

Data Science in Banking and Finance: The Case for AI

Originally published at <https://www.iflexion.com/blog/data-science-banking-and-finance>

Published: September 9, 2020 | By Martin Anderson



Understanding complex and interdependent financial systems has been one of the most tasking pursuits in human endeavor: the sheer volume of data involved, the administrative burden of maintaining and processing historical records, the high resources needed to exploit and make sense of that ever-growing lake of past data... all have conspired to make the financial sector arguably the most obvious beneficiary of [data science consulting](#) and the new revolution in machine learning.


Over the last decade, as GPU-driven machine learning systems and cloud-based processing power have reignited business engagement with these technologies, this most cautious of sectors has begun to adopt automation and AI-powered insights across a broad swathe of its portfolio of activities.

In this article, we'll examine some specific fields where data-fueled AI is making a difference in the banking and finance sectors, as well as taking a look at the nature of the challenges involved.


This is how data-fueled AI is making a difference in the banking and finance sectors.

Confidence in AI and Machine Learning for the Financial Sector

According to a recent report¹, industry respondents intend to increase their spending on AI/ML systems in the next 12 months by 62%, with a further 83% declaring that the development and deployment of such systems would form a core part of their business strategy in the near future.




Perhaps in line with the increased risks of buying into AI startups' offering², 77% of firms questioned intended to develop or augment their own internal AI teams, while only 38% planned to engage third-party services, either for a complete deployment solution or for specific provisioning of one section of a solution.



A report from the Bank of England in October 2019³ surmises that a representative majority of UK finance firms are utilizing [machine learning technologies](#), with most deployments now past the phase of initial testing. Firms surveyed report notable efficiency gains, enabling more customized products and greatly improved fraud detection, among other benefits.

The majority of surveyed firms across five sectors have already deployed machine learning in initial trials.



Currently, most of these deployments support customer relationships rather than implementing high-level automation for complex interactions and processes.

Some of the most advanced use cases involve the automation or semi-automation of anti-money laundering (AML) processes and checks, while others are increasing the extent to which machine learning contributes to insurance underwriting and claim processing.


The report additionally notes that the locus of AI/ML implementations is moving from back-office functions to customer-facing applications, particularly in market risk assessment, insurance underwriting, and credit risk.

The Financial Sector Engaging with Automation

In a report released in August 2020, the US Securities and Exchange Commission (USSEC) noted the 'benefits and risks' of algorithmic trading, and the increasing use of machine learning systems in highly data-driven equity markets. The report concluded that ongoing vigilance is necessary to oversee the autonomy and scope of AI-driven trading algorithms⁴.

According to USSEC, this circumspect attitude is necessary in the light of events such as the 'flash crash' of May 6, 2010⁵, when a trading algorithm dumped 75,000 S&P500 futures contracts, causing the biggest stock plummet in decades. The eventual inquiry findings observed that *'the interaction between automated execution programs and algorithmic trading strategies can quickly erode liquidity and result in disorderly markets'*⁶.

Within the UK banking sector, the most mature machine learning deployments are found in the risk management and compliance sector, with customer engagement close behind and treasury usage the most cautious among the sectors.



Over 57% of the surveyed firms using machine learning have based their models on pre-existing risk management templates, while several are establishing ethical principles to address AI explainability and 'model drift' (see below).

Meanwhile, Bloomberg reports that finance-based positions requiring AI skills increased by 60% in the period of 2018-2019⁷.

All these findings acknowledge general concern about the changing regulatory environment, and growing concerns around ethical AI. However, the biggest obstacle to rapid adoption remains the risk of creating autonomous AI systems that lack explainability and accountability.

The biggest obstacle to rapid adoption of AI in finance remains the risk of creating autonomous AI systems that lack explainability and accountability.

Conceptual Challenges in Developing AI for the Finance Sector

For obvious reasons, [black-box AI](#) is problematic for critical deployments in the finance sector. There is industry concern around explainability, model validation, and model drift⁸, where long-term fluctuations in data can make a machine learning model unstable.

