



Cutting Through the Smoke and Mirrors of AI on Wall Street

Artificial intelligence made lots of headlines in 2017. Alphabet (GOOGL) developed software that defeated the defending world champion in Go, then a few months later developed a new version that defeated the prior version 100 games to none.

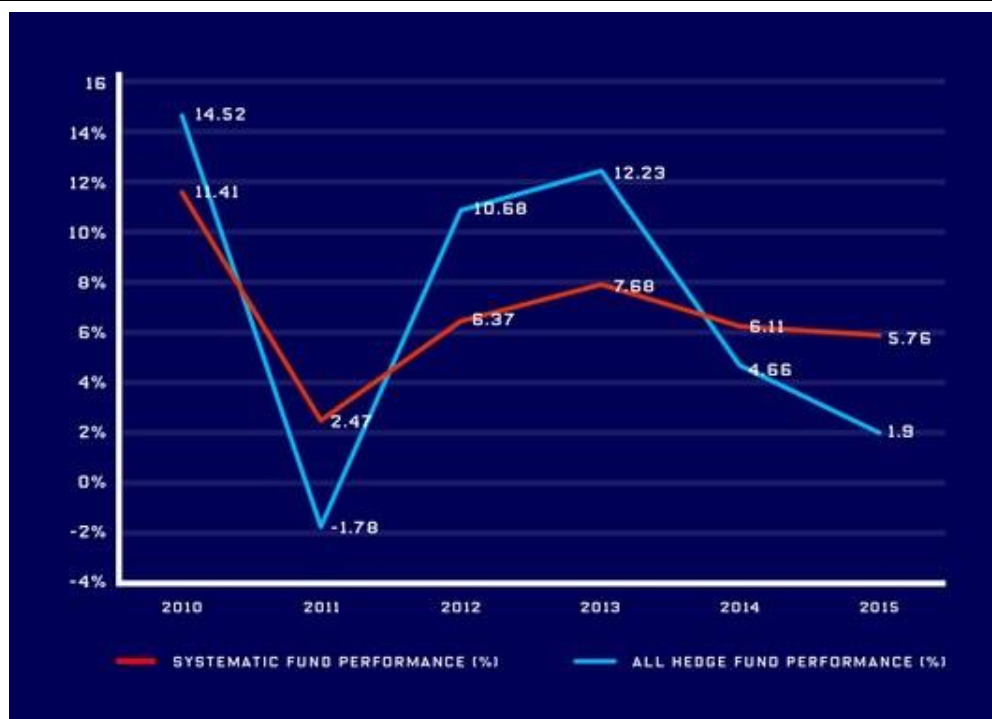
These developments have spurred predictions that “[AI Will Invade Every Corner of Wall Street](#).” Prognosticators see a world in which computers completely replace human investors.

“If computing power and data generation keep growing at the current rate, then machine learning could be involved in 99 percent of investment management in 25 years,” Luke Ellis, CEO of fund management company Man Group, PLC, told Bloomberg.

Despite this optimism, advances in artificial intelligence have not yet translated to superior returns. According to *Wired*, quant funds over the past few years have, on average, [failed to outperform hedge funds](#) (which have themselves failed to outperform the market).

Most people do not understand that AI, especially the AI used in finance today, lacks the application of deep subject matter expertise to create the clean data and relationships that are the foundation of any successful investment strategy or AI. Winning games is one thing, but the real world is not a game that follows immutable rules in a strictly defined space. In the real world, humans change the rules, break the rules, or the rules don’t even exist. Current AI is nowhere near navigating real world situations without a great deal of human intervention.

Figure 1: AI Is Overhyped and Misunderstood: Systematic Funds Underperform



Sources: [Pregin/Wired](#)



Finding the Talent(s)

One of the biggest problems with AI today is lack of interest or ability of those with adequate subject matter expertise to communicate with the programmers building the AI. The programmers don't understand the data they're feeding into their AI, and the analysts lack the understanding of the technology to communicate what programmers need to know to understand the source data and interpret the results.

This disconnect creates a number of well-publicized issues for the application of AI in finance and investing:

- Most AI firms end up [spending the bulk of their resources on data management](#) and data scrubbing rather than technology.
- Machines often discover [spurious correlations that don't work](#), or only worked in the past but aren't applicable in the future.
- Many AI systems turn into "[black boxes](#)" that spit out investment recommendations with no explicable basis or strategy. If the AI cannot articulate to humans how it "thinks," then how can investors trust it with significant sums of money?

