accenture consulting

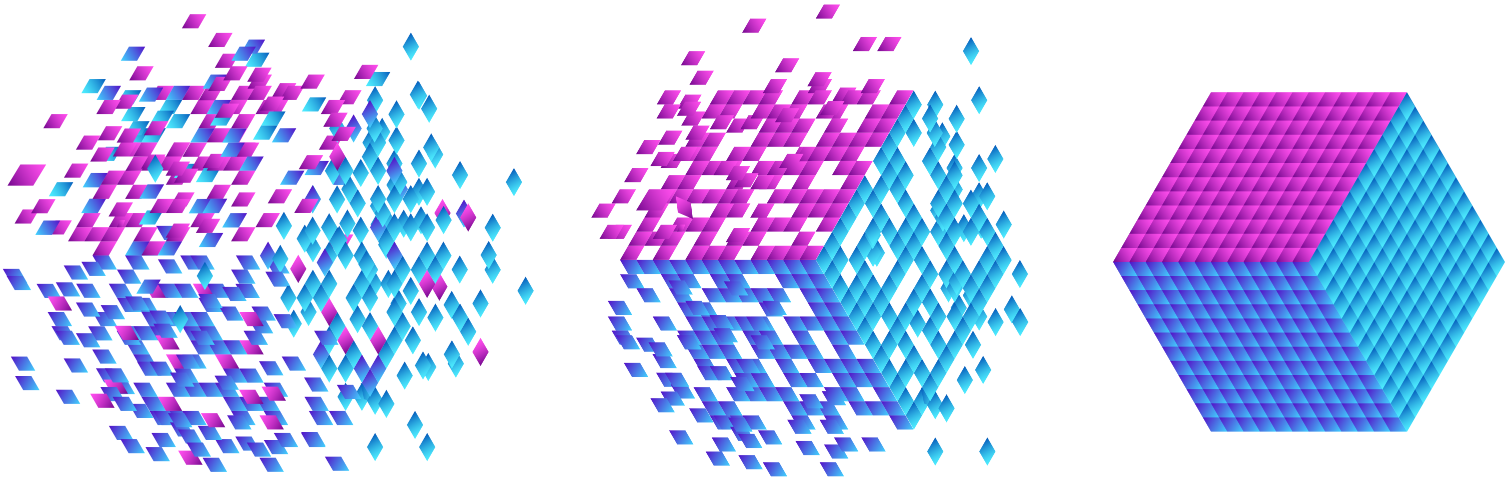

NEURAL NETWORKS

The Next Step for Artificial Intelligence in Financial Services

A perspective from Financial Services Technology Advisory

EXECUTIVE SUMMARY

The adoption of artificial intelligence (AI) across the financial services industry is advancing at an unprecedented pace.

As firms strive to transform their businesses for a digital world, realize efficiencies, improve the customer experience and revitalize their growth, they increasingly see AI and machine learning as vital enablers—and are investing in understanding how they can leverage them to greatest effect.

As machine learning becomes mainstream, the next wave of AI innovation in financial services is already emerging: artificial neural networks. Based on the multi-layered neuron structures that mimic the human brain, neural networks offer a step-change in the power of AI, opening up opportunities to automate ever more complex processes and decisions with the highest possible degree of accuracy.

Early AI proofs of concept (PoCs) involving neural networks implemented by financial services firms have yielded promising results. But while their potential is significant, they should be approached with care. Initially developed in academia, neural networks were designed to deliver the highest possible accuracy with little focus on explainability.

However, in regulated sectors like banking and insurance, where both regulators and customers often ask to know why a particular decision was made, an inability to explain AI reasoning can expose firms to risks ranging from legal challenges to a loss of customer trust.

The good news is that these risks can be addressed, and that—alongside neural networks—there are multiple machine learning options available. As a result, financial services firms can effectively harness the power of machine learning to solve some of the most complex problems they face. To do this, firms have to put in place appropriate processes, practices, tools and controls to make responsible and ethical use of these extremely effective capabilities and apply appropriate algorithms to the right processes in the appropriate way.