When COVID-19 began to significantly affect the economy in early 2020, a number of high-level commercial and finance-related machine learning systems found their definitions of 'spike', 'surge' and 'decline' so radically redefined that it was necessary in many cases to manually intervene⁹ and

account for the extraordinary new trends in the data.

Achieving 'elasticity without eccentricity' is a notable challenge for [AI applications](#) across the financial sector. Sometimes, as with COVID-19, the data skews so wildly from mission parameters that the model's accuracy can become severely compromised.

To further complicate the challenge, even a highly adaptable model may not always pick the best reaction to 'unexpected' data, depending on its design, scope and variables.

How a Fraud Detection Model Might Interpret a Change in Customer Behavior

For example, consider the hypothetical case of a 54-year-old surgeon who, having read that playing videogames is helpful to maintaining dexterity for surgeons¹⁰, buys a games console and attempts to purchase a downloadable game from an online store via the console.

This transaction stands out in stark relief from the surgeon's customary spending patterns and far outside her demographic profile. Therefore, the bank immediately freezes her card and requires further authentication that the transaction is genuine, perhaps via a live phone call or a smartphone-based authentication system.

After this, the surgeon downloads her game, and all is well. In future, she can probably buy more videogames from the online store without further interruptions.

But in what ways might a machine learning system interpret this event, in terms of protecting the surgeon in future?

- **Singular anomaly**

In the most cautious scenario, the model could add 'Buys PlayStation Games Online' to its understanding of the client's spending behavior. This may trigger additional freezes if the surgeon buys a game on any other platform, further impeding the customer experience. It's a granular solution that covers the incident but doesn't advance the model's flexibility, insight or autonomy.

- **Minor demographic extension**

Alternately, the AI might characterize the customer as more generally interested in video games and permit a slightly wider range of related purchases. Though this increases the attack area, it's a small increase and an acceptable and informed compromise between customer experience and customer security.

- **Large demographic extension**

What happens if the model is comparing the surgeon to typical histories within her age/status demographic? Having seen a number of people across a historical customer base pay off their mortgages, empty their nests, and indulge in a little mid-life atavism, the AI may begin to expect further such 'indulgent' purchases that are more strongly associated with a younger age-group and socio-economic status. This could make the AI more permissive in terms of allowing deviation from previous behavior, for instance with purchases traditionally associated with a younger age group. In such an eventuality, the attack surface is notably increased.

- **Baseline becomes renormalized**

In the case of a very poorly-configured model, this unexpected incursion from the 18-24 demographic into the 50-64 age range could cause a reassessment of the baseline expectations of the model, which will now flag purchases by the customer *that are typical of her demographic*.

- **Anomalies become acceptable**


By contrast, a less sophisticated model could deduce that the customer has become 'unpredictable', and begin to 'expect the unexpected' from her — probably the most disastrous result in terms of protecting her in the future.

Since these are all technically valid approaches to automated model revision, and since many other types of feedback-driven risk assessment models are susceptible to this kind of rogue logic, it's not difficult to understand the current cautious attitude to model validation and explainability in AI deployments for financial services.

Assessing Risk

From detecting fraudulent logins and transactions through to providing credit scores, insurance coverage and the anticipation of market fluctuations, nearly all data science applications in banking and finance are involved with some aspect of risk modelling. In the case of insurance, the application is clear — far clearer than the variables involved in calculating risk.

Yet, despite a more cautious adoption rate than the risk management and compliance sector, insurance companies that utilize AI to streamline and speed their underwriting processes have seen a 14% median increase in business over the last three years¹¹.



Another survey of [AI in insurance](#) asserts that over 90% of insurers plan to utilize or develop AI-driven processes to address and streamline claims and underwriting¹².

