

A Budding Romance: Finance and AI

Xiao-Ping (Steven) Zhang
David Kedmey
EidoSearch

Enthusiasm for artificial intelligence and multimedia information in the financial industry is at an all time high. Every leader in finance now feels the pressure to answer the question, “*What is your AI strategy?*” Start-ups are playing a key role in helping the financial sector determine

what AI can do and how humans and machines can work together. In this essay, we describe emerging trends and attempts by FinTech start-ups to apply AI and multimedia information processing techniques across a wide range of business needs.

Enthusiasm for artificial intelligence in the financial industry is at an all-time high.¹ The trend began approximately ten years ago when a small cohort of startups sensed an opportunity to apply machine learning and multimedia processing to finance. It began with an explosion of digitized multimedia data and cheap computing power—driving forces that are still underway to this day.

These conditions provided fertile ground for entrepreneurs who imagined entirely new and automated workflows. They envisioned products that combined domain expertise in finance with knowledge of machine learning.

Awareness of this trend spread slowly, lurking under the radar, until grassroots activity crossed a threshold with the spontaneous arrival of its own name: “FinTech” (financial technology) entered our lexicon. The inaugural class of the FinTech Innovation Lab² in New York City in 2011 was an auspicious introduction. According to CBInsights, global FinTech deals in 2018 are on pace for a record year, reaching \$20.3 billion in VC-backed equity funding in the second quarter alone.³

In the last several years, a surge of interest in the promise of AI joined up with FinTech, kicking off a new phase of exponential growth. From a technologist’s perspective, long-standing machine learning techniques were simply folded into the catchall phrase of AI. But then a major cultural shift occurred. The breakthrough successes of deep learning in multimedia, i.e., speech and image recognition, and AlphaGo, captured the imagination of people in finance and the general public.

Every leader in finance now feels the pressure to answer the question, “*What is your AI strategy?*” To provide answers, major resources are flowing to data science and AI groups within institutions and to the FinTech startups looking to serve them. Strategy depends on what aspect of the business you are addressing and the specific problem you are trying to solve. Financial institutions, like institutions in

any other industry, seek efficiencies in their operations: communications, customer service, human resource management, regulatory compliance, and fraud detection.

In this essay, we describe attempts by FinTech start-ups to automate aspects of business operations. We also address an intriguing problem that is in some ways unique to finance: the challenge of prediction in financial markets. The central challenge here—in a highly noisy and nonstationary system—is finding regularities in data that emerge from millions of traders and investors reacting to one another's decisions.

UNPACKING THE FINANCIAL INDUSTRY

Modern economies depend on a thriving financial sector, which in the U.S. accounts for 20 percent of GDP. Tens of thousands of banks, mutual funds, and hedge funds employ a vast army of financial professionals and technologists. This complex ecosystem serves one ultimate purpose that is to facilitate the flow of resources throughout the entire economy.

To understand the financial ecosystem, it helps to divide activities into two main functions. The brain of the ecosystem is an information-processing engine that detects where resources in the economy reside and where they should go. The body of the ecosystem is the infrastructure that supports this transfer. Below we organize the discussion by specific problems in operations (the body of the industry) and in making market predictions (the distributed brain) and the FinTech opportunities to address these challenges using AI.

A WORLD OF MULTIMEDIA SIGNALS

The financial industry continuously appraises signals in multimedia data: words spoken in audio and text recordings, consumer activity in the form of digital footprints, and economic dynamics from sensors and satellites that capture movement of goods and people. Processing this datum from raw form to actionable insight engages all of our multimedia techniques to help financial firms find relevant signals hiding in our midst. The proliferation of affordable drones and satellites is generating geospatial imagery that the company **Orbital Insight** processes to monitor and track global oil inventories—a key determinant of oil prices. Foursquare entered the arena of the Wall Street analysts several years ago, by translating its foot traffic data to accurately predict a new threshold of iPhone sales volumes. Foursquare was also able to anticipate plummeting sales at Chipotle.⁴ With daily check-ins at 8 million locations—a total of 12 billion over the past nine years—traditional prognosticators must adopt new forms of data cleaning, automation, and analysis to keep up.⁵

Entirely new ecosystems of data vendors and data analyzers are emerging. **ExtractAlpha**, partnering with alternative data vendors such as **alpha-DNA**, processes data for use in systematic money management. Online consumer behavior information across multiple websites and search and social media platforms is compiled and organized by alpha-DNA in near real-time, and transformed by ExtractAlpha into The Digital Revenue Signal, a stock selection score designed to forecast revenue surprises based on changes in consumer demand. And, lest you think politics has avoided algorithmic scrutiny, new measures of company behavior in the environmental, social, and governance domain are being devised and modeled.

The hunt is on to find out what information these datasets hold to improve decision-making and which new sources of data can be brought into the fold. AI experts, as well as the multimedia techniques we have applied in so many other domains, must play a central role in finance.