In this paper, we map out the various components and steps for an effective implementation of neural networks—the next stage of financial services' AI journey.

CONTEXT: THE EVOLVING LANDSCAPE OF AI IN FINANCIAL SERVICES

AI has come of age for financial services firms...

Across the global financial services industry, interest in the potential of AI is growing by the day. Organizations are eager to automate a growing range of processes across their operations, with a view to gaining several tangible business benefits—including lower costs, higher accuracy, an improved customer experience, and ultimately business transformation and competitive edge.

In search of these benefits, financial services firms have progressed in recent years from advanced statistical modeling to more complex algorithms and AI techniques, including combining robotic process automation (RPA) with cognitive AI technologies. To date, the industry's focus has been mainly on a major subset of AI: machine learning algorithms. These are a form of AI that can learn and improve automatically on the basis of experience—qualities that are well-suited to many types of decision-making in financial services.

AI in financial services: A rising tide of innovation —and savings

The momentum behind AI in financial services is now unstoppable. A recent survey by Intertrust Group B.V.[1] of 500 senior financial services decision-makers' views on disruptive technologies found that 77% think AI will play the biggest role in revolutionizing the industry over the coming five years, ahead of blockchain (56%) and robotics (27%). And financial services research house Autonomous Research LLP[2] estimates that financial services firms can expect AI to help reduce their operating expenses by about 20%, resulting in savings totaling US$1 trillion globally by 2030. Meanwhile, Accenture's own research[3] suggests that AI will add US$1.153 billion in value to the financial services industry by 2035.

...and neural networks are emerging as the next step

Having advanced from AI in general to machine learning, the industry's focus is now evolving further—to a specific subset of machine learning called neural networks. As the accompanying information panel explains, neural networks are highly complex machine learning algorithms that are structured as a number of interconnected layers of neurons. As well as reasoning taking place within each layer, the neurons within the different layers also interact and affect each other, creating a hugely powerful model for complex decision-making. The effect is that neural networks work in a way analogous to 3D, while simpler AI reasoning is more like 2D. This makes neural networks especially suited to complex "deep learning" applications that require the processing of massive amounts of data and high levels of domain expertise and judgment.

The power and wide applicability of neural networks and deep learning are seeing these technologies gain usage in financial services. This ongoing increase reflects factors including the availability of massive amounts of data (both internal and from third parties), ongoing academic advances, and algorithmic innovation supported by expanding computing power, especially on cloud platforms. These drivers have been supplemented by high-profile marketing and communications efforts from the global technology giants, further boosting awareness of AI and of neural networks in particular.

What are neural networks?

Artificial neural networks are inspired by the structure of the human brain. A typical neural network is made up of millions of artificial neurons arranged in a series of connected layers. During training of the model, the first layer serves as an example of input variables—such as documents, number or pixels from images—from which the model learns to create an association, with the output in the final layer, representing the target outcome. The '"hidden" layers in between the input and output layer convey the logic connecting the input to the output, cresting a network of reasoning that is much more complex than a relatively simpler linear algorithm.

A BROAD RANGE OF USE CASES

The attributes of neural networks—particularly their ability to handle highly complex decisions—makes them ideally suited to a wide variety of use cases in financial services.

The table below provides a snapshot of some of the common applications for neural networks in different sectors of the industry, across the three areas of predicting future value, extracting meaning from unstructured data such as language, and recognizing objects on images.

