Deutsche Bank
Corporate Bank

Unleashing the potential of AI in securities services

#PositiveImpact

flow
SPECIAL

Unleashing the potential of AI in securities services

Artificial Intelligence (AI) is fast becoming an essential tool for the financial services industry – not least within securities services. Here, it promises to herald a new era of competitive services based on digital intelligence, with AI already being deployed to improve a wide range of processes from front-office activities to compliance and operations. Tapping the transformative potential of AI, however, requires careful thought and preparation. Ensuring fit and proper governance for AI systems, grounded in a sound understanding of AI technologies, regulations, how ethical principles should be applied, and human domain expertise, remains a key priority. It's one that will enable industry participants, from broker dealers and global custodians to asset managers, to better harness the full value of AI.

Co-authors

Boon-Hiong Chan, Global Head, Securities Market Development & Technology Advocacy, Deutsche Bank Securities Services

Marcus Storm, Head of AI Products, Product Management, Deutsche Bank Securities Services

Deutsche Bank contributors

Peter Hung, Strategy and Innovation Network, Deutsche Bank

Ms Jiali Liu, Government and Regulatory Affairs, Deutsche Bank

Published June 2021

Contents

Foreword

Whether in furthering our understanding of diseases, delivering improvements to manufacturing methods or creating a fairer, more robust financial market, artificial intelligence (AI) continues to unlock new growth opportunities and competitive advantages for a diverse range of organisations and their stakeholders.

The rapidly expanding applications of AI are moving into the financial industry, where the deployment of AI is expected to unlock new growth opportunities and competitive advantages. AI is being implemented not only to improve front-office activities, such as client interactions and investment analysis, but also in other areas, such as compliance, risk management and operations, where it is helping to bolster productivity and scalability.

This paper focuses on the transformative potential of AI for one specific area of financial services: the post-trade securities services and custody industry. Here, AI is being progressively applied across three broad areas: identifying historical trends from existing ones, discovering up-to-the-minute trends using real-time data and analysis, and predicting future trends using a mix of both past and present data.

While the possible applications for AI technology are yet to be fully explored, a number of compelling use cases have already emerged. For example, AI is being applied to divide clients into discrete groups to reflect their similarities – a process known as client segmentation – that helps custodians better understand their clients and develop products and services tailored to the shared needs of specific segments. Elsewhere, the technology is also being used to detect and prevent settlement failures – an increasingly high priority for custodians, not only to avoid financial penalties, but also to ensure smooth and seamless service for their clients.

Yet, given the significant sway AI has on the decision-making process, the technology comes with a number of stringent governance requirements. The primary strength of AI – that it can rapidly process great swathes of complex data – is also the source of one of its greatest shortcomings: the complexity it manages also hides myriad opportunities for the model to go wrong.

This paper provides an introduction to AI in the securities post-trade industry – complete with use case examples, data and key algorithms that can be utilised in AI-based applications, as well as insight into governance considerations such as explainability and transparency. It concludes with key recommendations for industry participants to help facilitate the responsible development and adoption of AI.

Boon-Hiong Chan,
Global Head, Securities Market Development & Technology Advocacy,
Deutsche Bank Securities Services

Marcus Storm,
Head of AI Products,
Product Management,
Deutsche Bank Securities Services

1
AI in securities services

Artificial intelligence (AI) and machine learning (ML) represent a growing new wave of intelligence-based services in securities post-trade and custody. With multi-step work processes, numerous participants involved in each transaction and high volumes of transaction data, securities post-trade and custody is the kind of complex environment in which AI and ML can excel – helping service providers and clients to improve their product offerings and better manage market and operational risks.

1.1 Risk management

These risks include regulatory, operational, liquidity and counter-party considerations that are present in the typical two business days between the time a trade is executed and the time it is settled (T+2). While these risks can (and continue to) be addressed with other technologies, the challenges in managing them in a timely and responsive manner are amplified by a myriad of factors, including time-zone differences, incomplete communication across the entire value chain, the multiple participants involved and severe time pressures to settle a trade. Tackling these risks efficiently and effectively – either via AI or other means – presents banks and their clients significant competitive advantage.

At the heart of AI's power to achieve this are three drivers:

- Data: The data chosen must be suitable for the use case, as the data selected will influence the output. Securities market transaction data is generally depersonalised, so does not in itself contain protected attributes – although there can be other type of unfair bias at an institutional level. This means that the final models must still be rigorously tested for inferred bias, i.e. situations where discriminatory conclusions are drawn from statistical calculations;
- Algorithms: Algorithms and their parameters are the engine – a set of processes – that analyses the data, determines relationships, possibilities and results. Algorithms and their parameters play fundamental roles in determining explainability and transparency (see Section 3: Learning types, algorithms and governance);
- Processing power: Training a powerful machine-learning algorithm often means running huge banks of computers for days, weeks or even months. The algorithms of today require significant processing power as they run through terabytes of data in the task of creating the perfect model.[1]

1.2 Applying AI

AI can be applied to securities services to tackle a host of concerns, including those surrounding risk management, as well as to provide banks with a more rounded understanding of their clients' behaviours. There are three broad areas where these capabilities can be progressively applied, which can be split across three broad dimensions: past, present and future. For example, AI could be applied to existing data to identify historical trends, to current – or even real-time – data to identify current trends, or to a mix of both past and present data to predict future trends.

1. Past: Execution and insights mining:

— Automatically identifying possible breaches of service use.
— Analysing and improving client service through retrospective analysis of client. queries using natural language models and sentiment analysis.
— Using recommender systems to suggest 'next best' or alternative approaches (e.g. in time-stressed periods).
— Improving data quality by using new techniques to categorise, monitor and update data.

2. Present: Monitoring and diagnostic processes:

— More comprehensive monitoring of counterparty risks through more complete risk-related data.
— Client service improvements (auto-resolving common issues, 24/7 availability of help bots) using natural-language models.
— Quick, automated answers to queries and notification of reasons for non-straight-through processing reasons that can signal deeper operational issues.

