Artificial Intelligence Startups:

The Key to Real Returns in the Modern Economy?

November 9, 2017

EQUITYZEN

AI Startups | Executive Summary

AI (finally) ready for takeoff. We believe the ingredients for sustained progress in artificial intelligence are now in place: (1) Machine learning (ML), the most common AI technique today, has led to more promising breakthroughs vs. prior methods. (2) Data - the fuel of machine learning - is abundant with an estimated 2.5 quintillion bytes created daily. (3) Computing power and storage costs have dropped dramatically, enabling efficient data processing and analysis. (4) Capital availability has increased, with over $12B of venture capital invested in AI since 2011. Corporate players (notably Google, Amazon, Facebook, Baidu and Microsoft) have also been active through R&D and M&A.

The next technological revolution? Historically, technological revolutions have been triggered by a vital economic input becoming cheaper. AI reduces the cost of prediction -- a key input for business decisions across many economic sectors. Given AI's broad applicability, the potential impact of the technology is staggering, with two prominent consultancy firms pegging the eventual economic contribution at ~$14-16T.

Startups in AI. Crunchbase data shows 1,000 AI startups in operation currently. Most of these startups focus on technologies in specific verticals (autos, fintech, healthcare). A smaller number specialize in Core AI and application-specific AI, which are industry-agnostic.

In our view, AI will be pervasive in the modern economy and we believe investors will benefit from gaining exposure to this theme. Key positives for AI startups include significant market potential as well as strong M&A interest from corporates. We note that returns may take time to materialize, as most AI companies are still early-stage. Other risks include: challenges in developing competitive advantages against larger rivals, legal/ethical concerns and limited transparency (even by pre-IPO standards).

