Douglas Araujo
Douglas.Araujo@bis.org

Sebastian Doerr
Sebastian.Doerr@bis.org

Leonardo Gambacorta
Leonardo.Gambacorta@bis.org

Bruno Tissot
Bruno.Tissot@bis.org

Online annex for "Artificial intelligence in central banking"

This annex provides additional information on the use cases discussed in the text. A broad list of use cases can be found in Araujo et al (2022, 2023) and Beerman et al (2021). Table A1 lists a broader range of use cases by central banks, ordered by main method and area.

To provide high-quality micro data, central banks are progressively using machine learning techniques. Isolation forests are particularly suitable for the large and granular data sets typical of central banks, owing to their scalability and ability to identify outliers regardless of the shape of the data's distribution. For example, the Bank of Israel[1] uses isolation forest and other methods to detect anomalies in its extensive derivatives data set. Central banks have also surmounted a primary limitation of isolation forests – their restrictive use to numerical variables – by developing new ways to include other data types (eg the currency of a transaction or the counterparty's country). For example, given that the data underlying the calculation of the euro short-term rate (€STR) predominantly comprises non-numerical variables, the ECB[2] first applies an algorithm to convert categorical variables into numerical ones, and then makes full use of the data to find outliers.

Another example involves the Deutsche Bundesbank[3] collaborating with the German Research Center for Artificial Intelligence. Benchmarking against manually identified outliers, this project assesses the efficacy of isolation forests and other algorithms in detecting outliers in large data sets used for official euro area statistics. The project also investigates the benefits of a two-step approach: initially, the model autonomously identifies potential outliers, which are then reviewed by experts who provide feedback to refine the algorithm. This approach balances the value of domain expertise with the costs of human inputs. By analysing different methods to explain the outlier classification, the project aims to overcome the issue of "black box" machine learning models lacking "explainability". The findings indicate that several algorithms, especially when enhanced by human feedback, effectively detect outliers. Isolation forests, in particular, demonstrate strong performance across various data sets. Moreover, explainable machine learning methods provide experts with guidance on which data points warrant manual verification.

In the area of macro-financial analysis, the Bank of England[4] uses neural networks to break down services inflation into different components: past inflation dynamics, inflation expectations, the output gap and international prices. To estimate the contribution of each component, it employs four interlinked neural networks that jointly predict services inflation. Each of these networks considers multiple economic variables that are relevant to each component. For example, the input data for the output gap network are all related to economic activity, such as aggregate and sectoral unemployment rates and industrial production. The outputs of these networks are time series, each representing a component's contribution

1 S Kamenetsky Yadan, "Anomaly detection methods and tools for big data", *IFC Bulletin*, no 57, 2021.

2 M Accornero and G Boscariol, "Machine learning for anomaly detection in data sets with categorical variables and skewed distributions", *IFC Bulletin*, no 57, 2021.

3 T Cagala, J Hees, D Herurkar, M Meier, N-T Nguyen, T Sattarov, K Troutman and P Weber, "Unsupervised outlier detection in official statistics", *IFC Bulletin*, no 57, 2021.

4 M Buckmann, G Potjagailo and P Schnattinger, "Dissecting UK service inflation via a neural network Phillips curve", Bank of England *Bank Underground*, 10 July 2023.

to services inflation. These neural network models can process more input variables than traditional econometric models, allowing central banks to use granular data sets instead of aggregate data. Another advantage is neural networks' ability to reflect complex non-linearities in the data, which is of particular value during episodes of shifting inflation dynamics. This "neural network Phillips curve" provides insights into inflationary pressures, attributing the recent UK services inflation primarily to higher input costs and price inertia. Similarly, the Bank of Spain, BIS, Bundesbank, ECB and Federal Reserve Bank of Chicago[5] are harnessing neural networks to estimate a broader class of macroeconomic models with non-linearities and heterogenous agents.

The usefulness of neural networks for analysing periods of non-linear dynamics is further illustrated by the Bank of Korea's[6] nowcasting model. The model uses a neural network well suited for time series analysis (a so-called long short-term memory network) that takes as input relevant real-time economic data, together with results from a traditional econometric nowcasting model. The combination of a neural network and a traditional nowcasting model achieves significantly better forecasting performance, especially during tumultuous periods such as the Covid-19 pandemic.

To obtain real-time estimates of inflation expectations, recent work by the Bank of France[7] uses random forests on X (formerly Twitter) data. It analyses all tweets from users that have retweeted a post by the Bank of France, on the assumption that they have an informed interest in monetary policy and the economy. The first random forest model identifies tweets related to prices, which then feed into a second random forest model that classifies each tweet as reflecting inflation, deflation or other expectations. The difference in the daily tweet counts for higher versus lower inflation provides a gauge for inflation expectations. This indicator is highly correlated with traditional survey-based metrics and effectively predicts future inflation, thus serving as a dynamic index of inflation perceptions.