MERGING HUMANS AND BOTS

AI, machine learning, and natural language processing cannot replace people—at least not yet. But humans and bots can work together in harmony to automate tedious tasks and enhance the human connection between business professionals and their customers. Let us consider three examples in the financial industry: wealth management, banking, and research.

Wealth management. The impact of AI in wealth management will not be wealth creation for the typical investor. Sorry to disappoint. In the article *To the Victor Go the Spoils: AI in Financial Markets*, it is argued from first principles that the “vast majority of people—no matter their level of AI expertise—will not achieve large excess returns.”⁶ However, AI does bring efficiencies to operations. Wealth managers can serve more clients in less time and, in an increasingly digitized world, still deliver a personalized experience. To answer clients’ questions and make informed recommendations, wealth managers must consider hundreds of data points: a client’s investment portfolio holdings and transactions, trending market events, unexpected personal events, and evolving client needs. Start-ups such as **Forwardlane** address this need via AI, which prioritizes client Q&A from a natural language conversation interface, which helps wealth managers find precise answers to client questions.

Banking. Similar to wealth management, customers demand convenience and speed for banking services and loathe the time spent searching for answers. How can banks offer these services at scale? A start-up at the forefront in this domain is **Kasisto**, which handles frequent banking tasks and, in a CoBot-like approach, knows when to hand-off to a live agent for services that need a human touch. The goal is to power “human-like conversations” through a conversational AI platform fluent in finance. Banks are eager to deploy these types of intelligent virtual assistants that can care for customers at a fraction of the cost and reduce call center volume.

Research. In order to assess the value of stocks, analysts scour reams of unstructured news and structured data that streams in from the web and accumulates within an organization. This human activity of processing data itself generates information, specifically, trails of search activity, discussions over email, and written reports. Machine intelligence algorithms that watch what a research analyst is reading and writing can track these unfolding events, offering fertile ground for Bots to learn and make recommendations to analysts of what to research. For example, the start-up **Diffeo** has collaborative agents that on their own might not know where to look, but by watching research teams do their work can help direct research efforts. One application is to discover entity connections between companies and people that emerge from such activity. Backed up by contextual evidence, these connections could be critical in assessing the likelihood of important events in the life of a company, for example, an impending merger or acquisition.

DETECTING BAD BEHAVIOR

Financial institutions are highly regulated and are expected to scrutinize the knowledge and intentions of their traders. The survival and success of a financial institution, regardless of size, comes down in large part to effective risk mitigation. This requires meaningful alerts and noise reduction for false alarms through automatic analysis of communications data. These efforts are costly and require technology to operate at scale. McKinsey & Company estimates that CEO fraud, where imposters gain access to business email accounts and fool unsuspecting employees to send funds to bogus accounts has led to losses of more than \$2.3 billion over the past few years.⁷

Business leaders are held responsible for the actions of their employees’ actions. Take the following example given by the FinTech start-up, **Digital Reasoning**. When one trader messages “*Let’s take this offline*” to another, is this an innocuous request to advance a meeting agenda, or an attempt to collude in secrecy? It depends on semantics and context. Using a private knowledge graph and natural language understanding technology, digital reasoning helps financial firms combine analysis of behaviors, intentions, and emotions. By transforming multimedia communications data, including text and audio, firms attempt to identify threats and mitigate reputational risk.

INSURANCE AND AI

Insurance companies have been in the business of evaluating risk for centuries. Yet many businesses in need of insurance find the process of obtaining quotes painfully inefficient. McKinsey & Company forecasts that “in 2030, manual underwriting ceases to exist for most personal and small-business products across life and property and casualty insurance.”⁸ We are not there yet. Each year insurance underwriters receive applications from over 7.5 million small-medium sized businesses that fail to get an automatic quote. To solve this bottleneck, start-ups are ingesting data sources of a more granular

nature and applying machine learning techniques to permit analysis à la carte by geography and business classes. For example, **Open Data Nation** (ODN) aggregates billions of records published by city governments about commercial businesses and individual behaviors and builds machine learning models to anticipate issues. Insurance underwriters can then query ODN for an on-demand risk score better tailored to the unique risks of each business applicant.

USE MACHINES TO HIRE HUMANS

The most valuable asset of a financial firm is its workforce, so there is no more important process than the way a company attracts, selects, and retains talent. **Pymetrics** is a start-up that addresses this problem through the application of neuroscience games, which gather more insight about candidates and feeds this information through machine learning algorithms to increase the efficiency of hiring. More than ever before, financial institutions are in competition for tech talent. The firm **Untapt** is supplying AI algorithms to supplement human recruiters that review tech resumes. Using natural language processing, words on a resume are mapped relative to each other, each resume is sorted in a space with other resumes and their engine uses feedback to distinguish matches and nonmatches. By process millions of job scenarios, these algorithms learn to identify what qualities make for a successful candidate.

MARKET PREDICTION

The lifeblood of finance is information. Knowledge that provides more accurate and more *informative* predictions in financial markets translates directly into profit. Markets, however, are inherently noisy—much more so than traditional domains for AI. Changing relationships in markets pose unique challenges for AI researchers and, therefore, require unique solutions.