In certain cases, these new processes are being promoted as preventative or mitigating strategies for the end user: Tokyo-based insurance company Dai-ichi Life uses machine learning to aggregate the sum of a customer's habits (such as smoking) with their medical data to envision how their face will look later in life:



Improved Accuracy in Machine Learning Insurance Estimates

Insurance giant AXA has engaged significantly with machine learning over the last six years, gradually abandoning an explicitly Random Forest and Decision Tree-based approach in favor of an open-source supervised learning AI library from Microsoft¹³.

The AXA system implements a 'certainty index' that's lacking in traditional machine learning approaches. In research circles, this particular challenge is known as Regression Conformal Prediction, wherein a machine learning system will associate an 'accuracy probability' with its own predictions, instead of simply outputting the predictions.

This is a useful step away from black-box AI methodologies, in that a supporting rationale is provided by a parallel logic, with the model continuously compared to and adjusted by real-world feedback (sometimes from human-managed review processes) on its performance and accuracy.

The technique is becoming increasingly capable of generating premium pricing estimates in line with comparable human estimations and without any human supervision. AXA observes that it has used the new approach to notably reduce the person-hours required for the traditional underwriting process¹⁴.

Levels of Data Integrity and Trust in Risk Assessment

Financial AI systems must make use of three tiers of data, each with diminishing levels of accountability and trust.

- **Low risk: structured data**

This is likely to be first-party data developed by the company itself, or from actuarial and governmental sources. Such data is not always directly or ubiquitously accessible to the public (for instance, credit scores and information that the government may make available to privileged commercial sectors such as insurance and credit-scoring companies). This type of data will usually enter a machine learning process fully tagged and classified, with little human intervention and with the highest possible trust score. Such a data pipeline is closer to 'automation' of previous analogue methodologies rather than the 'intelligent autonomy' that has dominated headlines in recent years.

- **Medium risk: semi-structured data**

Here the data is gathered from external systems which have some structure (e.g., HTML, JSON) but lack the internal accountability of the first-party and privileged sources for structured data. The need for greater oversight and verification offsets the advantages of automation to a certain extent. Data in this tier could be scraped from websites or interpreted from documents that already exist as text (rather than requiring OCR). Sources will vary constantly in quality, and any document object model (DOM) on which the machine learning system might rely is subject to arbitrary change.

- **High risk: unstructured data**

Here the data is abstract and unstructured, such as a web-found image of an individual involved in risk assessment, wherein an external machine learning process such as [facial recognition](#) imposes meaning and classification onto a mass of pixels; or image-based text that must be transcribed through OCR and interpreted through [Natural Language Processing](#) (NLP). Though this ability to extract hidden relationships from abstracted data represents the most exciting potential of machine learning, in the case of critical financial applications it requires the highest level of oversight and verification by humans, indicating the limited potential for real-time AI-driven processes — at least for now.

The Bank of England report cited above asserts that structured data is used in more than 80% of machine learning use cases, while semi-structured and unstructured data are employed as secondary resources in two thirds of cases.

Wrapping Up

Because neural networks and other machine learning approaches are not always explicable or accountable, the financial sector is not in a position to pursue truly autonomous AI-driven processes at the same breakneck speed as many other industries. As noted in the Bank of England report, many companies are using machine learning systems to replicate their traditional analogue systems, rather than immediately exploring the new paradigms of data analysis offered by the latest research in artificial intelligence.

Nonetheless, the sector is currently experiencing great success and increased efficiency by incorporating robotic process automation (RPA) and explicable machine learning systems to automate administration, financial time series forecasting, and the gathering of intelligence for purposes of risk assessment.

The finance sector is also benefitting from less industry-specific machine learning trends, such as customer lifetime value (CLV) assessment, predictive analytics, and customer relationship management.

As customers ourselves, we can approve of this cautious approach, while looking forward to the greater benefits of higher-level machine learning systems for finance and banking, as firms gradually develop and deploy more sophisticated and proven analysis and risk assessment systems — with increasing amounts of autonomy being granted to the systems as they prove their reliability over time.

As AI systems prove their reliability in banking and finance over time, they are being granted with increasing amounts of autonomy.