Individuals with the skills and knowledge to bridge this divide are among the scarcest and most valuable people in finance. [Nine out of 10 financial services firms](#) have already started working on AI technologies, and they're all competing in this scarce labor pool.

As we wrote in "[Big Banks Will Win the Fintech Revolution](#)," the largest financial firms will be the biggest beneficiaries of technological advancements due to their scale and resources. Big banks can afford to pay the most for AI talent, and they have the biggest store of financial data to aid their new programmers.

A few banks are already making serious efforts to get the necessary talent. UBS (UBS) is on an [AI hiring spree](#), while Morgan Stanley's (MS) programmers and financial advisors have worked together to build "[Next Best Action](#)", a platform that uses machine learning to aid its advisors in offering personalized advice to clients.

These efforts should eventually pay off in a big way, but for now they remain in their infancy. Financial institutions still have a long way to go before they can truly implement AI in an effective way.

The Big (Data) Problem with AI

The total amount of digital data in the world [doubles every two years](#). As the volume of data grows exponentially, most of that data lacks the structures needed for machines to analyze it. As a result, AI projects, which are supposed to reduce the need for human labor, [require countless man-hours](#) to collect, scrub, and format data inputs.

Virtova founder, Sultan Meghji, told the Financial Revolutionists that many [AI startups spend at least half](#) their funding on data cleanup and management. Everyone wants to talk about teaching computers to think, but there's no short cut or substitute for curating the data sets that machines use to learn.

To train an AI, you need a training data set for it to learn from. Training data sets tend to be of two kinds. First, you have relatively small, accurate data sets that don't contain enough different kinds of examples to be effective. AI trained on these data sets become great at interpreting the training data, but they can't handle the variety and vagaries of the real world.

Other training sets are large but not very accurate. In these case, the AI gets to see lots of examples, sometimes with incorrect data, but it isn't being given clear and consistent instructions on how to respond. AI trained on these larger, inaccurate data sets often determine that there are few consistent things to be learned from the data and are capable of doing very little on their own.

For successful machine learning, training data sets need to be both accurate and widely representative. In other words, the training data needs to accurately represent what happens in as much of the real world as possible. How else can we expect the machine to learn anything consistently useful?

Herein is the AI challenge: machines can't learn without good training data sets and creating good training data sets requires more time than most realize from humans with deep subject matter expertise. Most humans with the depth of subject matter expertise required to curate a good training data set are not interested in such mundane work. An alternative approach is to have lots of humans with limited subject-matter-expertise do the work, but this approach has been [unsuccessful](#) so far.



The Big (Data) Problems Are Worse in the Finance & Investing World

In theory, curating training data sets should be less challenging in finance. After all, financial data is structured in the form of financial statements in official filings with the SEC. However, any layman can quickly see that there is not as much structure (humans do not always follow the rules) as one might presume in these filings. Plus, the structure that does exist is not all that useful for AI. In fact, it can be actively harmful.

Imagine a computer that wants to compare the financials of Coca-Cola (KO) and Pepsi (PEP). As the computer reads through the financial statements, how is it supposed to know that “Equity Method Investments” for KO and “Investments in Noncontrolled Affiliates” for PEP are the same? What about “Retained Earnings” vs. “Reinvested Earnings.” Industry groups have been trying to create a [standardized financial nomenclature](#) for years to solve this very problem.

In theory, the development of XBRL would solve this problem. In practice, [XBRL still contains too many errors and custom tags](#) to allow for fully automated reading of financial filings. Even the smartest machines need extensive training from humans with deep subject-matter expertise to be able to understand financial filings.

Without this pairing of sophisticated technology and expert analysts, any AI effort in finance is doomed to failure. As the saying goes, “garbage in, garbage out.” Dumping a bunch of unstructured, unverified data into a computer and expecting it to deliver an investment strategy is like dumping the contents of your pantry into the oven and expecting it to bake a pie. It doesn’t matter how good the machine is, it can’t function without the right preparation.

The Problem of False Positives

Even if the financial data is structured and verified, it may not be useful to a machine, and AI will struggle to tell what data is useful and what is not. The large volume of available financial data means there will inevitably be a large number of apparent patterns that are actually the result of pure randomness. This phenomenon is known as “overfitting,” and it’s such a recognized issue that it gets [its own lesson](#) in Stanford’s online course on machine learning.