3. Future: Proactive and predictive activities:

— Analysis of an existing client's transaction patterns to predict anomalies.
— Predicting outcomes to guide proactive remedial actions for clients.
— Optimising use of expensive assets, e.g. liquidity, securities borrowing.
— Predicting the likelihood of settlement failure for timely, proactive remediation and avoidance of penalties.
— Predicting volumes and potential issues in advance of portfolio flows and enabling proactive rebalancing to enable better preparations.

The following section will explore several of these use cases in greater detail.

"Over the past few years we have begun to deploy AI solutions across our Securities Services franchise. We have already introduced sophisticated client segmentation processes alongside our S-2 Predict tool for preventing settlement failures, while elsewhere self-executing bots are parsing natural-language messages as part of our client-facing chatbot, Debbie, which responds in real time to customer requests such as settlement status queries. Our team in China also operates with an AI-worker called "Yi" that can execute complex operations workflows flawlessly across different IT systems to reduce operational risk. The opportunities are endless, and we are not even yet scratching the surface. Going forward, we will continue pushing this emerging technology to the forefront – stay tuned!"

Paul Maley,
Global Head of Securities Services, Deutsche Bank

2

AI use cases

The scope of possible applications for AI technology in the securities services space is vast and not yet fully explored, but a number of compelling use cases have already emerged. The following sections focus on two of these: client segmentation and avoiding settlement failures.

2.1 Client segmentation

AI can be applied to client segmentation, i.e. the practice of dividing clients into discrete groups to reflect their similarities. This is an important process – helping custodians, for instance, to develop products and services that meet the shared needs of specific client segments. Here, AI can seek out patterns within client data quickly and on a huge scale – creating outputs that would otherwise be unachievable through manual means.

Applying AI to the client segmentation process has a broad range of applications, including:
- Improving internal business risk controls by automatically detecting counterparties that suddenly shift their behaviour from what is expected.
- Helping sales and relationship managers to proactively identify individual challenges and solutions for clients in the same segment based on the asks from their peers, or for clients with similar behaviours.
- Improving operational efficiency through the proactive use of historical data to help solve challenges particular to each client cluster.
- Finding new ways to cluster clients to enable groups of clients with other needs to be identified.
- Improving data quality through the identification of abnormal data points. Since market data quality can be variable, securities services can use AI to raise data quality by adding extra filters or layers of checks. Anomalies can then be identified and either addressed if significant, or removed from overall outputs if not.

2.1.1 Configuration is key

As with any AI solution, the value of the exercise is dependent on ensuring the algorithm is correctly tuned and focuses on the right parameters. Unguided, the AI can output data that incorporates hundreds of different parameters – placing clients into distinct clusters that are not necessarily very meaningful. Defining both the number and purpose of each cluster is where business knowledge and cooperation is crucial. For example, it is possible to run a client segmentation which, out of all the data attributes of a client, assigns a higher importance to its settlement activity, and separates out high-trading and low-trading intensity clients regardless of their segment – giving more precise, useful insights.

Ensuring that the solution generates truly insightful outputs is also contingent on the training and live data being high-quality and reliable. In addition, business stakeholders should understand the concepts used in order to provide use cases and clear objectives to guide the development process.

While the set-up of client segmentation must be carefully managed to ensure accuracy and appropriate weighting, once properly configured, the results can yield valuable insights. In the case study that follows, we walk through some of Deutsche Bank's own applications of client segmentation using AI.

Client segmentation using AI

Deutsche Bank leverages AI-driven client segmentation to improve and tailor its offering for clients. Figures 1 and 2 represent the visual outputs of fictional client segmentation runs. The points, which represent client entities, appear closer or further from each other depending on how similar their behaviour is across multiple parameters.

Figure 1: A 3D representation of the behaviours of Deutsche Bank's client clusters

This is a three-dimensional (3D) representation of the behaviours of client clusters, compressed from several hundreds of feature dimensions for illustrative and depiction purposes. To ensure data protection, all production data has been securely randomly scrambled, changed, and anonymised.

Each of the three dimensions is an approximation of 'client behaviour' as measured in different ways. Each point represents a client, and the closer a client appears geometrically to another, the closer their behaviour is across the various parameters.

For example, in this hypothetical illustration, we can see that certain broker dealer, insurance and global custodian clients all operate in a similar way. Deutsche Bank would then be able to take this information to drive proactive client servicing and scalable product developments for clients across different traditional segments.

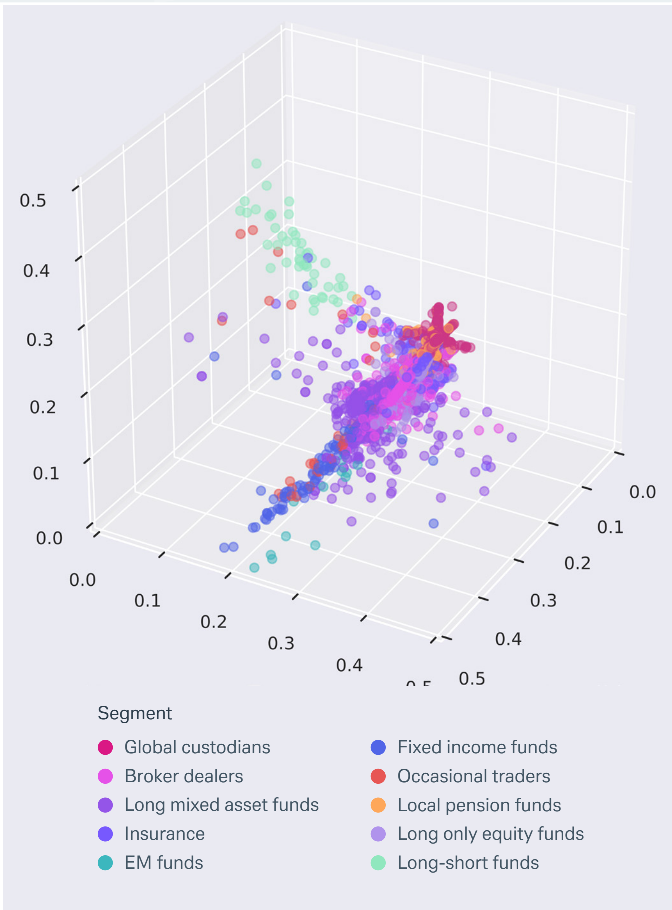

Source: Deutsche Bank Securities Services

Figure 2: A 2D representation of the behaviours of Deutsche Bank's client clusters

This is a two-dimensional representation of the fictional behaviours of client clusters, compressed from several hundreds of feature dimensions for illustrative and depiction purposes. To ensure data protection, production data has been securely randomly scrambled, changed, and anonymised.