Similarly, Bank Indonesia[8] uses machine learning models, including random forests and neural networks, to construct real-time indicators from text data to assess policy credibility. In a first step, over 12,000 sentences from financial news were classified by monetary policy experts as relevant or not to policy credibility perceptions, and, if relevant, as positive or negative. Four components of policy credibility were considered: (i) policy formulation; (ii) effectiveness; (iii) coordination with other authorities; and (iv) policy communication. The process then involves initial training of a model to identify credibility-related sentences, followed by another model that classifies these sentences into positive or negative perceptions. This sequence of models can then classify news at each date to measure credibility perceptions for each component over time. Periods with higher credibility indices correlate with inflation predictions which are more aligned with central bank targets, indicating better anchored inflation expectations.

Focusing on data with a long time series, a Federal Reserve Board[9] research project uses FinBERT, an open source LLM fine-tuned with financial news, to summarise economic conditions over time. This model processes anecdotal texts from interviews with business contacts, economists, market experts and other sources into a (positive or negative) sentiment value at each period. The interviews are taken from the Beige Book compiled for Federal Open Market Committee meetings since 1970 in order to compile a long-

[5] J Fernández-Villaverde, G Nuño, G Sorg-Langhans and M Vogler, "Solving high-dimensional dynamic programming problems with deep learning", *mimeo*, 2020; J Fernández-Villaverde, J Marbet, G Nuño and O Rachedi, "Inequality and the zero lower bound", *BIS Working Papers*, no 1160, 2024; and H Kase, L Melosi and M Rottner, "Estimating nonlinear heterogenous agents models with neural networks", *mimeo*, 2022.

[6] H C Yi, D Choi and Y Kim, "Dynamic factor model and deep learning algorithm for GDP nowcasting", *Bank of Korea Economic Analysis*, vol 28, no 2, 2022.

[7] J Denes, A Lestrade and L Richardet, "Using Twitter data to measure inflation perception", *IFC Bulletin*, no 57, 2021.

[8] M Abdul Jabbar, O Wibisono, A Widjanarti and A Zulen, "Machine learning for measuring central bank policy credibility and communication from news", *IFC Bulletin*, no 59, 2022.

[9] S Du, K Guo, F Haberkorn, A Kessler, I Kitschelt, S J Lee, A Monken, D Saez, K Shipman and S Thakur, "Do anecdotes matter? Exploring the Beige Book through textual analysis from 1970 to 2023", *IFC Bulletin*, forthcoming.

standing sentiment time series. The sentiment index can nowcast GDP and predict the probability of recessions. Moreover, by categorising texts into specific economic topics (eg consumer demand and credit conditions), the model can extract topic-specific sentiments. For instance, it identified that in 2022 negative sentiment due to higher prices and supply chain issues was partly offset by positive views on travel and consumer activity.

An approach for overseeing payment transactions involves the use of unsupervised learning methods to automatically single out transactions that are worth monitoring more closely. This can be particularly helpful to detect non-linear dynamics in the data, such as bank runs. Collaborative research by the Bank of Canada and De Nederlandsche Bank[10] explores the use of auto-encoder models within a Canadian retail payment system. Auto-encoders are neural networks where both the input and output layers are the same, so their hidden layers learn the characteristics of a "typical" payment transaction. The model effectively identified a simulated pattern of significant bank deposit withdrawals over several days. Similarly, the Central Bank of Ecuador[11] implemented auto-encoders in its interbank payment system. These models not only detected simulated bank runs but also identified a range of real-life anomalies, including operational disruptions in a major bank. These instances demonstrate the efficacy of machine learning techniques in providing actionable insights across various scenarios.

To consolidate the wealth of information in one place and facilitate supervisory analysis of unstructured data, the ECB[12] created the platform Athena. This suite of tools uses BERT[13] models fine-tuned on supervisory content together with NLP techniques to classify public and supervisory documents, undertake sentiment analyses and identify trending topics. Athena also deploys a neural network coupled with specialised open source tools that identify mentions of supervised entities. Athena thereby facilitates supervisors' access to relevant information scattered across millions of documents.