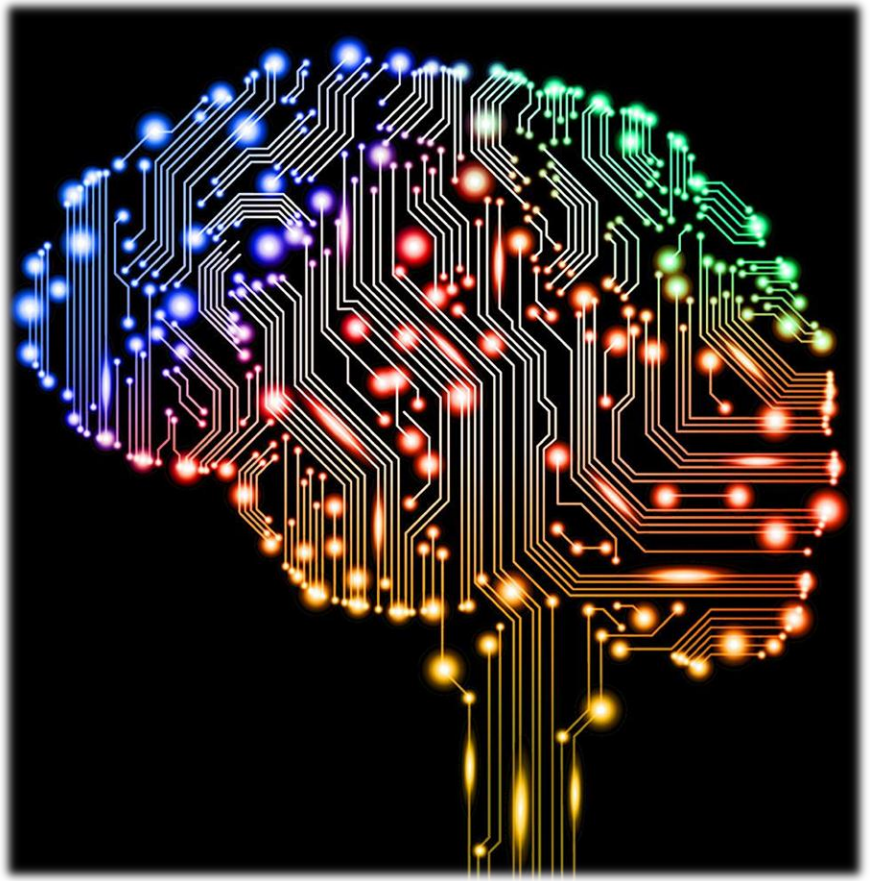

Some AI Basics

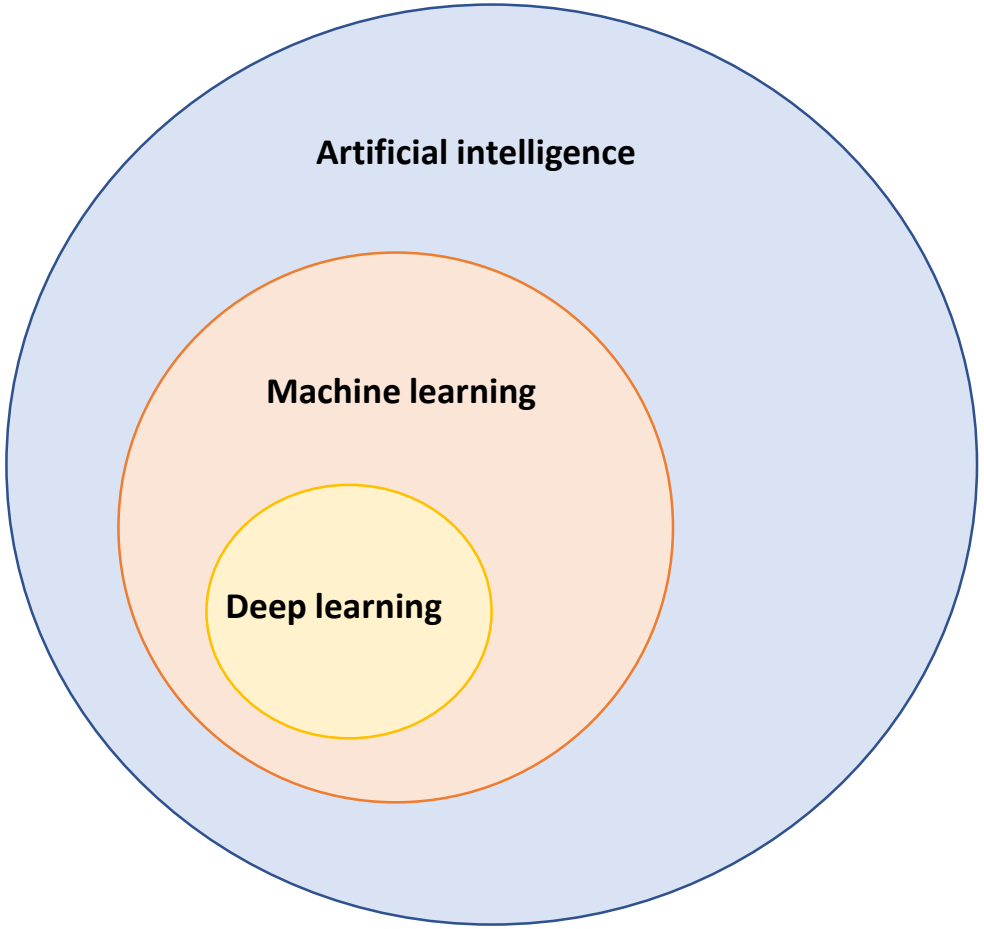

Artificial intelligence: Computers that can perform tasks that typically require human intelligence. There are two types of artificial intelligence – narrow and general. Narrow refers to non-sentient machines that are trained for a specific task (e.g. driving, text analytics, conversation). General AI, sometimes called the "holy grail" of AI, refers to a machine that can perform any task that a human can. The vast majority of applications of AI today are narrow AI.

Most AI research today is centered around Machine Learning and Deep Learning.

- **Machine learning:** Machine learning is a computationally heavy, data-driven method of training a computer so it can make inferences and predictions.
- **Deep learning:** A subset of machine learning that employs multi-layered neural networks to process information. This has become an increasingly important focus for AI in recent years.

AI (finally) ready for take-off

Breaking through stall speed. The short ~60 year history of AI is filled with booms and busts. In our view, the ingredients are finally in place for sustained progress in the field:

- *Machine learning.* Historical AI "busts" were often caused by techniques that failed to achieve promised breakthroughs. Today's AI research is generally focused on machine learning and deep learning. These techniques have shown more promise, powering key innovations including self-driving cars, speech recognition, and effective web search.
- *Abundant data.* Data is the fuel of machine and deep learning. The internet and increasing number of connected devices have created ever deepening pools of data – IBM estimates that 2.5 quintillion bytes of data are created daily.
- *Dramatically cheaper compute and storage.* Given its large data requirements, AI requires significant computing power. Recent studies suggest that over the past quarter century, computing power available per dollar has increased by a factor of ten every four years.
- *Strong capital availability.* Venture capitalists have poured over $12B into the sector since 2011 (see page 6). Corporates (notably Google, Amazon, Facebook, Baidu and Microsoft) have also been active in the space, completing over 250 acquisitions of AI startups over the past five years.

Computing costs over time

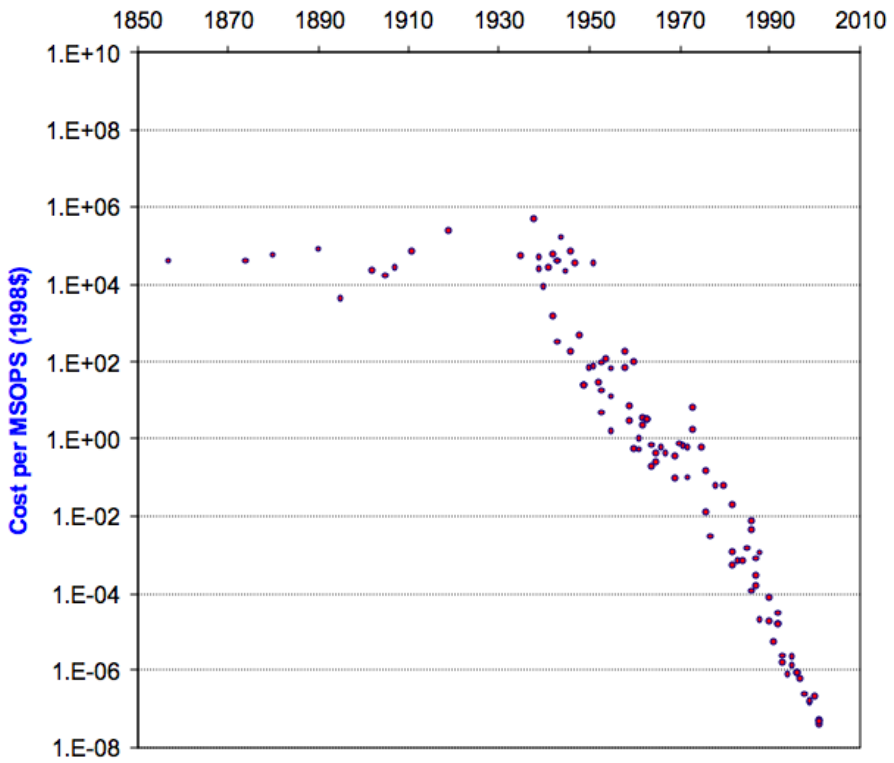

Source: Aiimpacts.com

The next technological revolution?