As with the previous figure, each dimension is an approximation of 'client behaviour' as measured by a variety of different factors. Each point represents a client and the closer a client appears geometrically to another, the closer their behaviour is across the various parameters.

In this figure, five separate major client clusters were identified from the dataset, and different client types with similar behavioural patterns were also identified – marked out by five different colours on the chart – to facilitate proactive client services and operational responses. For example, those red dots that overlap with pink dots or with the green dots represent sub-segments of a client cluster (represented by the red dots) that have similar behaviours like sub-segments from other client clusters (represented by the pink or green dots). This allows client services and products to be anticipated and delivered based on client's needs, and not on fixed client segment classification.

Source: Deutsche Bank Securities Services

2.2 Settlement failure detection

Settlement failure is costly for both banks and their clients. In addition to the substantial financial penalties levied by the Central Securities Depository Regulation (CSDR),[2] failures disrupt trading strategies, require manual interventions and act as a barrier to a smooth and seamless custody service. In addition, multiple failures can even begin to shake market resilience and reputation, and can precipitate other, much larger fall-outs.

Among the factors that can lead to settlement failure are:
— Insufficient balance in the account;
— Insufficient balance at the counterparty;
— Mismatched cycles;
— Issues at the central securities depositories;
— Day and time of transaction;
— Corporate actions or other market events; and
— Macroeconomic factors, such as volatility, affecting the trading environment.

Preventing and reducing failure rates is a top priority for custodians – one that, provided there is a clear goal and sufficient labelled data, can be solved using an AI-based approach. For example, custodians are leveraging AI technology to predict failure possibilities based on specific features and historical settlement data, including time, country, exchange, amount and asset type, as well as any combination thereof. In addition, they are also generating pre-trade predictions at the point a trade order is submitted – providing operations staff with a real-time view of the issues that can delay settlement.[3]

2.2.1 Preventing settlement failure: S2-Predict

It is often difficult to determine the root cause of settlement failures. In response, in 2019 Deutsche Bank Securities Services began work on a solution known as S2-Predict, designed to tackle settlement failure in three different ways:

Past: Historical analysis

S2-Predict tracks and shows historical failure-rate patterns and trends globally for a specific client, providing a benchmark against which current outcomes can be judged. This gives the client direct insight into whether the number of settlement failures they experience is going up or down compared to its historical average – allowing the team to proactively react to either maintain or improve their processes.

The solution also uses ML to provide deeper insights, such as a client's performance against similar firms. This means the client is also able to see how its settlement failures rates compare to those in the wider market. This can help protect the client from complacency: while the rate may be down compared to their own benchmark or client segment classification, it could still be much higher than those firms outside of its segment but with similar market behaviours, suggesting there are efficiencies being left unexplored.

S2-Predict is based on an algorithm that uses supervised learning (see Section 3.1.1). The algorithm is trained using historical data and calibrated with real-time information to detect anomalies, helping ensure the analysis is as accurate as possible, and can even support client decisions via human-operated expert systems.

Figure 3: A sample of the historical settlement rate dashboard

In this figure, production data has been securely randomly scrambled, changed, and anonymised for illustrative purposes. An advanced algorithm is used to detect outliers in the upper left-hand graph, with outliers shown as grey, to give clients insights into historical settlement rates.

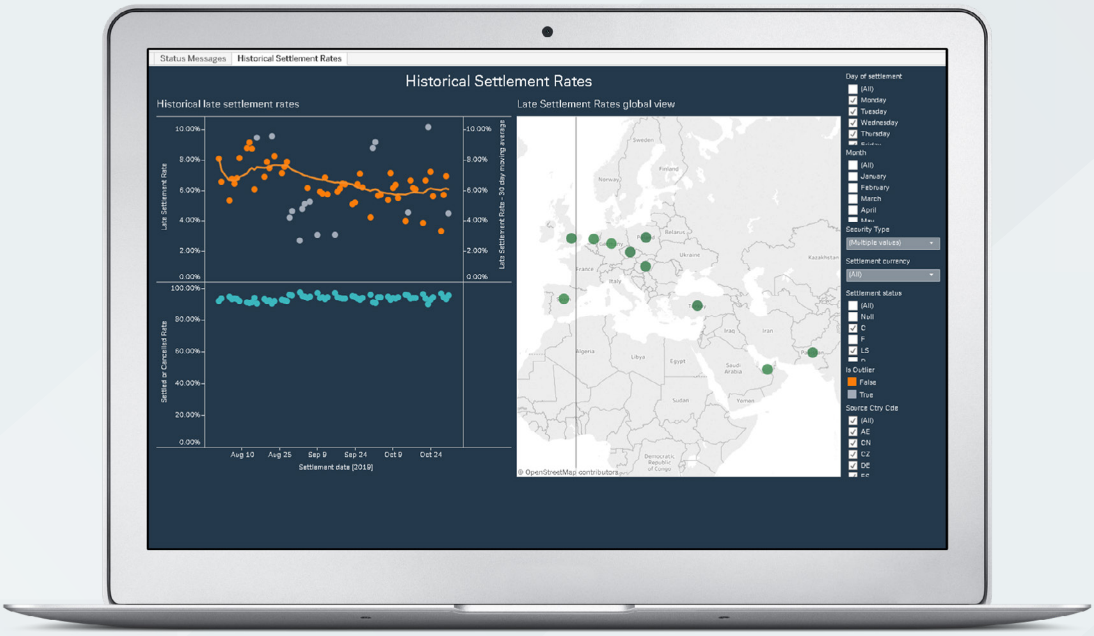

Deutsche Bank Securities Services Product Management

Present: In-flight transaction analysis

In addition to reviewing historical data, S2-Predict's AI system also monitors the status of in-flight transactions, including the probability of failure in each case, to give unique advanced warning to clients and for human-initiated actions. In order to provide advance warning on in-flight transactions, the solution's prediction models are also integrated with real-time data. The system constantly receives new inputs and uses them to generate up-to-the second outputs, generating forward-looking views on incoming transactions within five seconds of the data flowing in.