Overfitting is not just an AI problem. Humans have always struggled with seeing patterns where none truly exist (heuristics). At least, though, we can be conscious of this flaw and try to counteract it. Computers, for all their sophistication, cannot claim this same level of consciousness. When programmers design machines to find patterns, that’s what those machines are going to do.

As AI gets more complex, the problem of overfitting gets worse. Anthony Ledford, the chief data scientist at one of Man Group’s quant funds, recently [told The Wall Street Journal](#):

“The more complicated your model the better it is at explaining the data you use for training and the less good it is about explaining the data in the future.”

Many quant funds today are simply mining patterns from past data and hoping those patterns persist into the future. In reality, most of those patterns were either the result of randomness or conditions that no longer exist.

Again, we see the need for the pairing of AI with human intelligence. Machines can process data and find patterns more quickly and efficiently than any human, but for now they lack the intelligence to audit those patterns and understand whether or not they can be used to predict future results.

AI As a Black Box

Of course, to audit the results of AI, humans need to be able to understand how that AI thinks. They need some level of insight into the processes the machine is using and the patterns it discovers.

Right now, most AI is not transparent enough for potential users to trust it. All too often, the AI algorithms are a black box that take in data and spit out results without any transparency into their underlying machinations.

In part, this problem is unavoidable if we want the machines to operate with the scale needed for them to be useful. The code that goes into AI is so complex that few individuals could ever fully understand its inner workings.

In fact, software doesn’t even have to reach the complexity of AI to have these problems. Consider the unexpected acceleration problems that plagued the Toyota Camry about 10 years ago. So many programmers



had worked on the engine control software that it turned into “[spaghetti code](#),” a mass of unintelligible and often contradictory code that no one understood and caused great harm.

If the software to support human control of a car’s breaking and acceleration can become so complex, just imagine how much more confusing and susceptible to errors more sophisticated activities, like financial modeling, can be. One mistake in one line of code could alter the entire function of the system. The software wouldn’t break, it would just be performing a different task than intended without anyone realizing until, perhaps, it’s too late.

This problem is exacerbated by the divide between the people with adequate subject matter expertise in finance and the programmers. The finance experts don’t understand how the software works, while the programmers don’t understand how finance works.

Finance is far from the only sector to experience this problem. In “[The Coming Software Apocalypse](#),” *The Atlantic* detailed several examples of major failures that occurred because the coders didn’t properly anticipate all the potential uses of their software. These failures were prolonged because the people using the code didn’t have any idea how it worked.

As long as AI remains a black box, its utility will be limited. Eventually, the lack of transparency will lead to a significant and undetected failure. Even before that point, it will be difficult to get investors to commit significant money to a program they cannot trust.

The Way Forward

For all these challenges, AI will continue to expand its reach on Wall Street. There’s no other way for financial firms to meet the dual mandate of reducing costs and improving their service. Technology is the only solution for analyzing the huge volumes of corporate financial data filed with the SEC every hour and meeting the Fiduciary Duty of Care.

The firms that understand this fact and take concrete steps to invest in technology will have a significant advantage over their competitors, which is why [UBS](#) and [Morgan Stanley](#) are among our top picks in the financial sector.

This article is the first in a five-part series on the role of AI in finance. Over the next two articles, we will dig deeper into the challenges facing AI and how they can be overcome, while the last two articles will show how AI can lead to significant benefits for both financial firms and their customers.

This article originally published on [January 11, 2018](#).

Disclosure: David Trainer and Sam McBride receive no compensation to write about any specific stock, sector, style, or theme.

Follow us on [Twitter](#), [Facebook](#), [LinkedIn](#), and [StockTwits](#) for real-time alerts on all our research.



New Constructs® - Research to Fulfill the Fiduciary Duty of Care

Ratings & screeners on 3000 stocks, 450 ETFs and 7000 mutual funds help you make prudent investment decisions.

New Constructs leverages the latest in machine learning to analyze structured and unstructured financial data with unrivaled speed and accuracy. The firm's forensic accounting experts work alongside engineers to develop proprietary NLP libraries and financial models. Our investment ratings are based on the best fundamental data in the business for stocks, ETFs and mutual funds. Clients include many of the top hedge funds, mutual funds and wealth management firms. David Trainer, the firm's CEO, is regularly featured in the media as a thought leader on the fiduciary duty of care, earnings quality, valuation and investment strategy.

