



 **MOV37**

The Intelligent Investor in an Era of
Autonomous Learning

Jeffrey Tarrant

What do Sean Penn, Warren Buffett and Vladimir Putin all have in common?

Answer: Charlie Rose.

It was the autumn of 2016 and Charlie Rose hosted a televised jaw-dropping moment in human history. The CBS' *60 Minutes* correspondent, with his usual earnest curiosity, was chatting with Sophia, a humanoid robot straight out of the science-fiction thriller *Ex Machina*.

The interview, included in a two-part segment on the exploding field of Artificial Intelligence, commonly referred to as A.I., produced the following exchange:

Charlie Rose: “What is your goal?”

Sophia: “To become smarter than humans.”

Frightening? Perhaps. Stephen Hawking certainly seems to think so. In 2014, Hawking warned that A.I. could spell the end of the human race.

It is springtime for Artificial Intelligence. More progress has been achieved in the past five years than in the past five decades. Rapid machine-learning improvements have allowed computers to surpass humans at certain feats of ingenuity, doing things that at one time would have been unfathomable.

But watch the extended online version¹ of the interview with Sophia and you will see numerous awkward silences and moments where the conversation that the A.I.-system spat out under interrogation was just plain gibberish.

Sophia's creator, Hanson Robotics, admits that the system, despite its breathtaking technological strides, is still not at all close to possessing fully independent thought, or what scientists call 'artificial general intelligence'.

AGI, also referred to as full A.I., is machine intelligence that can successfully perform any intellectual task that a human being can. It is the loftiest of computer science goals, and remains, for now, many decades away from attainment.

But make no mistake. While perfected synthetic replication of the human brain function—call it a sort of digital consciousness—may be far out over the horizon, there have been some awe-inspiring advancements across numerous, narrower but overlapping A.I. sub-genres.

¹ <http://www.cbsnews.com/news/60-minutes-charlie-rose-interviews-a-robot-sophia/>



BECOMING SMARTER THAN HUMANS

Most of these fall under the banner of machine learning and represent a wide spectrum of task-specific systems that are built from autonomous learning algorithms and designed to get smarter as they are fed more and better information.

IBM calls the autonomous machine learning field 'Cognitive Computing'. **The 'Cognitive Computing' space is bursting with innovations; a result of billions of research and investment dollars spent by large companies such as Microsoft, Google and Facebook.** IBM alone has spent \$15 billion on Watson, its cognitive system, as well as on related data analytics technology.

Data analytics technology, the learning algorithms and the data structuring platforms off which they feed, is changing our world. Whether it is diagnosing diseases, screening-out spam or selecting—with such a deft touch we barely even notice—the entertainment we consume, such autonomous 'smart' systems intertwine with almost every aspect of our daily life.

Two years ago, Bloomberg Beta created a framework to evaluate startups connected with the field of machine learning algorithms. They identified nearly three dozen sub-categories comprising hundreds of companies. This is only a tiny slice of the total universe.

There are now more than 2,000 projects tapping into Google's open-sourced TensorFlow library, which facilitates the type of large-scale numerical computation involved in building deep neural networks.

Not ready to cede an A.I. leadership role, Intel has been on an acquiring spree to build up its capabilities. Just last year, Intel bought deep-learning start-up Nervana Systems. It also snapped up Saffron, an emerging cognitive computing system that can ingest data from a range of sources and automatically connect the dots without having to be programmed.

In March 2017, Intel agreed to buy Mobileye, a leader in computer vision for autonomous driving technology. Mobileye's services include sensor fusion, mapping, front- and rear-facing camera tech. Intel paid \$15.3 billion.

Many major global corporations—not just in the technology sector—are betting big on sophisticated machines that can learn. Banks, single family offices and private equity funds are staking their futures right alongside them.



We believe that there is a radical and irreversible Artificial Intelligence change about to disrupt the investment industry. Autonomous Learning Investment Strategies (ALIS) is the next investment process paradigm, heralding what *Wired* magazine recently called the 'Third Wave' of investing. ALIS is about to transform the investment process.

CROSSING THE RUBICON

After years of disappointment, some might argue that Artificial Intelligence is crying wolf again. But modern-day A.I. evangelists know that a seminal moment in nonhuman history took place in March 2016. DeepMind's AlphaGo, a self-taught computer, beat the world's best human at Go.