Future: Transaction analysis

For future or hypothetical transactions, S2-Predict incorporates an AI system that can estimate the chance of failure, with a breakdown of the factors contributing to that risk. Clients can run any transaction settlement as a hypothetical use case. If the settlement is likely to fail under the current scenario, the client can take the necessary steps to mitigate this risk before initiating the transaction – significantly reducing the number of failures overall.

Figure 4: A sample of the settlement failure prediction dashboard

In the this figure, production data has been securely randomly scrambled, changed, and anonymised for illustrative purposes. The number at the top of the sidebar on the right shows the probability – in this case 0.567 – that the settlement of the selected transaction will fail. The chart beneath it displays the different features contributing to the probability of failure, alongside the degree to which they contribute.

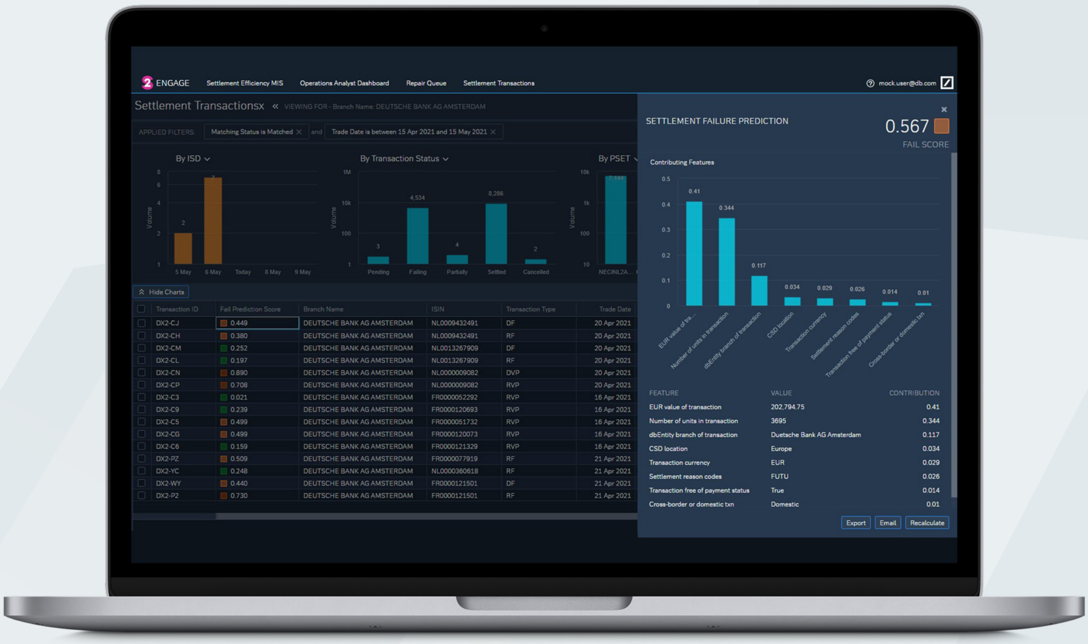

Deutsche Bank Securities Services Product Management

3

Learning types, algorithms and governance

In section 2, we introduced a range of AI-driven operational risk management services. These are underpinned by different learning types and algorithms, which play a key role in enabling the analysis, visualisation and interpretation that make AI services so valuable.

As these learning types and algorithms are not created in a vacuum, custodians looking to benefit from them must first consider how to configure their AI models to yield the right outputs.

This means human subject matter experts are involved in every key step – from determining which learning type and algorithm should be used, through to what parameter adjustments should be made, how to train the model and what data characteristics to focus on. This all serves to minimise a common trap of AI services – the so-called "AI black box", an AI system that delivers complex predictions, recommendations or other outputs that cannot be understood by the human operators (see Section 3.3: Governance). Understanding the different model types is, therefore, critical to generating useful and reliable outputs.

Figure 5: An overview of AI learning types

For the purposes of this paper, this figure serves as a broad-brush* taxonomy of the AI learning types and algorithms relevant to the securities services industry.