Table 1: Sample use cases for neural network in financial services

	Unlocking Value Prediction	Unlocking Unstructured Data	Unlocking Imagery
Banking	• Fraud detection for credit card transactions • Overdraft predictions based on the customer's transaction history	• Information retrieval from invoices to substantiate a transaction on business accounts • Verification of the plausibility of property prices from mortgage application by performing internet searches	• Using a facial recognition system to log into applications for corporate banking (already introduced by HSBC Holdings plc)[4] • Using satellite and street view images to verify the existence of a business as a part of know your customer (KYC) and anti-money laundering (AML) checks
Insurance	• By collecting and processing data from wearable devices, insurance companies can use neural networks to predict health problems and suggest lifestyle changes • Analyzing customers' interactions with the company to offer discounts to customers who wish to leave	• Accident severity evaluations based on the description in the claim • Assessing the level and type of risk associated with a customer based on their social media activity	• Car accident damage assessment • Risk prediction for home insurance based on images of the building and surrounding objects such as trees and rivers
Capital Markets	• Helping traders decide what price to quote when buying or selling bonds for their clients based on historic and real market data • Using electronic routing algorithms to find a match from stock brokers, stock exchanges or alternative trading systems that can fulfill the order	• Extracting information regarding profit or losses from financial reports to aid investment decision-making • Information extraction and summarization of legal documents	• Automation of site and due diligence checks

Source: Accenture, July 2019

ASSESSING THE PROS AND CONS OF NEURAL NETWORKS

Energized by the potential value created by this powerful technology, and the pressures to increase automated processes and meet aggressive accuracy targets, we are seeing more and more financial services firms piloting—or already deploying—neural networks. However, while the results are often promising, there are a number of pitfalls that firms should be aware of before pressing ahead with implementing neural networks in live production.

Neural networks' focus on accuracy rather than explainability...

Foremost among the possible challenges is the need for the decisions made by neural networks to be explainable. In highly regulated sectors such as retail banking and insurance, the industry's regulators—and indeed its customers—often ask to understand why a particular decision was made. And as AI decisions come to affect people's lives ever more profoundly, for example through the refusal of a loan, credit card or insurance cover, the demand for explainability and transparency around reasoning is expected to only increase.

However, despite billions of dollars of investment in the field of AI explainability,[5] AI models today mostly remain "black-box" solutions with no native way of articulating the reasoning behind their decisions. Accenture Labs has published a paper outlining the key requirements for making AI explainable[6]—and in it, the authors stress that the full promise of AI systems won't be realized unless people can understand the recommendations they make (see panel).

Some key considerations with neural networks

EXPLAINABILITY

This term refers to the ability of the algorithm to justify its decisions. While explainability is relatively transparent with simpler linear algorithms, the more complex and layered nature of neural networks make them a more opaque "black box" form of AI.

ACCURACY

This defines how correct the model should be—typically driven by business requirements and model capabilities.

DATA INSTANCES FOR TRAINING

Training different algorithms requires different amounts of data. Linear models can be trained reasonably well on a relatively small number of observations, and neural networks' complexity means they need far more instances to learn from.

DEVELOPMENT TIME

The need for a large number of parameters when training neural networks means their development time is typically longer than with simpler AI models.

Why has explainability not made greater progress in neural networks? The fact that they were originally created by the academic community means the core goal guiding their development was not explainability or algorithmic transparency, but the highest possible accuracy. So neural networks were not originally designed to answer explainability questions. And while they have been shown to solve remarkably difficult tasks, concerns remain over their high complexity and limited transparency – reflecting low awareness and understanding of how the neurons interact within multi-layered neural networks to arrive at a certain prediction.

...creates a number of additional concerns

A further barrier to explainability with neural networks is their complexity. Simpler models can reveal relatively easily why—for example—a customer has failed a KYC check, allowing a human investigator to validate the accuracy of the judgment and make the final call. The complexity of neural networks' reasoning makes this more challenging.

Beyond the issues around explainability, neural networks' underlying focus on accuracy rather than transparency can give rise to a number of other concerns. One of the main worries is that it may be difficult to spot bias that could manifest itself through discriminatory outcomes in the long term. This might lead to unfair treatment of certain groups of customers, and potentially legal action for alleged discrimination.