To fulfill the Duty of Care, research should be:

1. **Comprehensive** - All relevant publicly-available (e.g. 10-Ks and 10-Qs) information has been diligently reviewed, including footnotes and the management discussion & analysis (MD&A).
2. **Un-conflicted** - Clients deserve unbiased research.
3. **Transparent** - Advisors should be able to show how the analysis was performed and the data behind it.
4. **Relevant** - Empirical evidence must provide [tangible, quantifiable correlation](#) to stock, ETF or mutual fund performance.

Value Investing 2.0: Diligence Matters: Technology is Key to Value Investing With Scale

Accounting data is only the beginning of fundamental research. It must be translated into economic earnings to truly understand profitability and valuation. This translation requires deep analysis of footnotes and the MD&A, a process that our [robo-analyst technology](#) empowers us to perform for thousands of stocks, ETFs and mutual funds.



DISCLOSURES

New Constructs®, LLC (together with any subsidiaries and/or affiliates, "New Constructs") is an independent organization with no management ties to the companies it covers. None of the members of New Constructs' management team or the management team of any New Constructs' affiliate holds a seat on the Board of Directors of any of the companies New Constructs covers. New Constructs does not perform any investment or merchant banking functions and does not operate a trading desk.

New Constructs' Stock Ownership Policy prevents any of its employees or managers from engaging in Insider Trading and restricts any trading whereby an employee may exploit inside information regarding our stock research. In addition, employees and managers of the company are bound by a code of ethics that restricts them from purchasing or selling a security that they know or should have known was under consideration for inclusion in a New Constructs report nor may they purchase or sell a security for the first 15 days after New Constructs issues a report on that security.

DISCLAIMERS

The information and opinions presented in this report are provided to you for information purposes only and are not to be used or considered as an offer or solicitation of an offer to buy or sell securities or other financial instruments. New Constructs has not taken any steps to ensure that the securities referred to in this report are suitable for any particular investor and nothing in this report constitutes investment, legal, accounting or tax advice. This report includes general information that does not take into account your individual circumstance, financial situation or needs, nor does it represent a personal recommendation to you. The investments or services contained or referred to in this report may not be suitable for you and it is recommended that you consult an independent investment advisor if you are in doubt about any such investments or investment services.

Information and opinions presented in this report have been obtained or derived from sources believed by New Constructs to be reliable, but New Constructs makes no representation as to their accuracy, authority, usefulness, reliability, timeliness or completeness. New Constructs accepts no liability for loss arising from the use of the information presented in this report, and New Constructs makes no warranty as to results that may be obtained from the information presented in this report. Past performance should not be taken as an indication or guarantee of future performance, and no representation or warranty, express or implied, is made regarding future performance. Information and opinions contained in this report reflect a judgment at its original date of publication by New Constructs and are subject to change without notice. New Constructs may have issued, and may in the future issue, other reports that are inconsistent with, and reach different conclusions from, the information presented in this report. Those reports reflect the different assumptions, views and analytical methods of the analysts who prepared them and New Constructs is under no obligation to insure that such other reports are brought to the attention of any recipient of this report.

New Constructs' reports are intended for distribution to its professional and institutional investor customers. Recipients who are not professionals or institutional investor customers of New Constructs should seek the advice of their independent financial advisor prior to making any investment decision or for any necessary explanation of its contents.

This report is not directed to, or intended for distribution to or use by, any person or entity who is a citizen or resident of or located in any locality, state, country or jurisdiction where such distribution, publication, availability or use would be contrary to law or regulation or which would be subject New Constructs to any registration or licensing requirement within such jurisdiction.

This report may provide the addresses of websites. Except to the extent to which the report refers to New Constructs own website material, New Constructs has not reviewed the linked site and takes no responsibility for the content therein. Such address or hyperlink (including addresses or hyperlinks to New Constructs own website material) is provided solely for your convenience and the information and content of the linked site do not in any way form part of this report. Accessing such websites or following such hyperlink through this report shall be at your own risk.

All material in this report is the property of, and under copyright, of New Constructs. None of the contents, nor any copy of it, may be altered in any way, copied, or distributed or transmitted to any other party without the prior express written consent of New Constructs. All trademarks, service marks and logos used in this report are trademarks or service marks or registered trademarks or service marks of New Constructs. Copyright New Constructs, LLC 2003 through the present date. All rights reserved.