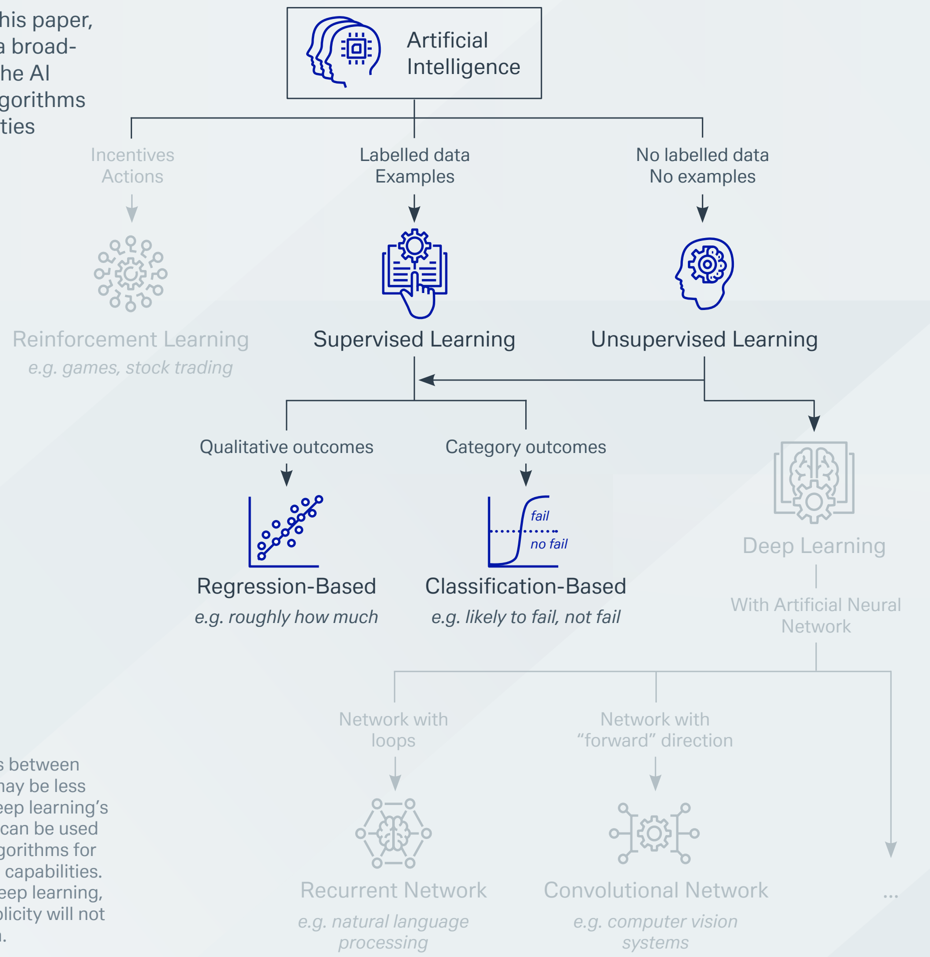

*In reality, the boundaries between different learning types may be less clear cut. For example, deep learning's artificial neural networks can be used in classification-based algorithms for semi-supervised learning capabilities. Networks are a form of deep learning, but in the interest of simplicity will not be covered in this section.

3.1 Learning types

Machine learning (ML) is an application of AI that enables systems to learn and advance based on experience without being specifically programmed for the new tasks, data or other variables. ML focuses on the development of computer programmes that can access data and use it for their own learning.

There are four types of ML: supervised, unsupervised, semi-supervised and reinforced. For simplicity, this section will only discuss the first three types.

3.1.1 Supervised learning

Supervised learning involves an AI model learning from examples and training data sets to produce desired outcomes. There is high transparency in both the data and the algorithm, as the large data sets need to be labelled and human subject-matter experts are required to train the ML algorithms. For data sets with a large number of attributes or "features", a manual procedure called feature engineering is employed to reduce or target specific data dimensions and to tailor the dataset to business-specific problems, which in turn improves the classification accuracy on unseen data.

3.1.2 Unsupervised learning

Unsupervised learning is an ML technique where the model looks for structure and relationships in unlabelled data sets. It is "unsupervised" because no external guidance or defined outputs are provided to the algorithm. However, this does not mean that there is no human participation in these learning algorithms. Experts are involved to adjust hyperplanes and parameters, determine materiality (see Section 3.3.2), and ultimately to decide if the outcomes can be explained and if they are relevant to the business goal.

Readers may be familiar with the term "deep learning" – a specialised form of unsupervised learning that imitates the workings of the human brain in processing data and creating patterns for use in decision making. It is also known as Deep Neural Learning or Deep Neural Network. However, while unsupervised algorithms are being applied to securities services, the applications of deep learning in the field have yet to be explored.

3.1.3 Semi-supervised learning

Unsupervised algorithms can be used alongside supervised algorithms to create semi-supervised learning models. Semi-supervised learning involves combining a small amount of labelled data with a large amount of unlabelled data during training. It provides the benefits of both supervised and unsupervised learning while avoiding the need to label a large amount of data.

In areas such as natural language processing, an interactive approach between the algorithm and the human subject matter experts is becoming increasingly popular. The process is as follows:

1. The trained model will produce a number of classifications for the unlabelled data.

2. The expert will "teach" the model by reviewing new classifications and providing the appropriate labels for those that are incorrectly classified.

3. The ML algorithm will retrain the model, sometimes in real time, to incorporate the newly labelled data.

Once complete, the cycle then restarts and continues. As the data sampling size increases, the classification accuracy can be expected to improve over time.

As many involved in this semi-supervised learning process are not ML specialists, these systems are often developed with intuitive user interfaces. This allows the human operator to easily update classifications and correct the algorithm. The ease and accuracy by which they are able to perform this function is central in ensuring the human operator does not unintentionally compromise the exercise by introducing errors.

3.2 Algorithm types: Recommenders, Predictors and Classifiers

Within these learning models are algorithms – a set of rules, calculations and steps that can be followed to solve a given problem. Some of the most common types of algorithms are:
- Regression-based algorithm: an algorithm that provides predicted output values based on the input data features.
- Classification-based algorithm: an algorithm that classifies or categorises new data into distinct groups.

These algorithms can be mixed and matched to perform different functions, such as providing recommendations, making predictions, classifying objects, and so forth. Choosing informative, discriminating and independent features is a crucial step for effective algorithms in pattern recognition, classification and regression.

> "Understanding the different types of algorithms is a critical first step for those looking to implement an AI solution. Everyone involved in the decision – be it business management, end users or developers – need to understand not only the different use cases of regression-based and classification-based algorithms, but also how they can be used in combination to meet a wide array of specific goals"
>
> Anand Rengarajan, Head of Securities Services, Asia

3.2.1 Predictors

Regression algorithms are selected depending on the number of variables, as well as the type of relationship between the variables.

Predictors – also known as "regression predictors" – attempt to identify relationships between the input variables in order to output a model that can then "predict". For example, a regression predictor could predict the value of a house, given historical transactions based on inputs such as size, number of rooms, and location. In the post-trade industry, one application of a predictor would be to "predict" estimated settlement penalties or fees across different markets, based on a given set of market conditions on which the model has been trained (see Section 2.2 for more details on Deutsche Bank's progress in this area).