> **"Going forward, AI promises to help us identify dangerous industrial sites, warn us of impending machine failures, recommend medical treatments, and take countless other decisions. But the promise of these systems won't be realized unless we can understand, trust and act on the recommendations they make. To make this possible, high-quality explanations will be essential."**

Source: Understanding Machines: Explainable AI. Accenture 2018.

Firms should consider the available data and infrastructure...

There are two other areas that firms considering implementing neural networks should examine carefully. One is that the data used for training these models is of sufficient quality, scale and diversity. While a simple algorithm might be adequately trained using less than a hundred data instances, with a neural network the number is likely to run into tens of thousands or even millions. Otherwise, once in live production, the system may either make some decisions with low confidence or return completely unpredictable results.

The other area to consider is IT infrastructure. Neural networks require vast amounts of processing power, and often the only way of accessing sufficient capacity is by tapping into the cloud.

...and how to build trust among stakeholders

Finally, as with any emerging technology, trust and confidence among stakeholdersis key to neural networks' long-term use. But the concerns over explainability, algorithmic transparency and bias can all put this trust under threat. While any gaps in computing power or data can be noticeable almost immediately, explainability issues may become apparent only at the later stages, often after the models have been fully built. This can delay go-live dates and create mistrust and skepticism around AI among end users.

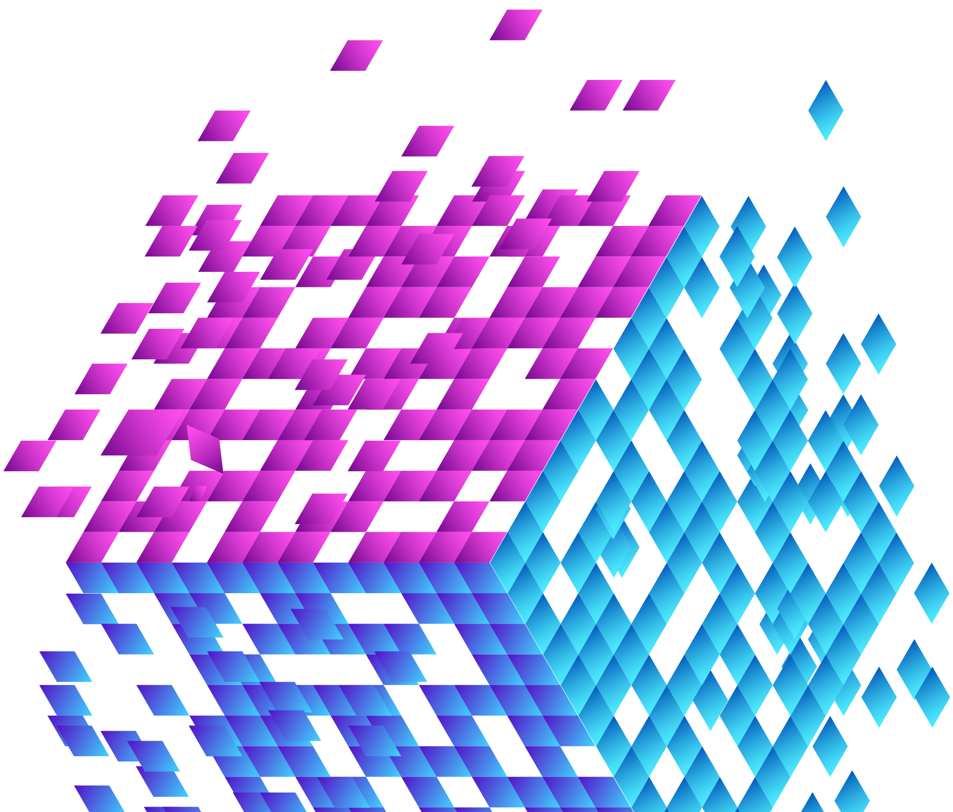

DIVERSE PERFORMANCE OBJECTIVES AND ALGORITHMS

While enhancing the accuracy of neural networks is an important target to aim for, we would recommend that firms also focus on other performance objectives and consider alternative algorithms as part of their AI suite.