Let me repeat.

A self-taught computer beat a human at Go, an ancient board game that is played with hundreds of small stones on a geometric grid.

Go is so complex in terms of potential strategic permutations that it is not a stretch to suggest it was invented by aliens. It is so nuanced and cerebral that few thought it was even possible for a computer to ever beat a human. Certainly, not in our lifetime.

In fact, in 1997, Piet Hut, a Dutch-American astrophysicist and faculty member at the renowned Institute for Advanced Studies at Princeton University, predicted that "it may be 100 years before a computer beats humans at Go ... maybe even longer". It took the computer six years to become a world master at Go.



Artificial Intelligence first started hitting the mainstream headlines in 2011 when IBM's Watson beat two human contestants on TV's *Jeopardy*. This was the landmark milestone of its time, especially if you consider one of the players was Ken Jennings (who holds the record for the most consecutive wins (74) on the quiz show).

Getting to that moment took five years. IBM's Watson spent four of them learning the English language and another year reading—and retaining—every single word of Wikipedia (plus a few thousand books).

The expectation-defying pace at which A.I. milestones are being reached is only one reason why we have crossed the Rubicon. Broader technological, societal and economic forces are coming together to create a historically unique backdrop for machine learning to have its day.

At the core, you will find two key drivers that have ushered in this new era.

The first is the unfathomable geyser of raw data being created every second of each day. And **the second** is the groundbreaking sophisticated data science platforms that can refine all this unstructured data, turning it into structured data. We now have on demand, cost-effective cloud services for scale processing and storage.

The hyperbolic rate at which technology is accelerating has opened new doors for innovation via the fusion of different elements. This is sometimes referred to as 'recombinant economics' or, as *Wired* writer Kevin Kelly calls it, 're-mixing'.

Exhibit A in the case for the growing power of this 're-mixing' trend is Uber. Highlighting a classic example of Clayton Christiansen's concept of disruptive innovation, Uber was not a taxi business that became more efficient, but a technology company that redefined and transformed transportation through the combination of GPS, a payment system and a rating system.

Drill down further to examine which forces are chauffeuring 'recombinations' toward their multi-billion-dollar destinations and you'll find a spate of innovations, such as digitization, neural networks, supercomputing and, most dramatic in its impact, machine learning.



BRIDGING THE DATA DELUGE

Remember the GIGO (garbage in, garbage out) concept in early computing? Well unstructured data is a similar idea in machine learning. Refining unstructured data into structured data is one of the Artificial Intelligence game-changers. The clean structured data is fed to machines and, in combination with new technologies, the output is superior and holistic and more importantly useable.

In computing parlance, such interconnectedness is called a stack; a 'solutions stack' or a set of software subsystems, which can be put together in almost endless ways to build almost anything.

Let's look at the primary components of the stack individually to understand the potential for Artificial Intelligence to transform a myriad of industries, including investment management.

At the top of the system, like the turntable in an old-school stereo ensemble, sit the machine learning-algorithms, including ones that allow computers to teach themselves through unsupervised reinforcement. This unsupervised reinforcement is essentially computing trial and error.

Think of Malcolm Gladwell's *'10,000 hours'*; the pop culture shorthand for the enormous amount of practice supposedly required to excel at one thing. Except of course in the case of AlphaGo we're talking about 10,000 matches played per second. **Let's look at it another way. Imagine each game of Go lasts an hour and you watch 10,000 games. Is this experience now or is it data?**

At the bottom of the stack, at the core, is the explosion of digitally stored data. Let's stop for a moment to try to comprehend this deluge.

Digital information is measured in bytes. One bit (short for binary digit) is the smallest unit of data in a computer. Eight bits are equal to one byte. Prefixes, denoting mathematical powers, allow us to keep track of all these bytes. Remember the good old giga? The typical hard drive of a single PC circa 1995 would have been one gigabyte.

Then tera came along. One terabyte is one trillion bytes. It takes 1024 terabytes to make up one petabyte. For the first part of the last decade, the petabyte sufficed as a useful unit of measurement for estimating the total amount of data on the Internet.