Predictors can also be used to recommend actions. Known as recommenders, this subset of predictors can be used to suggest 'next best' actions and alternatives based on patterns and relationships between potential choices.

3.2.2 Classifiers

Classifiers place results into categories and can be applied to situations with either two classes (e.g. fail/succeed), or with multiple classes/outcomes (e.g. tagging a transaction as: urgent; likely to fail; no penalty; or regulatory report needed). In securities services, given a set of conditions, classifiers can also be used to "predict" certain possibilities, such as settlement failure, or used to improve internal business risk controls by, for example, automatically detecting changes in the behaviours of counterparties (see Section 2.1 for more details on Deutsche Bank's progress in this area).

The feature that distinguishes a classifier from a regression predictor is its probabilities, which include false negatives and false positives. The importance of each of the "false" outcomes needs to be assessed according to its context and materiality, which requires expert human decision-makers to make judgement calls on parameter adjustments.

3.2.3 Classification vs Clustering

Within classification-based learning, there are two common methods used to recognise patterns:
— Classification uses predefined categories to group data. Some examples of Classification algorithms include the "Random Forest" approach (see Glossary).
— Clustering uses similarities between data as a basis to group or "cluster" them together. Some examples of Clustering algorithms include the Expectation-Maximisation Algorithm and K-Means Clustering (see Glossary).

3.3 Governance

3.3.1 Ensuring model accuracy and consistency over time

AI models – especially those that have been invested in heavily – carry significant clout in the decision-making process. This is a natural step as we move into a world increasingly defined by huge volumes of data that simply cannot be processed manually. Yet it also poses a governance challenge: AI is not perfect. In fact, given the complexity and scope of the undertakings AI is generally tasked with, there are a myriad of opportunities for models to go awry.

It is therefore crucial that AI solutions have robust governance measures in place to ensure that the outputs are accurate and that users are in a position to critically evaluate their conclusions. Ensuring this robust governance starts with understanding the algorithm and the model.

The choice of algorithm should be understood by business management, end users, and developers. Algorithms have different strengths and capabilities resulting from their parameters, which operate on the basis of relationships, structure and the statistical characteristics of data sets. The algorithms are differentiated by aspects such as processing speed, which, in turn, can be influenced by: the need for accuracy, the processing load, the ability to analyse and associate data (e.g. whether the algorithm can only classify two classes or many), accuracy, the number of parameters to adjust, the "F1 Score" (which shows the balance between precision and completeness) and other metrics.

A model consists of learning type, algorithm, parameters within the algorithm, and data/information input. A model is produced when the algorithm has processed the relationships found in the training data. The better the model is, the more accurate its output. If the model can be generalised well in new situations, the chosen technical approach is said to deliver quality and dependable results. If the model does not generalise well, as determined by human subject matter experts, it needs to be retrained or retooled with a new approach (such as new algorithms) before it is placed into production.

Care must be taken to monitor the model's relevance over time to prevent any occurrence of Model or Concept Drift. Model drift is where the relationship between two variables in the input data starts changing. For example, the relationship between two input variables might shift from multiplication to addition without this being labelled to the model. As a result, the accuracy of the model outputs will start to drift. This effect can be amplified by the fact that the model attempts to learn from the new inputs, leading to strange outcomes if not managed correctly. For instance, a business might look to forecast its sales based on historical data, but a major event such as the Covid-19 pandemic will suddenly change the inputs significantly, making the new forecasts unreliable both in the short term (because the historical data won't be representative of the current anomalous situation) and the long term (because the model will be learning from anomalous data that is unlikely to be directly predictive of future trends). Human experts who are in the loop can identify these changes and prevent the potential degradation of results they can cause. This may mean, for instance, that models need to be regularly retrained to ensure they retain their focus.

3.3.2 Materiality, explainability and interpretability

As noted, while AI can be finely tuned to divine insights invisible to the human eye, it can also get things wrong. Materiality refers to the impact of a wrong decision: the greater the impact of an AI mistake, the greater the materiality. For example, if AI is used in medical surgery, the materiality is high, since a mistake could result in permanent harm to the patient or even kill them. The materiality of an AI decision in securities services may not be life or death, but it is still significant, with trillions of dollars at stake across the industry.

Gauging the materiality of an AI decision or output is important because it corresponds to the level of controls that must be in place to ensure the output is correct. The most fundamental control in this respect is the concept of "explainability" – that is, how easy it is for humans to explain how the AI model works and derives its results. Put simply, the higher the materiality of an output, the higher the burden of explainability for the user.

Closely related is the concept of "interpretability", which denotes how easy it is for humans to understand what an AI output means in practice. Both explainability and interpretability should be high when materiality is high.

A comparatively lower number of parameters (variables used to calibrate models) in the machine-learning process contributes to ease of explainability and interpretability of the results. The context in which parameters are applied can also help to keep them comparatively simple, as the context draws a boundary on the range of outcomes.

If AI becomes a "black box", meaning that humans are no longer able to adequately explain or interpret the drivers of the results, this is a cause for ethical and governance concerns. There have been attempts to introduce ways of quantifying explainability on top of existing models. For example, using Shapley values (See Glossary) or deriving locally faithful explanations for predictions,[5] often have specific drawbacks, either in terms of processing speed or applicability to only certain types of algorithm.

In general, how an AI application is categorised from a risk-management perspective instead depends on the following factors: the context in which the application will be used; the sophistication of the algorithm(s); the complexity of the models; and the materiality of the application's outputs for people, communities and economies.[6]

If an application has any "black box" attributes, extra due diligence is needed. This can be achieved through a 'human-over-the-loop' structure, where a subject-matter expert monitors the model from the outside, as opposed to a 'human-in-the-loop' structure, which embeds a subject-matter expert within the AI process. This human-over the-loop approach ensures accountability for both the AI and the process in which the AI is embedded, as well as the overall legal responsibility for the process, which can be assigned to specific individuals or teams.[7]

If viewed from a technology risk management perspective, other factors would apply, including: model and software development lifecycle activities; data privacy and governance; cybersecurity regulations; and industry best practice requirements.