"Just as 100 years ago electricity transformed industry after industry, AI will now do the same"
- Andrew Ng (Machine learning pioneer, Chief Scientist at Baidu & Adjunct Professor at Stanford University)

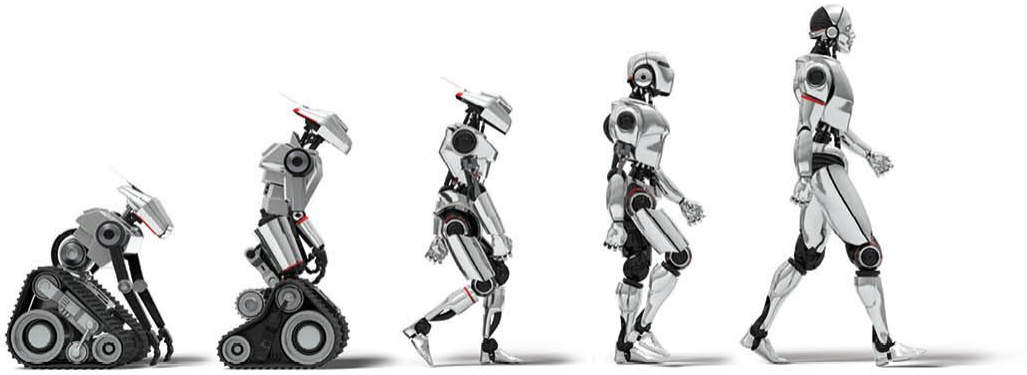

In our view, AI will be a transformative technology that will become increasingly pervasive in software applications across industries. Technological revolutions throughout history were sparked by a key input becoming cheaper – e.g. cheaper energy for the industrial revolution and cheaper communication for the information revolution. AI reduces the cost of prediction, which is a key input for business decisions across economic sectors. In addition to transforming existing business practices, cheaper prediction can also foster the use of the input in new ways – e.g. autonomous vehicles make decisions by predicting how humans behave behind the wheel.

We believe we are still in the early innings of the AI revolution. Incremental progress in AI is likely to lead to accelerating growth in the field, in our view. For example, some artificial intelligence projects are becoming sophisticated enough to write additional AI software, building on current knowledge and practice.

AI Startup Landscape

According to Crunchbase, over 1,000 artificial intelligence startups are operating currently. Unicorns (companies with valuations >$1B) include: Zoox, UBTECH, BenevolentAI, iCarbonX and insidesales.com. Most startups focus on narrow AI applied to specific industries and/or verticals (autos, fintech, healthcare). A smaller number specialize in the industry-agnostic Core AI and application specific AI segments.

Over $12B of venture capital has poured into the sector since 2011. Data Collective, New Enterprise Associates and Khosla Ventures are among the top VC backers. In addition, several corporates have raised AI dedicated funds, including Google (Gradient Ventures), Toyota and Salesforce.

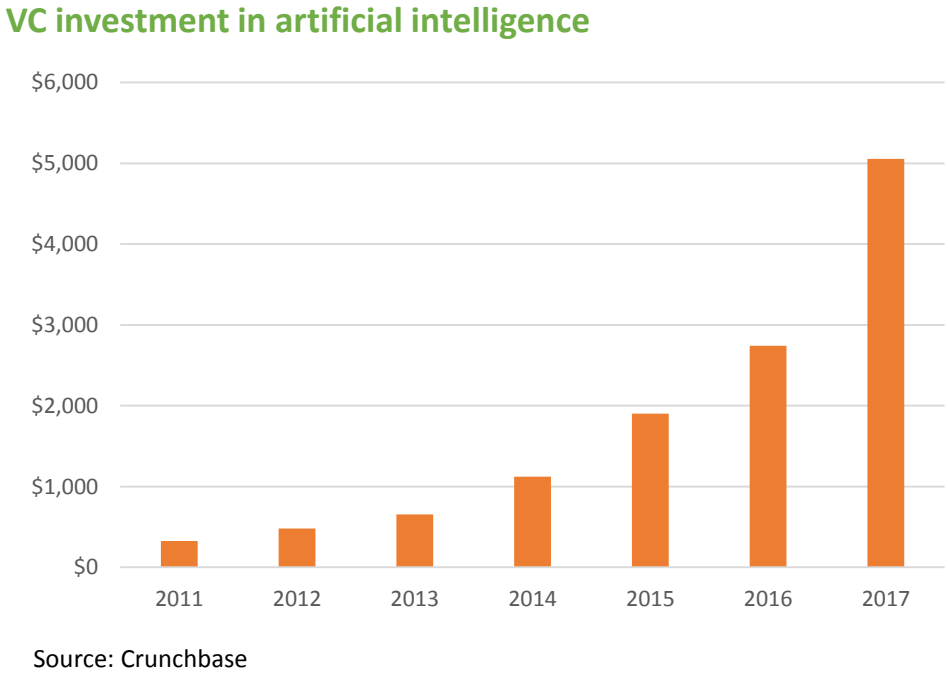

Source: Crunchbase

Investment Themes for AI Startups

Key Positives

> Transformative technology with large commercial opportunity

> M&A demand high from large corporates

Key Risks

> Nascent market with limited visibility; many companies still pre-revenue

> Difficult for small firms to gain competitive advantage

> Regulatory/ethical concerns could delay timetables for AI deployment

> Transformative technology with large commercial opportunity

Potential impact of AI is staggering. A recent study by PWC estimates AI will contribute $15.7 trillion to global GDP by 2030. Accenture, another consultancy, offers a slightly lower $14T forecast (by 2035). At these levels, we estimate AI would have a larger impact on GDP growth vs ICT (information and communication technology), which started the so-called "Information Revolution". Studies have pegged ICT's contribution to economic growth at ~0.2-0.6% annually.