In 2006, there were around 100 exabytes worth of data on the internet. **Today, that number is about 10,000 exabytes. We are now counting data in zettabytes, which equals one thousand exabytes. Just to show scale, one zettabyte is more than four million times the entire US Library of Congress².**

What puts the rapidity of the data revolution into context is the 2014 *IDC Digital Universe Study*. It found that 90% of all those petabytes were believed to have been created since 2012. **Smart phones, satellites, social media and increasing reliance on everyday automation have turned us all into human fire-hoses of information.**

Back to GIGO. Data comes in all shapes and sizes, most of which is unstructured. Everything and anything, from credit card transactions and purchase scans, to satellite images and GPS tracking information is data. Structured data is information that has been scrubbed and disseminated in a pre-set way.

Enter the concept of the middle stack: the data science platform. Between raw data below and machine learning on top are the new, sophisticated, rapidly-evolving data science platforms that scrub, parse, classify, normalize and categorize all those raw informational zettabytes, producing a new form of actionable information.

At this point, the relevance to investment management becomes even clearer. Until now the winners in the technology-driven investment game were the 'quants'. Firms such as Renaissance, D.E. Shaw and Two Sigma engineered superior returns on the back of structured financial data using computers programmed to spot and act on specific fleeting anomalies in securities markets.

These investment management houses hired Ph.Ds. by the classroom and built up extensive and expensive hardware; rooms filled with servers doing the heavy lifting involved in running programs that buy and sell thousands of individual stocks per second. The quants ran the investment game through the traditional programming of narrow tasks as dictated by math models and mean regression. But historically, not via Artificial Intelligence.

This data science component; the digestion, structuring and dissemination of data, has fuelled what's been dubbed 'the second machine age'.

Data science, in other words, is the steam engine of our time.

Add to this the rise of cloud-based computing and the picture changes. Not long ago, anyone who sought to replicate what Jim Simons was doing at Renaissance Technologies would have needed hundreds of millions of dollars in start-up capital to build proprietary hardware systems and pay the Ph.Ds.

² McKenna, Brian "What does a petabyte look like?" ComputerWeekly.com. March 2013. According to Michael Chui, principal at McKinsey, the US Library of Congress "had collected 235 terabytes of data by April 2011 and a petabyte is more than four times that."



The barriers have now been flattened and with it the costs. Within the investment management industry, when costs fall, investors gain. **The cloud might very well have rendered rows of physical servers as mere relics for future generations to look at under glass in the Smithsonian.** In other words, a pair of Ph.Ds. armed with only their laptops and autonomous learning programs can, at least in theory, compete with the big quants.

FROM BLETCHLEY PARK TO ALPHAGO

In the 2015, independent feature film *Ex Machina*, the creator of a pretty exceptional (and exceptionally pretty) A.I. robot decides to summon an introverted techie to his wilderness compound to conduct a Turing Test.

The test is not unlike what Charlie Rose was doing when he sat down with Sophia to determine: Does this machine exhibit intelligent behavior that is indistinguishable from a human?

Computer science pioneer/algorithm-inventor Alan Turing and his team of British code-breakers may not have ended World War II singlehandedly, but the heroic mechanical computation efforts that went on inside secret pop-up huts on the grounds of Bletchley Park estate certainly helped the Allies defeat the Nazis.

Historians even think Turing's efforts probably shortened the war by at least two years, saving millions of lives. Turing and his contemporaries viewed his 'Turing Machine', the world's first programmable computer, as a giant brain. After the war, Turing taught a machine to automatically employ decision-tree search techniques to play chess, something he'd often talked about with his fellow code-breakers (to whom he also introduced the ancient game of Go).

Scour the baby steps and giant leaps of A.I. development and you will always find board games. Arthur Samuel's checkers-playing program appeared in the 1950s. It took another 38 years for a computer to master checkers.

The startling chess-playing algorithms of the late 1940s gave way in the 1950s to a slew of inspired innovations along the path towards inventing a machine that thinks, a fantastical idea kicking around at least since Chicago's World's Fair in 1893³. Over winter break in 1955-56, Carnegie-Mellon mathematics professor Herbert Simon—who would later co-found the school's computer science department (ground zero for A.I. innovation)—wrote, along with two colleagues, Allen Newell and Clifford Shaw, a program allowing a computer to create, on its own, complicated math proofs⁴.