4

Key recommendations and conclusions

This paper has outlined two use cases of how AI can be used in post-trade custody and securities services to generate new value and benefits for the industry: client segmentation and dealing with settlement failure. The benefits range from speed, efficiencies and other improvements to the existing settlement lifecycle, to new roles and employment opportunities, with human expertise needed in every step of algorithm and AI lifecycle management.

As digital applications become more powerful and widespread, good governance and effective controls will play an increasingly important role. Custodians' 'right to play' comes from the new expertise in technology, knowledge in the technology's intersections with laws, regulations and best practices, the ability to reliably execute the activities of safekeeping assets, and to ensure investor and asset protection in a timely and cost-effective way. Bringing AI applications to the table without compromising these responsibilities will mean staying informed of how these models work, what they contribute to the decision-making process, what risks are involved and the process and how great the impact would be of any errors in the system.

Proactively identifying and addressing emerging issues – and working collaboratively to address them – will be a key part of any push forward for the industry. This is not a straightforward task. The issues are multidisciplinary and include legal and regulatory considerations, alongside the inherent complexities of AI/ML technologies, data, financial products and the large stakeholder ecosystem these solutions will feed into.

With this in mind, we have laid out some preliminary recommendations that we believe can help lay the foundation for continued successes in rolling out AI and ML applications in the in the industry.

> "As more and more AI solutions are implemented – improving everything from front-office activities and compliance, to risk management and operations – it is important that we have robust governance measures in place. One particularly noteworthy paper is this respect is the European Union (EU) AI High-level expert group's "Ethics guidelines for trustworthy AI", which outlines seven key requirements – from an ethical, governance and robustness standpoint – that AI systems should meet in order to be deemed trustworthy. This paper acts as a core foundation for Deutsche Bank's own set of AI ethical principles"
>
> Steven Hondelink, Head of Securities Services, EMEA

4.1 Foster a focussed AI agenda at the industry level

A focussed AI agenda at a securities post-trade industry level (similar to what has been achieved with central bank digital currencies or the shortened securities settlement cycle) stands to help identify constraints, risks and approaches and enable a more sustainable ecosystem of innovation.[8]

Regulators have already played a key role in fostering the right environment by creating 'regulatory sandboxes' for focussed innovation expert group dialogues with shared outcomes, such as those led by the Fintech AI Public Private Forum initiated by the Bank of England.[9] These are valuable and industry participants can leverage on these developments to share best practices for ethical controls and to learn by experience and action rather than by theory.[10] A global securities industry AI agenda can start to take shape within industry conference agendas.

4.2 Source and embed the right expertise

AI capabilities bring with them the potential for new non-financial risks. Unfair bias or poorly generalised AI models can lead to operational and financial costs that impinge on asset safety and investors. These concerns can be accentuated by a lack of expertise in tuning, monitoring and managing AI/ML risks. This, in turn, can lead to a multitude of perceived or actual issues, including a lack of control or accountability over these new, powerful engines of change.

Having the right expertise at hand to provide risk management controls is therefore a priority, with specialised skills needed right across the business, engineering and internal control functions. This pits the post-trade securities industry in a bid for talent against a number of other industries, with equally if not more attractive profiles.

Given the requirements to combine domain knowledge with AI expertise, the industry should consider training and certification programs that can upskill existing staff, while also playing an active role in attracting these new forms of expertise. Industry participants, associations and academic institutions can play a valuable role in delivering the right training to help ensure incoming specialists have the right blend of relevant AI and related risk-management expertise, contextualised to the securities markets.

These sorts of training initiatives can also address some of the negative perceptions of AI related to the perceived risk of job losses – instead showing how these developments can lead to more interests and discoveries with in the post-trade and custody industry. As the trend gathers pace, it could also result in the creation of more start-ups and entrepreneurs, further enriching the industry landscape.

4.3 Clarify auditing requirements

AI-related controls and processes will inevitably be subjected to thorough audits, with AI ethical principles as a reference basis. Having a set of agreed auditing criteria and guidelines across the industry can facilitate more objective and practical recommendations for participants looking to invest in the technology.

This could also facilitate the operationalisation of risk controls based on AI ethical principles. The guidelines can also address challenges, such as the need to scale control levels in relation to the risks associated with applications. Since the materiality of AI applications will vary hugely across different use cases, it would be unhelpful for the full range to be subject to the same simple and inflexible audit control checklist.

For example, if there is an audit for 'human-over-the-loop' or 'human-in-the-loop' accountability controls, how should control requirements and gaps be interpreted in order to ensure that they are also practical to the team's development and amenable to the careers of the individuals involved? These are issues that those in the industry can tackle together. Over time, refined operationalisation of controls of explainability, transparency and accountability should lead to a better understanding of AI-related risks, which can facilitate broader adoption of AI-based services and support industry confidence to advance on AI applications.

4.4 Bring it all together

The creation of new public goods, such as publicly available data sets, could prove highly valuable to AI endeavours, spurring market innovation. At the moment, a lack of real-time securities information from financial market infrastructure, or fragmented data pools, makes algorithm training harder and can force firms to use algorithms that are more complex than required. Making the data available sooner rather than later can eliminate some of these avoidable complexities.

Applying AI to large data sets also brings in legal and regulatory considerations such as data confidentiality, outsourcing and product liability – issues that can be complex, especially when multiple jurisdictions are involved. These will have to be managed, but the industry can streamline the challenge by identifying and defining a set of enablers and an ecosystem for AI-based applications, establishing a roadmap to bring all these together and lay the foundations for new growth.