"Data is the new oil". We believe AI's contribution to economic growth will be driven by three factors: (1) a direct impact to GDP from consumer demand for AI-powered products, (2) an increase in the productivity of current factors of production (i.e. capital and labor) and (3) an increase in overall innovation. In our view, the latter two items in particular make it difficult to accurately assess the impact of AI for the economy and may provide upside potential to current market projections.

Wide-ranging applications suggest that AI could transform large swathes of the economy. AI makes a core business function – prediction/data analytics – less costly. Current applications for artificial intelligence include: autos, cybersecurity, healthcare, fintech and business intelligence. We expect AI to become pervasive in software applications across the economy.

Estimated contribution of AI to economic growth by 2035

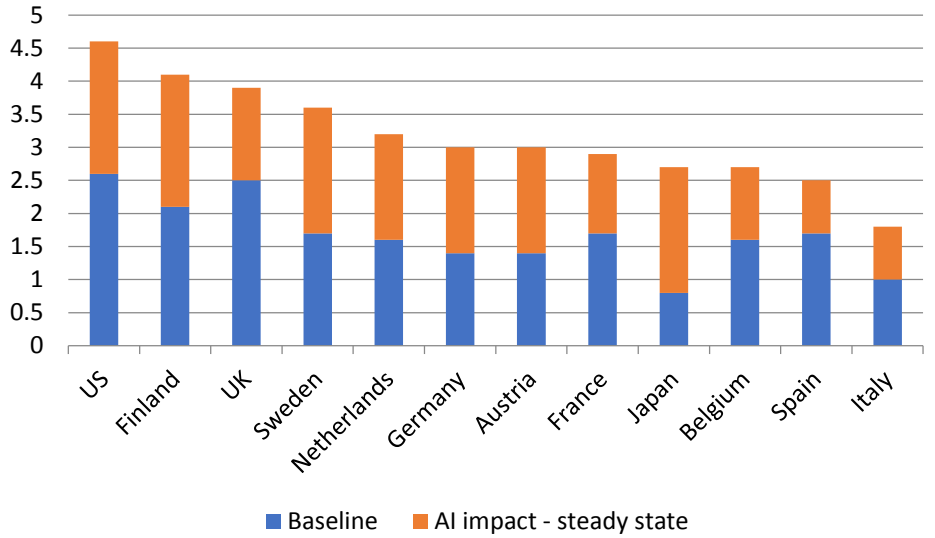

Source: Accenture

> M&A demand high from large corporates

Deal volume has accelerated since 2012 as deep-pocketed players continue to jockey for position. According to Crunchbase data, over 250 AI startups have been acquired over the past five years. Google has been the most active acquirer, with 12 deals completed since 2012. Microsoft, Amazon, IBM, Baidu, Apple, Intel and Facebook are also active. We expect corporate interest to remain high. Notably, mentions of artificial intelligence on public earnings calls has increased exponentially in recent quarters.

Talent & intellectual property (IP) likely the most important consideration for deals. According to Crunchbase data, most startup acquisitions were completed within four years of the company's founding, suggesting acquirers are less focused on purchasing fully-developed, proven technologies. IP accumulation is likely more critical given nascent state of market overall. Picking up scarce AI talent is another key motivator, in our view. Element AI estimated that fewer than 10,000 individuals globally are qualified to perform AI research.

AI M&A Transactions

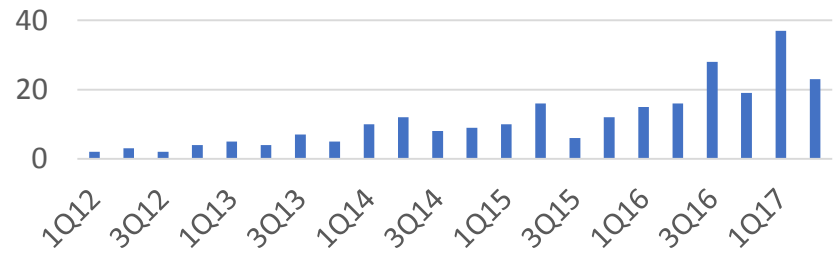

Source: Crunchbase

AI mentions on company earnings calls

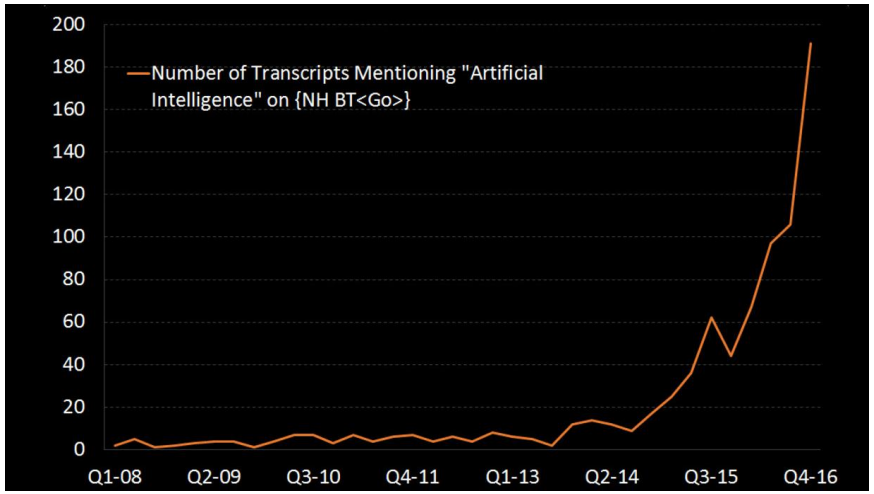

Source: Bloomberg

> Opaque investment environment

The AI startup investment environment is opaque – even by pre-IPO company standards:

- **Many startups are still pre-revenue.** A business model is not proven until it generates sustainable profits and most AI startups are still far from hitting this mark. Of the ~1,100 startups involved with AI in Crunchbase (excludes those that have exited already), less than 40 are later stage ventures. Further, many of these companies have not yet released commercially viable products. That said, some industry-specific AI subsectors do feature more mature companies with immediate revenue/growth prospects (e.g. cybersecurity, sales tech).
- **Very limited transparency.** AI startups tend to be very secretive – a trend that is likely to continue near-term given the market's relative immaturity. Financial data and projections are scarce. In addition, other metrics that may indicate the underlying health and viability of the business (customers, reviews, case studies, etc.) are generally not available publicly. Large companies pursuing AI (Google, Facebook, etc) tend to be muted on details behind their AI strategies as well.

AI startups by funding status

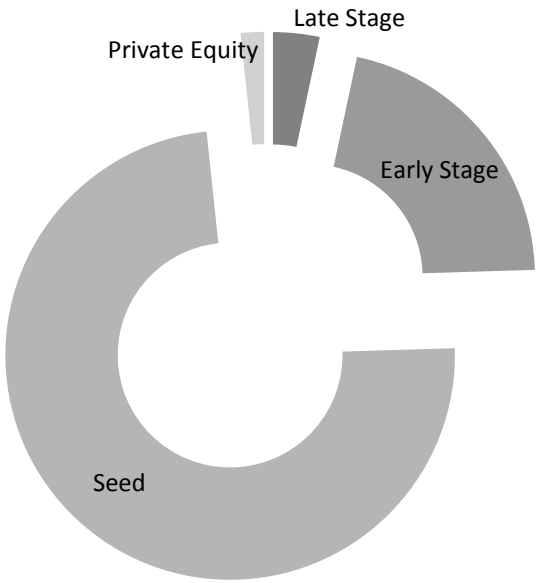

Source: Crunchbase

> Difficult for small firms to develop a competitive advantage

Data requirements favor larger enterprises. Machine learning applications require access to troves of data to hone algorithms. Larger technology firms are at a clear advantage here. Estimates suggests that Google stores 10 EB of data (1EB = 10^9GB) while Facebook collects 500 TB (1TB = 1,000 GB) daily. Outside of the tech world, Walmart receives 2.5 petabytes of data an hour (1PB = 1M GB). Most startups are data-poor by comparison.

Open source has democratized AI algorithms. Artificial intelligence boasts a vibrant open source community. Among companies that have released open source libraries are Google (TensorFlow), Facebook, Yahoo (CaffeOnSpark), IBM, Microsoft and Amazon (Gluon). The availability of these free tools make it difficult for startups to monetize algorithm development.

AI data requirements

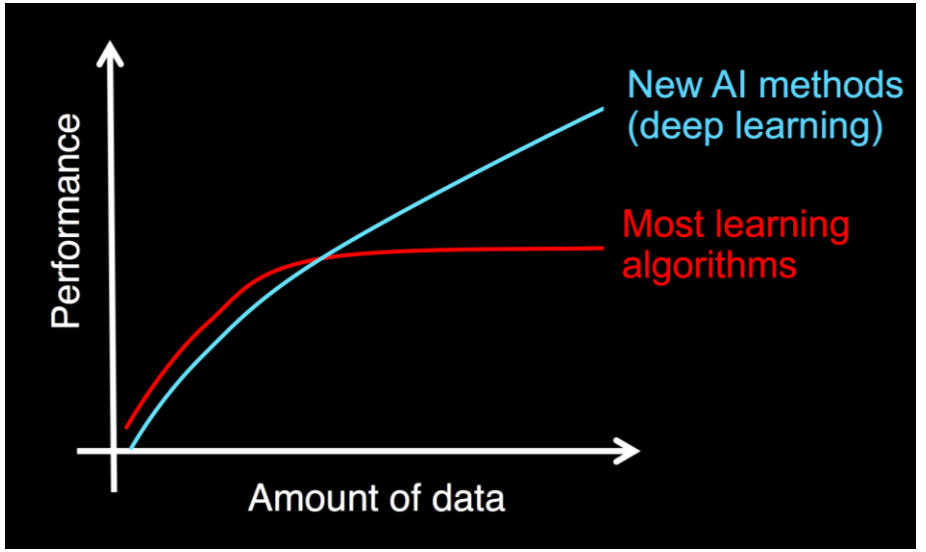

Source: Andrew Ng

> Regulatory/ethical concerns could delay timetables for AI deployment

Ethical and legal concerns could slow widespread adoption of AI. Most AI research does not currently fall under the purview of a regulatory body and no framework has been established to address the integration of artificial intelligence into the economy and legal system. Among the concerns to address, in our view:

- *Sociological and economics impacts from workers displaced by AI.* As a new factor of production, AI could replace certain categories of human labor. Concerns over employment effects have historically stunted (though not stopped) other economic forces such as globalization; we believe AI could face similar obstacles. In our view, thoughtful legislation will likely be necessary to pave a path for sustainable, widespread adoption of AI into the broader economy.
- *Who is responsible for harm caused to humans from AI?* Autonomous vehicles will cause human deaths (as human-driven vehicles currently do). Further, past research shows that AI can develop racial and gender biases. Companies will likely look seek legal frameworks and protections before they deploy their AI technologies widely.
- *Robot rights.* As AI becomes more independent from its human creators, it may require a legal status (similar to how corporations currently). The European Parliament has already debated the rights of "electronic persons"; we believe this will likely need to occur in all countries where AI is deployed.