Some examples of these algorithms are shown in the table below, complete with their respective requirements in terms of data, explainability, accuracy and processing power.

Table 2: A comparison of different AI algorithms and their requirements[7]

	Data Requirements	Explainability	Accuracy	Processing Power
Linear regression	◔	●	○	◔
Logistic regression	◔	●	○	◔
Naïve Bayes	◑	◕	◔	◑
Support vector machines	◑	◕	◔	◑
Decision trees	◑	◑	◑	◕
Random forest	◑	◑	◕	●
k-means	◑	◑	◑	◕
Neural networks	●	○	●	●
Ensemble learning	◕	○	●	●

Source: Accenture, July 2019

FROM “BLACK BOX” TO “GLASS BOX”

Start with the process to be automated—not the technology...

For any firm considering implementing AI to automate a process, the first step should be to focus not on the technology, but the specific process to be automated. By evaluating, identifying and quantifying the level of need to explain the model’s decisions in the process—and the related risks in areas such as bias—the business can define explainability upfront as an overarching business requirement, and balance it with the goal of enhancing the accuracy of the model. This opens the way for developers to build in explainability from the beginning, allowing the neural network model to become a “glass box” rather than a “black box”.

To support this approach, firms should create ethical guiderails that broadly define the use cases for which neural networks—and AI more generally— should and should not be used. As a rule of thumb, neural networks are more suitable for the processes where accuracy is more important than explainability, and the impact of incorrect decisions on people’s lives and firms’ reputation and financial stability is relatively low. However, as explainability improves in years to come, usage of neural network may expand into more areas.

...and apply ethics as a matter of policy

It’s significant that world-leading AI pioneers such as Google LLC[8] and Microsoft Corporation[9] already have ethical frameworks in place regulating their AI use cases. In Accenture’s view, financial services firms should learn from their example by imposing ethical criteria and requirements that reflect their own values and address wider legislation for algorithmic transparency at a policy level. And at an operational level, these criteria should be made a part of every opportunity qualification framework—along with technical readiness and business benefits—during the process assessment and algorithm selection processes.

Figure 1: A process assessment

Where are you at?

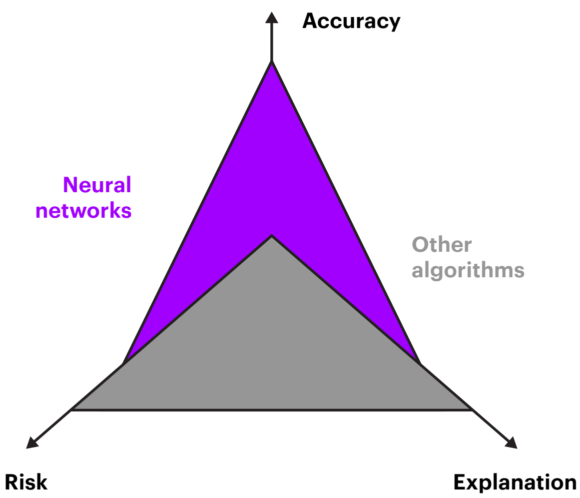

Source: Accenture, July 2019

Trading-off explainability and accuracy

That said, explainability and accuracy can sometimes be traded off against each other (see Figure 1)—and there are certain processes and scenarios where explainability becomes "good to have" rather than "must have". These are scenarios where the problem is well understood, and the recommendation can be approximated and verified with alternative techniques, or alternatively where there is a human agent in the loop, who is competent and authorized to validate and override the decision made by AI. These are the cases where the risk is relatively low, and accuracy more important than reasoning—with bank product suggestions being one example.

Finally, there are some processes where the system's decision might be misinterpreted or misused by end users. For example, staff in bank branches might notice that it is easier for customers to have their mortgage application approved if they do not currently have any outstanding loans. As a result, they might advise a customer who wishes to get both a mortgage and a loan to apply for the mortgage first—as the mortgage application is the more stringent process, while the loan is more straightforward and easier to get approved. If staff did behave this way, it could result in a higher risk of defaults on mortgage payment in the future.

