

Real potential for AI

Analysis of the artificial intelligence & machine learning market

PitchBook is a Morningstar company. Comprehensive, accurate and hard-to-find data for professionals doing business in the private markets.

Credits & Contact

Analysts

CAMERON STANFILL Analyst
cameron.stanfill@pitchbook.com

Contact PitchBook

pitchbook.com

RESEARCH

reports@pitchbook.com

EDITORIAL

editorial@pitchbook.com

Key Takeaways

- As a general-purpose technology, artificial intelligence (AI) and machine learning (ML) have potential use cases in virtually every industry and the ability to reshape the way people live and do business. Breakthroughs in deep learning in the past decade have engendered a proliferation of artificial intelligence applications into daily life and paved the way for further advancement in the field.
- VC investment in the vertical is on an extended growth trend to levels 12x above what we saw in 2008. 2017 recorded \$6 billion invested across 643 VC deals in AI/ML. Similarly, after years of negligible exit activity, the last two years represented a substantial uptick in liquidity and a shift to a new stage of the AI/ML exit environment.
- For now, almost all commercially successful ML applications use supervised learning, which encompasses a vast number of applications but is limited to areas that have clean, labeled data. Startups will face stout competition from low-cost options available via the cloud from technology giants, but they can excel by focusing on more niche areas or datasets.

Contents

Key Takeaways	1
Introduction	2-3
VC Trends	4-5
Exits	6
Limitations	7
Outlook	8

Introduction

There has been a great deal of excitement surrounding the AI/ML field over the past decade, with rapid advancements and significant commercial integration of applications. The concept of AI as commonly recognized today has been on scientists' minds since the 1950s, which saw experimentations with building an artificial brain and the "Turing Test." Several periods of stagnation have occurred within the field since, however. These so-called "AI winters" often arise due to limitations in technology, infrastructure or methodology. The most recent breakthroughs in AI/ML began around 2006 with advancements in deep learning, a subset of ML that is based on layered neural networks (NNs) inspired by the biology of the human brain. Deep learning techniques are behind nearly all the current frontier research and successful applications of AI/ML (e.g. image recognition, search engines, drug discovery, deep reinforcement learning). This progress was made possible by the new research approaches to ML problems, combined with huge advances in computing power and the rampant increase in data digitization and availability.

Artificial intelligence (AI) is the area of computer science that focuses on the creation of an intelligent machine that can perceive its environment and make decisions to maximize the chances of reaching its goal. Machine learning (ML) is a subfield of AI and data analysis that is working to give computers the ability to learn iteratively, improve predictive models and find insights from the data without being explicitly programmed. As individuals, we already interact with AI/ML applications every day when we talk to a voice assistant, use facial recognition technology, receive movie or restaurant recommendations, and numerous other instances. ML can be broken into the three subcategories based on the method in which the algorithm is trained:

1. Supervised learning

Supervised learning is a method in which the machine learning algorithm is trained on a set of labeled data (e.g. training an image recognition system with labeled pictures to identify pictures of dogs).

2. Unsupervised learning

In unsupervised learning, the algorithm trains with unlabeled data and must determine the underlying structure of the dataset and how to group it (e.g. recognizing handwritten numeric digits by classifying them into 10 groups).

3. Reinforcement learning

Reinforcement learning starts with no training data, which means the machine must learn through experience and iterative trial-and-error to perform its task while maximizing long-term reward.

In the near term, we think there will continue to be a proliferation of supervised learning applications in AI/ML developed and refined from a consumer-facing approach to focus on the automation of tasks, almost to the point of ubiquity.

In the near term, we think there will continue to be a proliferation of supervised learning applications in AI/ML developed and refined from a consumer-facing approach to focus on the automation of tasks, almost to the point of ubiquity. Over the next year, these companies will produce the most commercially viable AI/ML products and will represent many acquisitions in the vertical. However, we believe the largest AI/ML market prospects will derive from a potential expansion into a host of enterprise applications and new industries.

Opportunities in unsupervised and reinforcement learning present some exciting scenarios at the frontier of the vertical. While there are still a fair number of obstacles to proliferation of these techniques, the possibilities could offer solutions to an increasingly large array of problems. One of the most significant early achievements in reinforcement learning was AlphaZero, originating from Google's DeepMind. This algorithm mastered the games of chess, Go and Shogi purely via simulated games against itself. Further potential use cases where reinforcement learning can improve on more traditional solutions include resource allocation problems, a plethora of personalized user interfaces (healthcare, content such as headlines, ads, etc.), and applications in robotics and autonomous vehicles. While there are hurdles to bringing these techniques into commercial products, proof of concept at the bleeding edge of AI/ML can attract significant investment. This has been demonstrated by the amount of capital flowing into the autonomous vehicle space. In the end, the most adaptable VC-backed companies and investors will succeed in the AI-centric world.