In addition to the inherent challenge of market prediction, data analysis needs have grown exponentially with the rise of alternative data. Investment banks and hedge funds are building new teams of data scientists to clean, structure, and analyze the fire hose of data streaming from every corner of the economy. Hundreds of alternative data start-ups are offering this new content. The list is long. It includes sentiment from social media, mobile data content, online reviews, and web searches; transaction data from e-commerce and credit cards; and new data from sensors such as satellite and geolocation data.

The basic question money managers are trying to answer is “*what and how much can traditional and alternative data tell us about the future?*” Also, an equally important question required to generate profit: “*Is it unique? Does the information tell me something about the markets that others do not know?*”

Helping to answer these questions, the start-up **EidoSearch** is a probability intelligence company that created a new type of AI to systematize the investment process and quantify prediction uncertainty.¹ Its numeric search engine finds conditions in data through a technique called data-incident based modeling, which takes advantage of multimedia signal processing and content-based retrieval technologies and is uniquely suited for nonstationary systems such as financial markets. Where deep learning and other forms of machine learning have fallen short, the EidoSearch method jettisons the need for a model with functional form. Instead, current events are automatically matched to similar data incidents, and their associated outcomes are used to generate a dynamic, model-free distribution forecast. This method, subjected to historical testing, has enabled hedge funds to detect new sources of profit. Each scenario tested is a “ProBot,” which is a probability forecasting robot. Also, new evaluation measures for accuracy and informational uniqueness have been developed to select the most skilled (and profitable) forecasters among the ProBots.

PLANNING FOR THE UNFORESEEN

The growth of knowledge is inherently unpredictable, and awareness of what is unforeseen needs to be considered. The financial industry has great enthusiasm for AI and is wisely investing for the long term by bringing in experts from academia and the tech industry to help lead the way. These new hires are

critical points of contact for start-ups to market their AI services and to work collaboratively with in-house teams to both define and solve the pressing problems named in this essay. Using state-of-the-art techniques is imperative. Two recent examples: Dr. Li Deng, a former chief scientist of AI at Microsoft, was recently hired as a Chief AI Officer at Citadel, one of the largest hedge funds in the world; and Dr. Manuela Veloso, on leave from Carnegie Mellon University, where she was the Head of the Machine Learning Department, recently joined the world's largest financial institution, JP Morgan Chase, to create and head an AI Research Center.

Dr. Deng is a leader in the speech recognition industry using large-scale deep learning—the successes of which served as a major impetus for the massive wave of interest in AI. Dr. Veloso, coming from the world of autonomous robots, is a particularly interesting and revealing choice. A CoBot, as her group defines it, is a robot that “*follows a novel symbiotic autonomy, in which the robots are aware of their perceptual, physical, and reasoning limitations and proactively ask for help from humans.*”⁶

How is this relevant to finance? Replace the concept of navigating the physical world with the demands of navigating information flows within finance, i.e., a “virtual” overlay of our physical economy. New paths ahead in the world of finance are being forged—in concurrence with creative start-ups—as humans and machines learn to work symbiotically.

REFERENCES

1. M. Kolanovic and R. T. Krishnamachari, “Big data and AI strategies: Machine learning and alternative data approach to investing,” JP Morgan Chase & Co., New York, NY, USA, Tech. Rep., May 2017.
2. FinTech Innovation Lab. [Online]. Available: <http://www.fintechinnovationlab.com/>
3. CBInsights, “Global Fintech report Q2 2018,” New York, NY, USA, 2018.
4. J. Wiecezner, “Foursquare just predicted Chipotle’s sales Will plummet 30%,” *Fortune*, Apr. 15, 2016.
5. E. Stinson, “Foursquare may have grown up, but the check-in still matters,” *Wired*, Aug. 9, 2017.
6. X.-P. Zhang, “To the victor go the spoils: AI in financial markets,” *IEEE Signal Process. Mag.*, vol. 34, no. 6, pp. 171–176, Nov. 2017.
7. J. Corbo, C. Giovine, and C. Wigley, “Applying analytics in financial institutions’ fight against fraud,” McKinsey & Company, New York, NY, USA, Apr. 2017. [Online]. Available: <https://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/applying-analytics-in-financial-institutions-fight-against-fraud>.
8. R. Balasubramanian, A. Libarikian, and D. McElhane, “Insurance 2030—The impact of AI on the future of insurance,” McKinsey & Company, New York, NY, USA, Apr. 2018. [Online]. Available: <https://www.mckinsey.com/industries/financial-services/our-insights/insurance-2030-the-impact-of-ai-on-the-future-of-insurance>.
9. CoBot Robots. [Online]. Available: <http://www.cs.cmu.edu/~coral/projects/cobot/>

ABOUT THE AUTHORS

Xiao-Ping (Steven) Zhang is a Co-Founder of EidoSearch. Contact him at xzhang@eidosearch.com.

David Kedmey is a Co-Founder of EidoSearch. Contact him at dkedmey@eidosearch.com.