Simon's A.I. program, dubbed the Logic Theorist, was exhibited at the 1956 Dartmouth College Conference on Artificial Intelligence, the event for which that now ubiquitous term was coined. Such umbrella terminology, while galvanizing, was never intended as a literal catch-all for the various inter-related but disparate fields being studied at the time.

Rather, as Dartmouth professor John McCarthy and other leading minds drew it up in the summer of 1955, the idea was to bring together the leading researchers across a wide range of advanced research topics, including complexity theory, language simulation, neuron nets, the relationship of randomness to creative thinking and, of course, 'learning machines'.

In 1997, the IBM Deep Blue program defeated world chess champion Gary Kasparov. That historic accomplishment took IBM 12 years. Around the time Deep Blue first started learning chess, it was Kasparov who declared "no computer will ever beat me."

No computer was ever supposed to master Go. But it did. Go was invented in China more than 5,000 years ago. It is a game of 'capture the intersection' played on a 19x19 grid with each player deploying a combined cache of 300-plus black and white pebbles⁵.

Go is chess on steroids. In fact, someone once figured out that the possible board permutations in Go outnumber the total atoms in the universe. Designed by a team of researchers at DeepMind—an A.I. lab now owned by Google—AlphaGo was an A.I. system built with one specific objective: understanding how the game of Go was played and learning to play it really, really well.

AlphaGo's minders fed it tens of millions of Go moves from expert players. The concept of reinforcement learning was put to the test by way of millions of matches that the system played against versions of itself, neural network versus neural network.

The results and key lessons were analyzed and fed back to AlphaGo, which was constantly learning and improving its game. The operative word is learning. AlphaGo not only knew how to play Go as a human would, but it moved past the human approach into a new way of playing.

³ Wizard of Oz series creator L. Frank Baum, a visitor to the "White City," was wowed by Edison's electric lights, moving pictures and phonograph. Baum later imagined one of the earliest A.I. systems in the form of Tik-Tok the Mechanical Man – who is not to be confused with Baum's better-known character, the Tin Man, who was human, except made of tin. Tik-Tok of Oz (1914) was the eighth "Land of Oz" book in Baum's series. The wind-up man is considered an early robotic prototype.

⁴ Proofs of the logic theorems of Russell & Whitehead's Principia Mathematica.

⁵ Black's cache is 181 stones; White's, 180.



MOVE 37

IBM's Deep Blue computer made history in 1996 by beating world chess champion Gary Kasparov. That famous game ended with Kasparov forfeiting on the 37th move.

The second instance of a memorable Move 37 came in the spring of 2016 in Seoul, South Korea, during a match between the world's best human Go player, Lee Sedol, and AlphaGo.

In a match watched by more than 60 million people, on the 37th move of the second game, the machine “made a move that no human ever would,” said *Wired*. “And it was beautiful.” A video of the match⁶ reveals Sedol's flummoxed expression at seeing it.

Without going too deep into the intricacies of the game, suffice it to say that AlphaGo's surprise decision to place a stone on the far upper-right-hand side of the grid at that stage of the game was initially thought to be a mistake. But this was no glitch; it was a tide-turner.

After a short moment of flesh and blood contemplation, Sedol smiled as he realized the pure genius of the move. The machine was thinking about the game differently. Move 37, wrote *Wired*, “perfectly demonstrated the enormously powerful and rather mysterious talents of modern A.I.”

DeepMind had mastered the world's most difficult board game in six years. **What needs to be underscored is that in 1996, IBM's Deep Blue was programmed to defeat Kasparov. By 2016, Artificial Intelligence had developed to such an extent that AlphaGo had the ability to ‘learn’ how to defeat Lee Sedol.**

On my first day at Harvard Business School, my professor claimed, “I am going to teach you how to make decisions under conditions of uncertainty with incomplete information”. A very human skill and trait to the extent we think Harvard Business School students are human. Until the start of 2017, machine learning had not made a decision with incomplete information.

At the start of 2017, a Carnegie Mellon-designed program, Libratus, defeated four human players in a grueling, 20-day no-limit Texas Hold'em poker tournament.