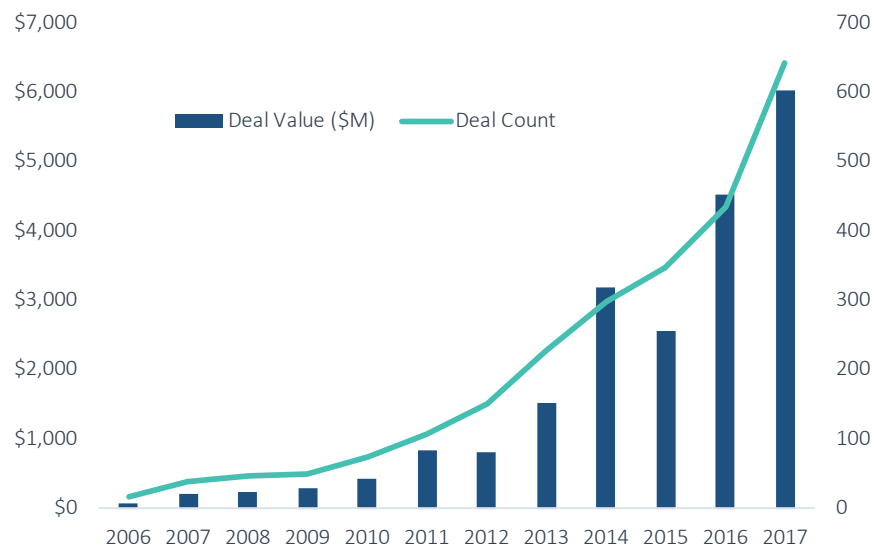
That said, many of these applications will take more time to come to fruition due to the complexity of the problems and the computation power that they require. In response to these issues, advances are being pursued in supporting hardware—quantum/high performance computing and hybrid computing (GPUs/CPUs/FPGAs/TPUs)—as well as decomposition techniques to break down complex problems into manageable segments.

VC Trends

Venture investment in the AI/ML space tells a clear story about the advancement and increasing interest in the vertical over the last decade. 2017 has recorded \$6 billion invested across 643 VC deals in AI/ML.

Deal counts have swelled at a nearly exponential pace to levels 12x above what we saw in 2008.


US venture activity in AI/ML



Source: PitchBook

Paralleling contemporary technological breakthroughs, deal counts have swelled at a nearly exponential pace to levels 12x above what we saw in 2008. Deal counts across all stages have pushed higher, but the vast majority of growth has come through angel & seed and early-stage deals. While mainstream media exposure of AI/ML has ratcheted up recently, it's clear that the vertical is still in the early innings. Even still there is a segmentation in the market. Some larger companies are already executing on commercial products—online lender Avant and used goods marketplace letgo, for example—while many others are working to refine tools as well as expanding their applications and product offerings. Theoretical and practical research have made huge steps in proving functionality, but we are still in the initial period of determining feasibility and use cases for many businesses. Furthermore, due to the uniqueness of each dataset, execution of AI and ML techniques in a business usually requires a high level of customization that hinders early widespread adoption.


US venture activity in AI/ML by stage



Source: PitchBook

AI companies have exhibited rapid valuation growth similar to that of the broader VC market. The exception is the late stage, where medians have flattened out after a pop in 2013 and 2014, unlike the exceptional growth in 2016 and 2017 experienced by late-stage VC companies in aggregate. One reason for this anomaly is that a host of large AI/ML companies raised funds in 2013 and 2014, including multiple rounds raised by Palantir, and have yet to return for sizable follow-on rounds. Additionally, the vertical is still dominated by earlier-stage companies, which tempers the median numbers. This is a bit surprising given the acceleration of demand for AI/ML deals in the last couple years, although valuations at minimum doubling over 10 years isn't tepid growth.