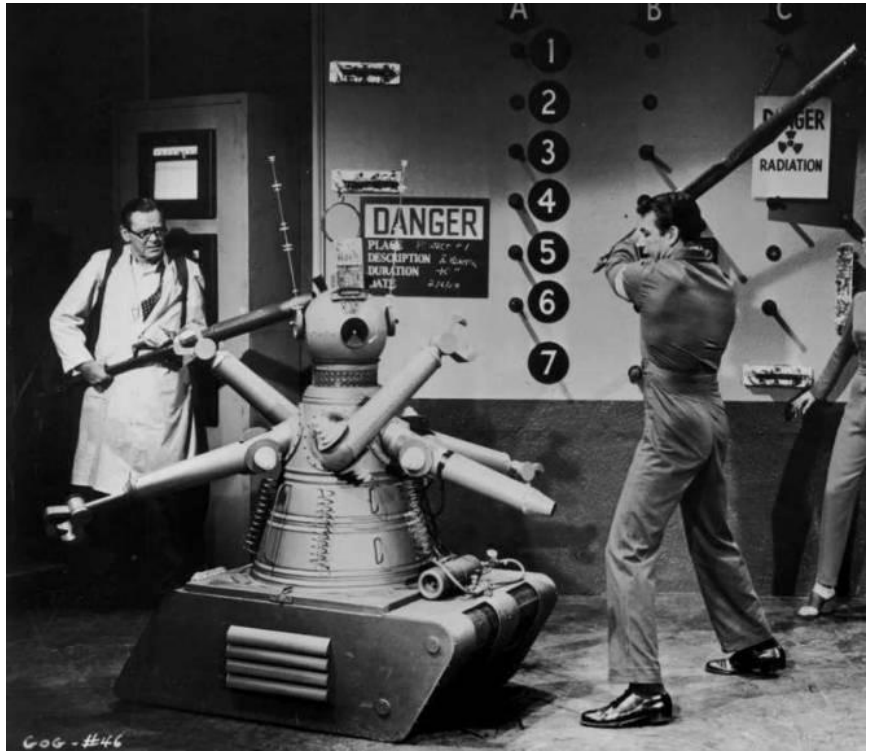

Appendix A: AI Sub-segments

> Core AI

> Application-specific AI

> Industry-specific AI

> Core AI

What is it?	Core AI companies develop foundational technology that facilitates AI deployment (e.g. data processing and modeling). These technologies are generally sector agnostic and target data scientists and engineers.
Why is it important?	Core AI is a key element powering the industry as a whole. Projects vary from researching specific tools to others reaching for the holy grail of the industry – general AI.
Market commentary	Many core AI companies, including almost all general AI companies, do not yet have commercially viable products. Given the broader application potential of core AI technologies, startups in this category may have higher potential payoffs. These are typically longer-term investments, however, and gauging winners from losers can be difficult.
Key Players^	AYASDI, Digital Reasoning, sentient technologies, vicarious

^ Not a comprehensive list. Companies may compete in more than one sub-sector. Key players in a category can specialize in different security areas and listed companies may not directly compete against each other.

> Application-Specific AI

What is it?	Application-specific AI companies develop technologies that help other businesses with a specific application of AI (i.e. analyzing text/images and conversational AI).
Why is it important?	Application-specific AI can be relevant to multiple sectors and businesses.
Market commentary	Startups in this category are generally closer to commercializing their technologies vs. core AI (some already have).
Key Players^	 textio NarrativeScience

^ Not a comprehensive list. Companies may compete in more than one sub-sector. Key players in a category can specialize in different security areas and listed companies may not directly compete against each other.

> Industry-Specific AI

What is it?	Industry-specific AI companies research and develop technologies that apply to a specific industry or vertical. Most AI startups today fall into this category.
Why is it important?	Industry-specific AI companies generally have more immediate prospects for generating revenue and growth; however, the applicability of their products is generally narrow and confined to their specific market/industry.
Market commentary	A meaningful number of these startups have proven technologies, though some are still in the early stages of commercializing these technologies. Cybersecurity, Healthcare and Ad/Sales Tech are among the most well-funded. Key items that investors need to consider are (1) how well does the AI work and (2) how compelling is the underlying market [different verticals have different growth prospects].
Key Players^	BenevolentAI, drive.ai, ZOOX, KENSHO, zymergen

^ Not a comprehensive list. Companies may compete in more than one sub-sector. Key players in a category can specialize in different security areas and listed companies may not directly compete against each other.

Appendix B: Select Startup Profiles*

> Ayasdi

> Digital Reasoning

> Drive.ai

> Insidesales.com

> Sentient Technologies

> Vicarious

> Voyager Labs

> Zoox

* Includes select companies with $50M+ of total funding

> Ayasdi

AYASDI

Description	Ayasdi has built a mathematics driven AI platform that powers the design, development and deployment of enterprise scale, intelligent applications. Their solutions are used by clients in the financial services, government and healthcare industries.
Business Model	Ayasdi's solutions include anti-money laundering, counter fraud, regulatory risk, denials management, population health and clinical variation management.
Management Team	Bob Griffin — CEO Gurjeet Singh — Executive Chairman & Co-founder Gunnar Carlsson — President & Co-founder Harlan Sexton — VP Research & Co-founder
Total Funding: $106M	**Key Investors** IVP Institutional Venture Partners; KPCB Kleiner Perkins Caufield Byers; GE Ventures; khosla ventures

Key Investment Positives	(1) Many reputable customers including Lockheed Martin, Credit Suisse, DARPA, DHS, Siemens, Citi, and Merck
Key Risks	(1) Market still in nascent stages; possible for a competing technology to capture wider adoption.