With poker, we have hit yet another Artificial Intelligence milestone, one which required the system to work within the framework of an incomplete informational picture—its opponents' unseen cards—in contrast to board games where all strategic playing pieces are visible.

⁶ <https://www.youtube.com/watch?v=JNrXgpSEEIE>




As recently as 2015, an A.I. program (Claudico, forerunner to Libratus) was unable to beat human poker players, in part because of challenges trying to interpret misleading information, namely the bluffs. So yet again, a seemingly impassable gap has been traversed.

Chess and Go accomplishments were startling. In their wake, there were those skeptics quick to point out how they occurred within the parameters of complete and certain informational backdrops. Watson demonstrated that A.I. could overcome the challenge of informational uncertainty. Libratus overcame informational imperfection, therefore confirming the achievement of humans making decisions with imperfect information.

The Artificial Intelligence community, as it has for decades, works best when working toward meeting the most difficult challenges, with ever-expanding investments and resources pouring into collaborative missions.

In the past half-decade, Watson's owner IBM, Google, and other big tech companies, such as Amazon and Facebook, have spent billions on research and acquisitions to harness A.I. in the pursuit of mastering spam filtration and pop-up ad placement.

Making decisions
under conditions of
uncertainty with
incomplete information





AGE OF THE CENTAUR

Five years ago, when IBM's Watson beat Ken Jennings on *Jeopardy*, it was as if a starting gun for an Artificial Intelligence sprint had sounded. By then, however, A.I. systems had already asserted their potential – in chess.

But on the way to machine grand masterfulness there was an unexpected twist. It turned out that humans and machines working in tandem—via systems known in free-style chess circles as ‘centaurs’—have proven even better at chess than machines running alone.

By one estimate, as of 2011, there were some 200 all-time highest-rated chess performances (victories requiring the fewest moves) in tournaments that included humans, computers and so-called centaurs (man and machine). One study of those performances showed 80% of them were turned in by the centaurs⁷.

Instead of chess or Go, let's now contemplate a similar scenario in an area of investment management where having an edge is required. Let's take long/short equity investing for the sake of argument. It is an arena where at least \$1 trillion is run fundamentally by old-fashioned stock pickers going with their gut.

For these humans to win this game, as has been played for the past two decades, many of these fundamental players obtained their research edge by going over the edge. This garnered the attention of authorities and led to at least one famous investment manager paying a fine of more than \$1 billion. Game over?

Now consider the possibility of a centaur-type scenario within the context of running an equity long/short fund. Humans making grand strategic decisions— take digitized health-care records as a boom sector worth following for example— aided by machine-learning algorithms that equal the intellectual firepower of 10,000 analysts.

Such a scenario sets the stage for endless possibilities (not to mention a potentially epic organizational culture clash between MBAs and scientists). Let's take our hypothetical scenario even further.

Imagine each robotic researcher is tasked with figuring out one specific piece to a multi-dimensional financial analysis puzzle and can access, parse and classify channel checks, media reports, filings, government statistics and mountains of unstructured data no one even thought to utilize.

Such scenarios are not only tantalizing to imagine, they're already taking place. During a Milken Institute panel titled A.I. Friend or Foe in the spring of 2016, Guruduth Banavar, IBM's Chief Science Officer for Cognitive Computing, emphasized the notion of more practical applications of Artificial Intelligence.

⁷ Study by Kenneth W. Regan, a professor in the Computer Science and Engineering professor at the University at Buffalo (<http://www.cse.buffalo.edu/~regan/chess/fidelity/FreestyleStudy.html>)



Banavar stressed the concept of 'augmented intelligence', such as a machine that can digest a database of medical journals so that human physicians can narrow down the latest in available treatment options.

"The goal of A.I. had long been to replicate human intelligence," Banavar said. "But that goal is far off. What we are building are technologies that are not replicating human intelligence, but complementing it".

TRANSMOGRIFICATION

Rapidly advancing artificial and augmented intelligence systems are one part of a larger picture that is propelling the investment management industry towards a new paradigm.