Median post-valuations of AI/ML companies by stage in US



Source: PitchBook

MOST ACTIVE ACQUIRERS OF US VC-BACKED AI/ML COMPANIES


FIRM	ACQUISITIONS
Alphabet	5
Microsoft	4
Intel	3
Apple	3
Yahoo	3
Amazon	2
Facebook	2
Salesforce	2
Samsung	2
Airbnb	2
HubSpot	2
IBM	2
LexisNexis	2
Nuance	2
Splunk	2
Xilinx	2

Source: PitchBook
 *Data as of 2006 to 12/31/2017

Exits

After years of negligible exit activity, the last two years represent a shift to a new period in the AI/ML exit environment. The relative dearth of AI/ML exits until recently is part of the typical VC cycle, as VC investment usually flows into a vertical as it begins to grow, and it subsequently takes years for businesses to mature and for strategic acquirers to enter. This is key because AI/ML exits are currently occurring almost exclusively via strategic acquisitions. We expect this trend to continue as current tech giants work to bolster their in-house AI offerings. Additional acquisitions will be fueled as corporations operating outside the software/internet sphere realize the potential of integrating AI into their business.

US venture-backed exit activity in AI/ML



Source: PitchBook

The largest recorded exits in both 2016 and 2017 came in the autonomous driving space, with GM's purchase of Cruise Automation for \$1 billion and Aptiv's acquisition of nuTonomy. It's also important to point out that the Cruise acquisition represented nearly 40% of recorded AI/ML exit value for 2016, an important reminder of the effect of outliers within the AI/ML market as it stands.

Much like no one had the job title of “Social Media Manager” or “YouTuber” before the proliferation of internet applications, the advent of AI/ML will spawn previously unconceived industries and occupations.

Limitations

Within the vertical, the dominant sentiment waxes positive; however, there are limitations that could slow further penetration of AI/ML across all industries. The effect that AI/ML will have on our lives is often accompanied by grandiose claims that overpromise short-term viability, with the implicit assumption that failure is impossible. This mindset can be dangerous, as realistic expectations of ramp-up timeline and effective failure management should be an integral part of implementing AI/ML techniques in enterprise. For instance, accidents or injuries involving an autonomous vehicle or other algorithmic mistakes can damage public perception of the technology and cause an array of other problems including delays in implementation.

Another common issue with AI/ML occurs when a machine’s actions can’t be explained by humans—a situation known as the “black box” problem—which obscures auditing and the culpability for the actions of the machine.

AI/ML should ease the burden of both menial and difficult tasks, but there are broader societal impacts that may accompany broad adoption. One of the most covered is the idea that AI is going to take jobs from humans. While this may end up being true for some occupations, advancements will likely change the nature of potential jobs, as opposed to eliminating them entirely. Much like no one had the job title of “Social Media Manager” or “YouTuber” before the proliferation of internet applications, the advent of AI/ML will spawn previously unconceived industries and occupations.

Key Opportunities

While AI/ML has the potential to effect almost every industry, we see a few as ripe for extensive transformation or disruption. One prime opportunity for AI/ML applications is healthcare, which can provide a great deal of value and more importantly save lives. There are three basic areas where these technologies will be employed. Firstly, the automation of menial tasks for doctors and other healthcare workers—especially the tedium of manually inputting patient data into electronic health records (EHRs). New drug discovery is also an exciting use of AI/ML technology—from treating previously incurable diseases, to testing the potential outcomes of gene manipulation in CRISPR treatments. Finally, the ability to provide personalized healthcare for every patient will improve patient experience as well as solve a myriad of problems including drastically decreasing the thousands of deaths every year caused by adverse reactions to medication.

Outlook

As a general-purpose technology, AI/ML have potential use cases in virtually every industry and the ability to reshape the way people live and do business. For this reason, AI has been referred to as “the new electricity” by deep learning pioneer Andrew Ng because of its potential to revolutionize business and our daily lives, similar to the internet or database technology. Although AI/ML companies and investment has proliferated, many large markets are still waiting to be explored.

For now, almost all commercially successful ML applications use supervised learning, which addresses a large market but is limited to areas that have clean, labeled data. Startups face stout competition from low-cost options available via the cloud from behemoths like Google, Amazon, Microsoft, Baidu and Salesforce that have access to data on a massive scale and highly developed distribution channels. This is especially important for brute force applications (e.g. speech/object/facial/image recognition and translation) that thrive off large datasets. As a result, AI/ML startups focusing on fringe or specific areas or datasets may be the best positioned to compete and thrive. For example, recent advances in data digitization of many medical records and scans make healthcare especially well positioned to see further penetration by AI/ML applications. Even with the resources available today, however, many potential uses of AI/ML have been hamstrung by a lack of actionable learning data.

To review PitchBook's Artificial Intelligence & Machine Learning Market Map, [click here](#).