> Digital Reasoning

Description	Digital Reasoning provides cognitive computing software that understands human communication in many languages to assist in decision making, surveillance, and insight discovery. Their solutions are used by clients in the financial services, government and healthcare industries.
Business Model	Digital Reasoning's sells access to its main product, Synthesys, that enables users to read, understand and analyze large volumes of structured and unstructured language (web, email, social media, instant messages, complaint logs etc.). L12M revenues estimated at ~$18M.
Management Team	Brett Jackson – CEO Tim Estes – President & Founder Prakash Ramachandran – CFO
Total Funding: $74M	**Key Investors** SILVERLAKE SUMERU, Nasdaq, Lemhi Ventures, Goldman Sachs, IQT IN-Q-TEL, Credit Suisse

Key Investment Positives	(1) Reputable customer set including Nasdaq, Goldman Sachs, UBS and Point72 (2) Increasing public emphasis on cyber security and insider trading regulation may provide a tailwind for demand
Key Risks	(1) Proven core concept with initial customers but unclear whether this can scale to wider market

> drive.ai

drive.ai

Description	Drive.ai creates AI software for autonomous vehicles using deep learning. The company has pivoted its business model to focus on refrofit kits that can be used to add self-driving capabilities to existing commercial and business vehicle fleets.
Business Model	Drive.ai is aiming to have its first pilot kits to market by end of year. Given complex installation, the kits will be marketed towards commercial fleets (not consumers).
Management Team	Brody Huval — Co-Founder Joel Pashayampallil — Co-Founder Tao Wang — Co-Founder Sameep Tandon — CEO
Total Funding: $62M	**Key Investors** 北极光创投 northern light VENTURE CAPITAL; NEA; GGVCAPITAL

Key Investment Positives	(1) Many potential customers likely already own large car fleets; retrofit kits may gain faster adoption among these clients given lower upfront capex vs. buying autonomous fleets outright. (2) Board incl. NEA chairman Carmen Chang and Baidu AI alum Andrew Ng, who could help drive.ai expand to China.
Key Risks	(1) Pre-revenue company. Business model not yet proven. (2) Potential for cost overruns and time delays with pilot kit. (3) Market still in nascent stages; possible for a competing technology to capture wider adoption.

> Insidesales.com

INSIDESALES.COM

Description	InsideSales.com sells SaaS sales technology that leverages machine learning to predict and recommend optimized sales activities. The company claims is customers experience revenue growth of up to 30% in as little as 90 days.
Business Model	InsideSales' software platform includes: playbooks, lead & account scoring, predictive pipeline, sales communications, email & web tracking and gamification. Subscriptions range from $95-295/month for each user.

Management Team	
Dave Elkington	Chairman, CEO and Founder
Ken Krogue	President and Funder
Steve Brain	CTO and Chief Product officer
Elliot Smith	CFO

Total Funding: $251M

Key Investors

Polar Capital

salesforce ventures

POLARIS PARTNERS

Key Investment Positives	(1) Key customers include Fidelity, Microsoft, Groupon, McGraw-Hill, Marketo ADP, Apttus, BMC, Caesars Entertainment and Century Link. (2) Could be an attractive acquisition candidate for companies looking to broaden their CRM capabilities.
Key Risks	(1) Limited financial disclosure. (2) Deep-pocketed competitors including SAS, SAP and Salesforce (though neither offer an identical product currently)

Denotes a unicorn company (valuation >$1B)

> Sentient Technologies

Description	Sentient Technologies delivers AI software solutions powered by a large, distributed processor network to run deep learning processes based on natural selection. They deploy solutions across many disciplines including e-commerce, marketing, and finance.
Business Model	Sentient has deployed its software to several major clients: - Sentient Aware: visual search tool (Skechers, Tumi, Sunglass Hut) - Sentient Ascend: conversion optimization (Cosabella, Nexway) Financial Information not available
Management Team	Antoine Blondeau — Co-Founder & Co-Chairman Babak Hodjat — Co-Founder & CEO Gurmeet Lamba — COO & GM, Sentient Ascend
Total Funding: $144M	**Key Investors** TATA COMMUNICATIONS, ACCESS TECHNOLOGY VENTURES, Horizons Ventures 維港投資

Key Investment Positives	(1) Well positioned to benefit from strong tailwinds in marketing optimization technology, "85% of customer interactions in retail to be online by 2025" – *Gartner*. (2) Well funded with high quality board representatives e.g. Mary Meeker (KPCB)
Key Risks	(1) Lots of external interest but ability to execute on mass market retail adoption remains to be seen (2) Market still in nascent stages; possible for a competing technology to capture wider adoption.

