidu	KITT.AI			7/5/2017
Baidu	xPerception			4/13/2017
Baidu	Beijing Raven Tech			2/15/2017
Cisco Systems	MindMeld	$125M		5/11/2017
Cisco Systems	ParStream GmbH			10/26/2015
Cloudera	Sense			3/22/2016
Ebay	SalesPredict SW	$40M		7/11/2016
Ebay	Expertmaker			5/5/2016
Etsy	Blackbird Technologies	$32.5M		9/19/2016
Facebook	Ozlo			8/1/2017
Facebook	Zurich Eye			11/10/2016
Facebook	Masquerade Technologies			3/9/2016
Facebook	Wit.ai			1/5/2015
Facebook	Mobile Technologies			8/13/2013
Facebook	Face.com	$60M		5/29/2012
Ford Motor Company	Argo AI	$1B		2/10/2017
Ford Motor Company	SAIPS			8/16/2016
Fortinet	AccelOps	$32M	5.08x	6/7/2016
General Electric	Wise.io			11/15/2016
General Electric	Bit Stew Systems	$153M		11/15/2016
Genpact Ltd.	Endeavor Software	$14M		9/30/2015
Google	Halli Labs			7/12/2017
Google	Kaggle			3/8/2017
Google	Api.ai			9/19/2016
Google	Apigee Corp.	$625M	6.1x	9/8/2016
Google	Moodstocks			7/6/2016
Google	Timeful			5/4/2015
Google	Granata Decision Systems			1/23/2015
Google	Vision Factory			10/23/2014
Google	Dark Blue Labs			10/23/2014
Google	Jetpac			8/16/2014
Google	Emu			8/6/2014
Google	DeepMind Technologies	$500M		1/27/2014
Google	DNN Research			3/12/2013
HubSpot	Kemvi			7/25/2017
IBM	Explorys			4/13/2015
IBM	AlchemyAPI			3/4/2015
IBM	Cognea			5/20/2014

Buyer	Target	Price	Revenue Multiple	Date
Infor	Birst	$100M	1.2x	4/25/2017
Intel	Movidius	$400M		9/6/2016
Intel	Nervana Systems	$400M		8/9/2016
Intel	Itseez			5/27/2016
Intel	Saffron Technology			10/26/2015
Intel	Indisys			9/13/2013
Intuit	Level Up Analytics			10/23/2013
LogMeIn	Nanorep	$50M		8/1/2017
Meltwater Group	Wrapidity			2/21/2017
Meltwater Group	Encore Alert			3/29/2016
Meltwater Group	OCULUSai			3/18/2013
Microsoft	Maluuba			1/13/2017
Microsoft	Genee			8/22/2016
Microsoft	SwiftKey	$250M	31.7x	2/19/2016
Microsoft	Volometrix	$250M		9/3/2015
Microsoft	Equivio	$200M		1/20/2015
Microsoft	Netbreeze			3/20/2013
NetApp	Plexistor	$32M		5/24/2017
NICE Systems	Nexidia	$135M	4.0x	1/11/2016
NICE Systems	Causata			8/7/2013
Nokia	Medio Systems			6/12/2014
Nokia	Desti			5/30/2014
Oracle	Palerra			9/18/2016
Oracle	Crosswise	$50M		4/14/2016
Oracle	Blue Kai			2/24/2014
PROS	SignalDemand	$13.5M		12/16/2013
salesforce.com	BeyondCore	$113M		9/1/2016
salesforce.com	MetaMind	$48M		4/4/2016
salesforce.com	PredictionIO			2/19/2016
salesforce.com	Tempo AI			5/29/2015
salesforce.com	RelateIQ	$390M		8/12/2014
ServiceNow	DxContinuum			1/18/2017
Software AG	Zementis			12/2/2016
Sophos	Invincea	$120M	12.2x	2/8/2017
Sophos	Barricade.io			11/2/2016
Splunk	Cloudmeter	$21M		12/10/2013
Spotify	Niland			5/17/2017
Teradata Corp.	StackIQ			7/13/2017
Teradata Corp.	Big Data Partnership			7/25/2016
Trimble	Innovative Software Engineering			6/13/2017
Twitter	Magic Pony Technology	$150M		6/20/2016
Twitter	Whetlab			6/17/2015
Twitter	TellApart	$479M		4/28/2015
Twitter	Madbits			7/30/2014
Twitter	Marakana			8/13/2013
Twitter	Lucky Sort			5/13/2013
Uber Technologies	Geometric Intelligence			12/5/2016
Uber Technologies	Otto	$680M		8/18/2016
Vmware	Cetas Software			4/24/2012
Workday	Platfora			7/21/2016
Yahoo!	Incredible Labs			1/30/2014
Yahoo!	SkyPhrase			12/2/2013
Yahoo!	LookFlow			10/23/2013
Yahoo!	IQ Engines			8/25/2013

Disclaimer

Certain statements in this Report (the "Report") may be "Forward-looking" in that they do not discuss historical facts but instead note future expectations, projections, intentions, or other items relating to the future. We caution you to be aware of the speculative nature of forward-looking statements as these statements are not guarantees of performance or results.

Forward-looking statements, which are generally prefaced by the words "may," "anticipate," "estimate," "could," "should," "would," "expect," "believe," "will," "plan," "project," "intend" and similar terms, are subject to known and unknown risks, uncertainties and other facts that may cause actual results or performance to differ materially from those contemplated by the forward-looking statements.

We undertake no obligation to update or revise any forward-looking statements, whether as a result of new information, future events or otherwise, except as required by law. In light of these risks, uncertainties, and assumptions, the forward-looking events discussed might not occur.

This Report has been prepared solely for informational purposes and may not be used or relied upon for any purpose other than as specifically contemplated by a written agreement with us.

This Report is not intended to provide the sole basis for evaluating, and should not be considered a recommendation with respect to, any transaction or other matter. This Report does not constitute an offer, or the solicitation of an offer, to buy or sell any securities or other financial product, to participate in any transaction or to provide any investment banking or other services, and should not be deemed to be a commitment or undertaking of any kind on the part of Navidar Holdco LLC ("Navidar") or any of its affiliates to underwrite, place or purchase any securities or to provide any debt or equity financing or to participate in any transaction, or a recommendation to buy or sell any securities, to make any investment or to participate in any transaction or trading strategy.

Although the information contained in this Report has been obtained or compiled from sources deemed reliable, neither Navidar nor any of the Company affiliates make any representation or warranty, express or implied, as to the accuracy or completeness of the information contained herein and nothing contained herein is, or shall be relied upon as, a promise or representation whether as to the past, present or future performance. The information set forth herein may include estimates and / or involve significant elements of subjective judgment and analysis. No representations are made as to the accuracy of such estimates or that all assumptions relating to such estimates have been considered or stated or that such estimates will be realized. The information contained herein does not purport to contain all of the information that may be required to evaluate a participation in any transaction and any recipient hereof should conduct its own independent analysis of the data referred to herein. We assume no obligation to update or otherwise revise these materials.

Navidar does and seeks to do business with companies covered in Navidar Research. As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of Navidar Research.

Navidar and its affiliates do not provide legal, tax or accounting advice. Prior to making any investment or participating in any transaction, you should consult, to the extent necessary, your own independent legal, tax, accounting and other professional advisors to ensure that any transaction or investment is suitable for you in the light of your financial capacity and objectives.