As this scenario underlines, it is important to think carefully when defining the level of detail in the explanation of an AI decision, and the list of users to whom the explanation should be displayed. Figure 2 illustrates the spectrum of use cases for which neural networks generally should not be used, and those for which they should only be used in combination with strict controls and governance.

Figure 2: Use cases for which neural networks can be used with due consideration

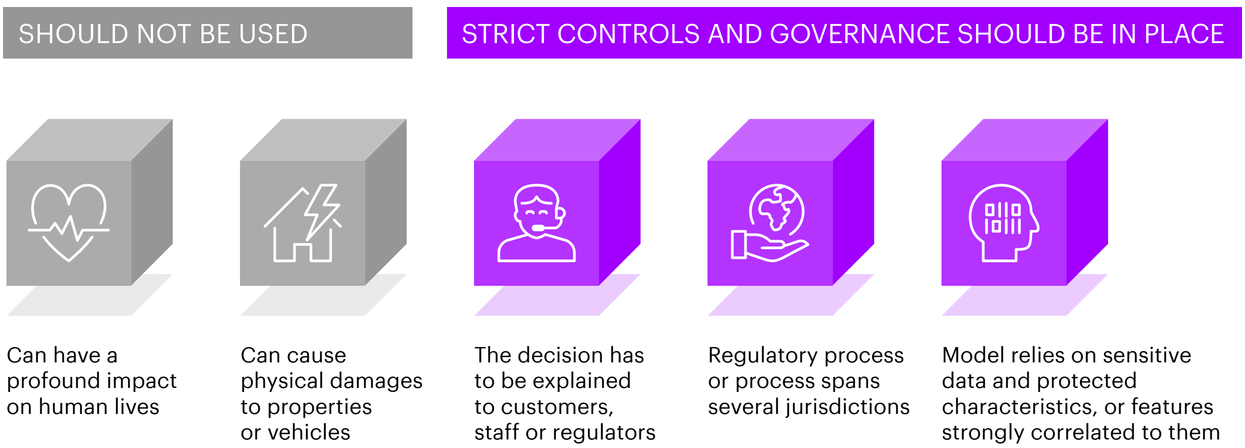

Source: Accenture, July 2019

SIX STEPS TO THE EFFECTIVE IMPLEMENTATION OF NEURAL NETWORKS

As we've highlighted in this paper, for firms to select, develop and implement neural network algorithms that are not only effective and accurate but also compliant, we believe they should first put in place the appropriate IT infrastructure with sufficient processing power—generally from the cloud—and then overlay this with a robust set of frameworks, controls, tools and ethical principles.

Building on these solid foundations, here are six steps for an effective implementation.

1. Have your training data ready

Adequate training of neural networks requires far more data instances than simpler, linear AI models. To avoid problems with insufficient training, potential issues over data quality and volume should be addressed at various levels in the organization. At an enterprise level, the Chief Data Officer (CDO) should provide a data lake that is curated and appropriately governed, striking a balance between providing greater insight and value while providing compliance with data protection regulations.[10] Also, all source data for training should adhere to the three data-focused tenets: provenance (the history of the data throughout its lifecycle); context (the circumstances around data use); and integrity (data security and maintenance).[11]

During exploratory data analysis, the team of data scientists should consider not only the volume of available data, but also how diverse the sample is, whether all groups of customers are equally represented, and whether the human agents in any training scenarios made decisions that were fair to all customers. Focusing on these issues can allow the data scientists to apply appropriate data treatment strategies to reduce any embedded bias in the data, while also providing enhanced coverage (the model's ability to make decisions) and accuracy (the model's ability to make correct decisions).