In a similar vein, the Federal Reserve system[14] applies a NLP tool called Language Extraction (LEX) to help supervisors find relevant information amidst large collections of different documents. By training on a large body of text combined with an expert-defined lexicon of words and clauses relevant to different tasks, LEX automates the discovery of excerpts containing information on risks. LEX allows for text analysis, including document summarisation, on a much broader scale than the previous process of manually perusing documents. It has proved particularly successful at discovering sentences that examiners missed. As the results are readily available to supervisors in a variety of formats, including as visual interfaces, LEX halves the time spent reviewing document submissions. Key use cases include finding relevant material to conduct supervision on specific areas such as climate change, cloud adoption and cyber risk.

Beyond boosting supervisors' productivity, text models can also improve the consistency of supervisory communication with banks. The Central Bank of Malaysia[15] had experts classify sentences of past letters to banks into categories such as "forceful" or "neutral". These labelled data served as input to a BERT-like LLM that first learns to classify unseen sentences and then informs bank supervisors whether draft letters match the desired tone. A related tool, powered by a version of BERT that focuses on understanding text at the sentence level, facilitates the search of sentences from past letters based on similarity, offering the possibility of using similar text adapted to a new context.

[10] L Sabetti and R Heijmans, "Shallow or deep? Detecting anomalous flows in the Canadian automated clearing and settlement system using an autoencoder", *DNB Working Papers*, no 681, 2020.

[11] J Rubio, P Barucca, G Gage, J Arroyo and R Morales-Resendiz, "Classifying payment patterns with artificial neural networks: an autoencoder approach", *IFC Bulletin*, no 57, 2021.

[12] ECB, "Suptech: thriving in the digital age", *Supervision Newsletter*, 15 November 2023.

[13] BERT stands for Bidirectional Encoder Representations from Transformers (BERT), a LLM based on the transformer architecture.

[14] K Beerman, J Prenio and R Zamil, "Suptech tools for prudential supervision and their use during the pandemic", *FSI Insights on policy implementation*, no 37, 2021.

[15] J Tan, C K Shum and M Akmal Amri, "Supervisory letter writing app: expediting letter drafting and tone consistency", *IFC Bulletin*, no 57, 2021.

The Central Bank of Brazil[16] created ADAM to identify individual borrowers for which lenders underestimate potential credit losses. ADAM uses three classification models that take in about 300 data points from borrowers and were trained on almost 11,000 supervisory analyses. The first model identifies borrowers likely to result in a 100% loss. It is a collection of decision trees that learn sequentially from each other's errors through so-called extreme gradient boosting. ADAM uses neural networks that include the first layers of a network trained on the 100% loss sample to improve identification of the smaller set of borrowers with expected losses of 70% or 50%. ADAM can analyse three million borrowers in one day. Its main output is a list of underprovisioned borrowers, with an explanation for each prediction and how correct provisioning would affect a lender's capital ratio.

Selected list of central bank use cases of machine learning — Table A1

	Application type			
Main method	Information collection	Macro/financial analysis for monetary policy	Payments oversight	Supervision
Tree-based methods	Banco de Portugal, Bank of Israel, Deutsche Bundesbank, ECB, Magyar Nemzeti Bank	Bank Indonesia, Bank of France, Reserve Bank of Australia, Reserve Bank of New Zealand	Central Bank of Iceland	Bank of France, Bank of Italy, Bank of Japan, Banco de Portugal, Bank of Spain, Central Bank of Brazil
Neural networks	ECB	Bangko Sentral ng Pilipinas, Bank Indonesia, Bank of Canada, Bank of Korea, Central Bank of Chile, Central Bank of Malaysia, Bank of Canada, Bank of England, Bank of Spain, Deutsche Bundesbank, ECB, Federal Reserve	Bank of Canada, Bank of Italy, Bank of Thailand, Central Bank of Ecuador, De Nederlandsche Bank	Bank of France, Bank of Greece, Central Bank of Brazil, Deutsche Bundesbank, Hong Kong Monetary Authority
Large language models	Deutsche Bundesbank	Bangko Sentral ng Pilipinas, Bank Indonesia, Bank of Korea, Deutsche Bundesbank, Federal Reserve	Bank of Korea	Central Bank of Malaysia, ECB, Federal Reserve
Other techniques	De Nederlandsche Bank, Deutsche Bundesbank	Bank of Italy, Czech National Bank, South African Reserve Bank		Bank of Canada, Bank of Slovenia, Bank of Spain, Bank of Thailand, Central Bank of the Republic of Austria,[1] ECB,[1] Federal Reserve, Monetary Authority of Singapore[1]

[1] Specific technique not disclosed publicly.

Sources: Araujo et al (2022, 2023); Beerman et al (2021); national central banks; IFC regular stocktaking exercises.

[16] K Beerman, J Prenio and R Zamil, "Suptech tools for prudential supervision and their use during the pandemic", *FSI Insights on policy implementation*, no 37, 2021.