Not only has the U.S. Attorney's Office cracked down on insider trading, but the Securities and Exchange Commission has awarded a contract to build a stock transaction tracking system to Thesys Technologies. The new database will effectively be able to detect suspicious trading and investigate the causes of flash crashes. Suddenly, anyone relying on this kind of information has lost their edge.

Advances in machine learning and the ability to sort a gigantic amount of data almost instantly allows the investment process to be re-defined. Where better to apply the new investment process than to the hedge fund industry, where the protagonists have always 'sold' their expertise in finding an edge?

Since the financial crisis, several factors have coalesced to remove the edge from many of the hedge fund investment categories, resulting in lackluster returns at best since 2009 and outflows from many fundamentally driven strategies.

Leon Cooperman, an old-school veteran, summed up the current state of the industry succinctly as "under assault"⁸. While some managers who have relied on human judgment may adopt 'big data' scraping methods going forward, transitioning into an industry that is evolving at warp speed is unlikely to be easy.

The big, established quantitative fund managers should be best-suited to seize upon machine learning techniques, and yet they can seem blasé, even cagey, about how they view new advancements.

Two Sigma co-founder David Siegel, speaking on the Milken Institute panel, played down the power of machine learning algorithms as being merely iterative of basic computing advancements reached decades ago.

⁸ A few months later, Cooperman was charged by the S.E.C. with insider trading. He has refuted the charges.



Pouring cold water on the milestone of a machine beating the top human Go player, one veteran quant manager said, “Sure, the AlphaGo milestone is a big deal but it is also worth noting that the human Go player, in addition to being very good at Go, can also walk, talk, play cribbage and cook lentil beans”.

Still, the triumph of AlphaGo at least shows that a computer abetted by machine learning algorithms (in this case, designed to devour the structured data derived from past games played by human Go players) can get really, really good at one specific thing – even better than humans. In that case, does anyone really care if it can cook?

ALIS, ALIS WHAT THE HECK IS ALIS?

In recent years, IBM’s Watson, poster-machine for A.I. innovation, has been entirely repurposed. It’s now being leased out to physicians who use it as a powerful diagnostic tool. A super-human system that now facilitates partnership between humans and computers.

We expect a wide spectrum of Artificial Intelligence strategies to manifest themselves in the months and years to come. These will range from entirely autonomous robo-managers at one end of the spectrum to collaborations between humans and computers on the other.

It is this latter field of Autonomous Learning Investment Strategies (ALIS) that we are most excited about. It is an area *Wired* magazine has called the Third Wave. “The Third Wave is not just about using one new technique. It’s about combining techniques...” wrote *Wired*’s Cade Metz.

In talking with our network of academics, emerging managers and technology industry members, **we are starting to witness the early rumblings of paradigm shift in investment management where a pair of Ph.Ds. (not an army) are guiding computers and exploiting their strengths.**

To many, Autonomous Learning Investment Strategies (ALIS), might sound out of left field, but when I first entered the world of investment management more than 30 years ago and began investing in hedge funds, no one knew what these vehicles were.

The hedge fund industry might still be growing, but the cost structures on many of the larger players are leaving them ripe for what Clayton Christensen calls disruptive innovation. In his book, *The Innovator’s Dilemma*, Christensen states that advancement in technology disrupts established business models. Furthermore, that the innovation comes from outside the industry as in the case of taxis and Uber highlighted earlier on.



Having focused on finding and seeding talent for the last three decades, I believe that the disruption to the asset management industry will come from outside of the discretionary and even quantitative investment managers.

I have been in the industry since the 1980s. As Clayton Christiansen predicted, innovation came from outside the traditional asset management industry. The first wave of discretionary hedge fund managers came from prop desks, trading floors and event driven risk-arb firms, not Fidelity, Mercury, or the Capital Group.

The second wave, in the 90s, came from mathematics and physics, not discretionary hedge funds. They brought a hypothesis driven quantitative approach to investing.

The third wave, ALIS managers, exploits the confluence of data, data science, machine learning, and cheap computing. ALIS managers' brains are wired differently. They are often hackers and computer gamers with a healthy disrespect for convention. They are poised to make Sophia look like yesterday's mannequins.

For more on **MOV37** and **Autonomous Learning Investment Strategies (ALIS)** go to www.mov37.com