2. Change your culture

Bias can creep into a neural network not only through the training data, but also at the fine-tuning stage. For example, the model developer might discover that manually increasing the weighting attributed to certain factors can lead to more accurate results—but this might have an unintended negative impact on fairness. So it is important to change the culture of the data science teams by providing them with ethical and anti-discriminatory training, thereby increasing their awareness of protective characteristics, bias and ethicality. Leaders should also look to create diverse teams of data scientists and promote collaboration and peer reviews.

3. Keep experimenting

When building a model, the data scientists should compare several algorithms and keep experimenting with them to evaluate their relative levels of suitability for the task at hand. In doing this, they should bear in mind that a neural network may not be the best choice for use cases where training data is scarce and where more transparent algorithms could play the same role.

4. Test and validate

In addition to standard machine learning testing on 20% of the "holdout sample" or validation set, it is also advisable to conduct cross-validation, since this can reveal a model's ability to generalize—which means making decisions on previously unseen cases. To allow this, the data sample used should be carefully selected to represent both frequently occurring scenarios and also boundary or "outlier" conditions. In applications such as fraud detection, where the amount of "live" fraudulent data is usually much smaller than non-fraudulent transactions, anomalous transactions should be deliberately selected or created for testing purposes. Finally, fairness testing should be carried out to complement the technical tests, and make sure that the model does not discriminate against or favor certain groups of customers.

5. Keep on top of documentation

Each step of the model design process should be recorded, and the documentation about each model maintained and updated continuously throughout its entire lifecycle. The documentation should also be strictly governed while being made accessible to all relevant and authorized stakeholders.

6. Create a center of excellence for responsible AI

To attain and sustain leadership in the use of neural networks, each financial services firm should set up a dedicated function to act as a center of excellence in the responsible use of AI, including neural networks. This central resource would establish preferred practices and standards, drive AI-focused thought leadership and research, and impose risk controls and governance so all use of AI across the organization is ethical and compliant.

> **Brakes help a car go faster.**[12]
>
> Dr. Rumman Chowdhury, Accenture Labs

WHAT NEXT?

Effectively implementing a neural network is not the end of the story. As any AI model operates, learns and makes decisions, the world around it continues to change, and its performance may degrade over time.

All of this means that controls and quality assurance should not stop after deployment but should be maintained and reviewed throughout the model's entire lifecycle. Ongoing considerations include the following.

Regulatory changes: Continuous monitoring of the regulatory environment is critical, both with respect to automated processes (such as changes in trade or AML laws) and AI legal requirements (such as explainability obligations or data handling restrictions).

Accuracy and bias monitoring: Ongoing monitoring of accuracy and bias has to be maintained, to allow the neural network's performance not to drop below defined thresholds and that its decisions do not systematically favor certain groups.

Security and re-training data screening: Maintaining strict cyber security provisioning is vitally important. Also, as firms re-train their models to keep pace with changes in the wider environment, they should take account of the risk that "poisoning" training data with adversarial examples can skew neural networks' reasoning, resulting in the models taking decisions that are incorrect and potentially illegal. And malicious access to the model itself—whether from inside or outside the firm—can compromise competitive advantage, and even lead to criminal misuse of the model, such as playing on stock market predictions.

Quality assurance and audit: Many financial firms have already established a framework for traditional model validation, and in many cases have been using this for years. However, while this framework might be adequate for relatively simple models, firms should consider introducing adjustments and refinements to accommodate the higher complexity and ambiguity of neural networks and deep learning models.

Plan B: There should always be a back-up model that the system can turn to automatically, in the event that the front-line model displays signs of degradation.

CONCLUSION: A POWERFUL TOOL— TO BE USED WITH CARE

There's no doubt that neural networks are a powerful tool for financial services firms as they seek to improve efficiency, enhance the accuracy of their decision-making, boost revenues and improve the experience for customers.

But like any other powerful tool, neural networks should be used with due care and consideration—and in the light of the fullest possible information.