Just as car safety belts and collective road rules are part of the same ecosystem that make the highways safer and better for all drivers, a collective focus that incentivises the use of AI in a safe and responsible manner is a key catalyst for the adoption of AI in the securities services industry. By establishing and embracing these principles, the securities post-trade and custody industry can step forward into an exciting future defined by digital intelligence-based competitiveness.

5

Glossary

AI is a technical and complex field, involving a good deal of specialist language. We have aimed to keep this to a minimum in the interests of readability. Where we have used technical language, we have endeavoured to define it below. Readers should note that, as in any fast-moving and complex field, definitions often vary, overlap, and evolve, but the below should serve as broad-brush definitions for the purpose of understanding this paper.

AI black box. A situation where humans are no longer able to adequately explain or interpret the outputs of an AI model. This type of situation typically represents grounds for ethical and governance concerns.

Algorithm. A set of rules, calculations and steps that automatically processes input data and the relationships between the data to provide a pre-defined form of output designed to address a given issue.

Artificial Intelligence (AI). The simulation of human intelligence in machines to automate, accelerate and increase the complexity of tasks that would usually have to be performed by people.

Central Securities Depository Regulation (CSDR). The regulation relating to securities settlement and central securities depositaries (CSDs) that entered into force on 17 September 2014. It applies to CSDs that are based in the European Union and their participants. The principal objectives of the CSDR are to:

— Establish an enhanced level playing field among CSDs;
— Harmonise the different CSD rules in the European Union; and
— Increase the safety and efficiency of securities settlement and the settlement infrastructures in the European Union (EU).

Notably, the CSDR involves a penalty fee for participants whose transactions are not settled on time.

Classification vs. clustering. Classification uses pre-defined categories to group data; clustering analyses unclassified data to identify similarities between data points and group or "cluster" them accordingly.

Classification-based algorithm. An algorithm that classifies or categorises new data into distinct groups.

Classifiers. Classifiers place results into categories and can be applied to situations with either two classes (e.g. fail / succeed), or with multiple classes/ outcomes (e.g. tagging a transaction as: urgent; likely to fail; no penalty; or regulatory report needed).

Client segmentation. The practice of dividing clients into discrete groups to reflect their similarities.

Expectation-Maximisation Algorithm. A type of cluster algorithm that produces maximum-likelihood estimates for the presence of certain variables. It does this by first estimating the values for the variables, then optimising the model, then repeating these two steps until convergence. It is an effective and general approach that is most commonly used for density estimation with missing data.

Explainability. A measure of how easy it is for humans to explain how an AI model works and derives its results.

F1 Score. A measure of an AI model's performance, based on the balance between precision (the proportion of the outputs are correct) and completeness (the proportion of the overall relevant data set that is successfully converted into an output).

Gradient-boosting. A type of algorithm designed to "boost" the accuracy of AI outputs by predicting cases where an existing algorithm will perform poorly.

Hyperplanes. The boundaries established between different classifications of data points. Data points falling on either side of a hyperplane can be attributed to different classes.

Inferred bias. The potential for discriminatory conclusions to be drawn from statistical calculations and other attributes based on limited data.

Interpretability. A measure of how easy it is for humans understand what an AI output means in practice.

K-Means Clustering. One of the simplest and popular unsupervised machine learning algorithms. The objective of K-means clustering is to group similar data points together and discover underlying patterns.

Machine Learning (ML). A subset of artificial intelligence where models automatically learn from and adapt to new data without being assisted by humans.

Materiality. The degree of impact of a wrong decision. The greater the impact of a mistake generated via an AI model, the greater the materiality. For example, if AI is used in medical surgery, the materiality is high, since a mistake could permanently harm the patient or even kill them.

Model or Concept Drift. A phenomenon where the relationship between two variables in the input data of an AI model starts to change, causing irregularities in the output.

Predictors (*also known as* "regression predictors"). A type of algorithm that attempts to identify relationships between input variables in order to output a model that can then "predict". For example, a regression predictor could predict the value of a house, given historical transactions based on inputs such as size, number of rooms, and location.

Random Forest. A type of classification algorithm that works by combining multiple decision trees with a view to generating a more accurate outcome than relying on just one.

Recommender. A sub-type of predictors that outputs best next actions and alternatives based on patterns and relationships between choices.

Regression-based algorithm. An algorithm that provides predicted output values based on the input data features.

S2-Predict. An AI solution developed by Deutsche Bank Securities Services to predict and minimise settlement failures.

Semi-supervised learning. A type of machine learning that trains AI models using a small amount of labelled data with a large amount of unlabelled data during training.

Shapley value. A measure devised by US mathematician and economist Lloyd Shapley and used to explain the predictions of a complex predictive model or "black box". Shapley values correspond to the contribution of each of a model's features towards pushing the prediction away from the expected value.

Supervised learning. A type of machine learning that involves an AI model learning from examples and training data sets to produce desired outcomes.

Unsupervised learning. A type of machine learning where the model looks for structure and relationships in unlabelled data sets.

References

1. https://www.wired.com/story/ai-great-things-burn-planet/
2. https://cib.db.com/insights-and-initiatives/white-papers/csdr-settlement-discipline-regime-client-toolkit.htm
3. https://www.finextra.com/newsarticle/36727/deutsche-bank-to-provide-early-warning-system-for-settlement-fails
4. Russell, Stuart, Norvig, Peter, “Artificial Intelligence: A Modern Approach”
5. Zhang, Z., Yang, F., Wang, H., Hu, X., “Contextual Local Explanation for Black Box Classifiers”
6. https://www.db.com/news/detail/20191112-the-responsible-use-of-ai-in-financial-services-from-principles-to-practice?language_id=1
7. Storm, M, “Modern Britain, Global Leader in Ethical AI”
8. https://flow.db.com/cash-management/when-will-digital-currencies-become-mainstream
9. https://www.bankofengland.co.uk/events/2020/october/fintech-ai-public-private-forum
10. https://www.bankofengland.co.uk/speech/2020/dave-ramsden-opening-remarks-for-the-launch-of-aippf

www.db.com/flow

flow SPECIAL