> Vicarious

Description	Vicarious is creating software that thinks and learns like a human brain, artificial general intelligence, for robots. As opposed to deep learning techniques that require exposure to thousands of training examples, Vicarious' AI learns and retrains much faster.
Business Model	Vicarious aims for its technology to be available in robot hardware across industries ranging from manufacturing to agriculture to medical sciences.
Management Team	D. Scott Phoenix — Co-Founder & CEO Dileep George — Co-Founder
Total Funding: $134M	**Key Investors** FOUNDERS FUND khosla ventures Mark Zuckerberg, Elon Musk, Peter Thiel, Jeff Bezos, Ashton Kutcher

Key Investment Positives	(1) Roster of top tech investors including Elon Musk, Mark Zuckerberg, Jeff Bezos, and Peter Thiel (2) Awarded accolades from World Economic Forum and Goldman Sachs (3) General intelligence has much greater potential than current deep learning models
Key Risks	(1) Pre-revenue company. Business model not yet proven. (2) Company is very secretive, little information on customers or financials

> Voyager Labs

Description	Voyager Labs has developed an artificial intelligence software engine that analyzes unstructured public data to provide dynamic and real-time insights into user behavior. Their solutions are used by customers from the retail / ecommerce, finance, consulting and public sectors.
Business Model	Voyager Labs sells access to two core products: (1) Voyager Analytics: Insights for investigation, vetting, and other security-related work. (2) VoyagereCommerce: Engine for more personalized recommendations, personalized offerings and more precise targeting campaigns.
Management Team	Avi Korenblum — Founder & CEO Ron Aplboim — CFO Jay Klein — CTO Maya Racine Netser — COO
Total Funding: $100M	**Key Investors** Horizons Ventures 維港投資 ORACLE Ronald Cohen Lloyd Dorfman

Key Investment Positives	(1) Recognized by Gartner in 2017 as a “cool vendor” (2) Recently emerged from stealth mode with ample funding of $100M
Key Risks	(1) Likely still a pre-revenue company. Business model and core technology not yet proven.

> Zoox

ZOOX

Description	Zoox is a robotics company developing the first ground-up, fully autonomous vehicle fleet. The company aims to pioneer the next generation of mobility-as-a-service in urban environments.
Business Model/ Financials	Zoox is in stealth mode; as such, no financial data is available.
Management Team	Tim Kentlry-Klay — Co-Founder and CEO Jesse Levinson — Co-Founder and CTO James Philbin — Senior Director of Computer Vision
Total Funding: $290M	**Key Investors** DFJ

Key Investment Positives	(1) Large total addressable market. Global Market Insights projects car-sharing services will reach $16.5B of revenue by 2024.
Key Risks	(1) Pre-revenue company. Business model not yet proven. (2) Technological and regulatory feasibility.

Denotes a unicorn company (valuation >$1B)

Appendix C: Key AI Milestones

1950	• Alan Turing publishes ***Computing Machinery and Intelligence***. His proposal for evaluating a machine's ability to exhibit intelligence later comes to be known as the ***Turing test***. • Isaac Asimov published his ***Three Laws of Robotics***. His work was popular, thought-provoking and visionary, helping to inspire a generation of roboticists and scientists.
1955	• John McCarthy, Marvin Minksy, Nathaniel Rochester and Claude Shannon propose a study of artificial intelligence which culminates as a conference the next year at Dartmouth College. Their initial work is generally considered the official birth of AI (they also coined the term).
1957	• Frank Rosenblatt develops the ***Percepton***, an early artificial neural network.
1958	• John McCarthy develops ***Lisp***, which becomes the most popular programming language for AI research
1965	• Joseph Weizenbaum develops ***ELIZA***, a program designed to carry teletyped conversations with humans (similar to SIRI today).
1969	• Researchers at Stanford Research Institute (SRI) came out with ***Shakey the Robot*** after spending six years on developing it, demonstrated combining animal locomotion, perception and problem solving.
1973	• After AI progress stalls, funding for the industry was slashed, ushering in what became known as the AI winter.
1981	• The first successful commercial expert system, known as the ***RI***, began operation at the Digital Equipment Corporation helping configure orders for new computer systems. By 1986 it was saving the company an estimated $40m a year.
1986	• ***Backpropagation*** ("backprop") – the single most important algorithm in the history of machine-learning becomes mainstream. The algorithm was originally developed in 1969. Also, the first driverless car is developed in Munich (from a Mercedes-Benz van).
1988	• Judea Pearl publishes ***Probabilistic Reasoning*** in Intelligent Systems, which introduced Bayesian networks and marked a large step forward in the field of AI.

Appendix C: Key AI Milestones

1997	• IBM's Deep Blue Supercomputer beats world chess champion Garry Kasparov in a chess battle. The IBM-built machine was, on paper, far superior to Kasparov - capable of evaluating up to 200 million positions a second.
2002	• Rodney Brook's spin-off company, i***Robot***, created the first commercially successful robot for the home – an autonomous vacuum cleaner called Roomba.
2004	• NASA's robotic exploration rovers ***Spirit and Opportunity*** autonomously navigate the surface of Mars.
2009	• Google builds self driving car.
2010	• Microsoft launched ***Kinect for Xbox 360***, the first gaming device to track human body movement, using just a 3D camera and infra-red detection, enabling users to play their Xbox 360 wirelessly.
2011 – 2014	• IBM's ***Watson*** computer defeated television game show Jeopardy! champions Rutter and Jennings. • Apple's ***Siri*** (2011), Google's ***Google Now*** (2012) and Microsoft's ***Cortana*** (2014) are smartphone apps that use natural language to answer questions, make recommendations and perform actions.
2012	• Andrew Ng and a team of computer scientists achieve a major deep-learning breakthrough while at Google's X lab with a neural network that identifies cat videos on the internet.
2016	• Google DeepMind's ***AlphaGo*** defeated Lee Sedol 4–1. Lee Sedol is a 9 dan professional Korean Go champion who won 27 major tournaments from 2002 to 2016.[55] Before the match with AlphaGo, Lee Sedol was confident in predicting an easy 5–0 or 4–1 victory
2017	• Google DeepMind's ***AlphaGo*** won 60–0 rounds on two public Go websites including 3 wins against world Go champion Ke Jie.