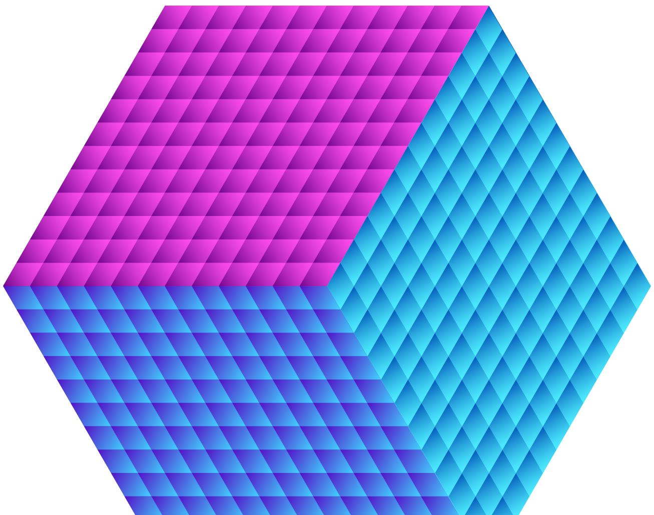

Authors

Sabyasachi (Saby) Roy
Managing Director
Financial Services Technology
Advisory Artificial Intelligence
Lead UKI
sabyasachi.a.roy@accenture.com

Elena Khmeleva, PhD
Technology Consultant
Financial Services Technology
Advisory Intelligent Automation UKI
elena.khmeleva@accenture.com

Additional contacts

Elodie B. de Fontenay
Director
Financial Services Technology Advisory
Offering Development Global
elodie.b.de.fontenay@accenture.com

David C. Jones
Managing Director
Financial Services Technology Advisory
Intelligent Automation Global
david.c.jones@accenture.com

Tim Broome
Managing Director
Financial Services Technology Advisory
Intelligent Automation APAC
tim.broome@accenture.com

About Accenture

Accenture is a leading global professional services company, providing a broad range of services and solutions in strategy, consulting, digital, technology and operations. Combining unmatched experience and specialized skills across more than 40 industries and all business functions—underpinned by the world's largest delivery network—Accenture works at the intersection of business and technology to help clients improve their performance and create sustainable value for their stakeholders. With more than 482,000 people serving clients in more than 120 countries, Accenture drives innovation to improve the way the world works and lives. Its home page is **www.accenture.com**

Stay Connected

Accenture Financial Services Technology Advisory

www.accenture.com/us-en/services/financial-services/technology-advisory-index

Connect With Us
www.linkedin.com/showcase/16197660

Follow Us
twitter.com/TechAdvisoryFS

Disclaimer: This document is intended for general informational purposes only and does not take into account the reader's specific circumstances, and may not reflect the most current developments. Accenture disclaims, to the fullest extent permitted by applicable law, any and all liability for the accuracy and completeness of the information in this document and for any acts or omissions made based on such information. Accenture does not provide legal, regulatory, audit, or tax advice. Readers are responsible for obtaining such advice from their own legal counsel or other licensed professionals.

References

1 "Global financial services industry bullish towards disruptive technology but widescale adoption yet to come," Intertrust, November 28, 2018.

2 "#Augmented Finance & Machine Intelligence," Autonomous Next, April 2018.

3 "How AI Boosts Industry Profits and Innovation," Accenture, 2017.

4 "HSBC rolls out facial recognition for mobile," Financial Times, May 8, 2018.

5 "DARPA Plans To Spend $2 Billion Developing New AI Technologies," Forbes, September 7, 2018.

6 "Understanding Machines: Explainable AI," Accenture 2018.

7 "Machine learning algorithm cheat sheet for Azure Machine Learning Studio," Microsoft, March 3, 2019.

8 "Responsible AI Practices," Google, 2019.

9 "Our approach: Microsoft AI principles," Microsoft, 2019.

10 "What's Your Data Strategy?," Harvard Business Review, June 2017.

11 "Redefine Your Company Based on the Company You Keep," Accenture, 2018.

12 Accenture Labs, March 11, 2019.

191109