BIS Bulletin

No 84

Artificial intelligence in central banking

Douglas Araujo, Sebastian Doerr, Leonardo Gambacorta and Bruno Tissot

23 January 2024

BIS Bulletins are written by staff members of the Bank for International Settlements, and from time to time by other economists, and are published by the Bank. The papers are on subjects of topical interest and are technical in character. The views expressed in them are those of their authors and not necessarily the views of the BIS. The authors are grateful to Bryan Hardy and Galo Nuño for comments, Ilaria Mattei and Krzysztof Zdanowicz for excellent research assistance, and to Louisa Wagner for administrative support.

The editor of the BIS Bulletin series is Hyun Song Shin.

This publication is available on the BIS website (www.bis.org).

ISSN: 2708-0420 (online)
ISBN: 978-92-9259-738-2 (online)

Douglas Araujo
Douglas.Araujo@bis.org

Sebastian Doerr
Sebastian.Doerr@bis.org

Leonardo Gambacorta
Leonardo.Gambacorta@bis.org

Bruno Tissot
Bruno.Tissot@bis.org

Artificial intelligence in central banking

Key takeaways

- *Central banks have been early adopters of machine learning techniques for statistics, macro analysis, payment systems oversight and supervision, with considerable success.*
- *Artificial intelligence brings many opportunities in support of central bank mandates, but also challenges – some general and others specific to central banks.*
- *Central bank collaboration, for instance through knowledge-sharing and pooling of expertise, holds great promise in keeping central banks at the vanguard of developments in artificial intelligence.*

Long before artificial intelligence (AI) became a focal point of popular commentary and widespread fascination, central banks were early adopters of machine learning methods to obtain valuable insights for statistics, research and policy (Doerr et al (2021), Araujo et al (2022, 2023)). The greater capabilities and performance of the new generation of machine learning techniques open up further opportunities. Yet harnessing these requires central banks to build up the necessary infrastructure and expertise. Central banks also need to address concerns about data quality and privacy as well as risks emanating from dependence on a few providers.

This Bulletin first provides a brief summary of concepts in the machine learning and AI space. It then discusses central bank use cases in four areas: (i) information collection and the compilation of official statistics; (ii) macroeconomic and financial analysis to support monetary policy; (iii) oversight of payment systems; and (iv) supervision and financial stability. The Bulletin also summarises the lessons learned and the opportunities and challenges arising from the use of machine learning and AI. It concludes by discussing how central bank cooperation can play a key role going forward.

Overview of machine learning methods and AI

Broadly speaking, *machine learning* comprises the set of techniques designed to extract information from data, especially with a view to making predictions. Machine learning can be seen as an outgrowth of traditional statistical and econometric techniques, although it does not rely on a pre-specified model or on statistical assumptions such as linearity or normality. The process of fitting a machine learning model to data is called *training*. The criterion for successful training is the ability to predict outcomes on previously unseen ("out-of-sample") data, irrespective of *how* the models predict them. This section describes some of the most common techniques used in central banks, based on the regular stocktaking exercises organised in the central banking community under the umbrella of the BIS Irving Fisher Committee on Central Bank Statistics (IFC).

Tree-based methods are flexible machine learning algorithms that can tackle a wide range of tasks. *Decision trees* group individual data points by sequentially partitioning data into finer categories according to specific characteristics of interest. For example, a tree may first sort houses (the input data) into those with more than three rooms and those with at most three, and then partition houses in each of these subgroups into those built before 1990 and those built after, and so on. The resulting finer partitioning of houses can then be compared with a particular dimension of interest (the output) to see how well the partitioning matches an attribute of interest. For instance, capturing how house prices vary across the finer partitioning would be a way to group similar houses in terms of their price.

Random forests combine several trees trained on different slices of the same data to improve prediction out of sample while guarding against the risk of overfitting the training data sample. Random forests and related models can be seen as a more flexible form of regression analysis, as they predict output from the explanatory variables of interest (Athey and Imbens (2021)). In addition, tree-based methods can serve as an exploratory tool to glean patterns in the data without imposing a model structure. For instance, they can classify data points into similar categories. In the same spirit, forests can be deployed in identifying outliers by means of *isolation forests*, a method that singles out the data points that can be isolated from others.

Neural networks are perhaps the most important technique in machine learning, with widespread uses even for the latest generation of models. Their main building blocks are *artificial neurons*, which take multiple input values and transform them in a non-linear way to output a single number – like logistic regressions. The artificial neurons are organised to form a sequence of *layers* that can be stacked: the neurons of the first layer take the input data and output an activation value. Subsequent layers then take the output of the previous layer as input, transform it and output another value, and so forth. This way, similar to neurons in the human brain, an artificial neuron's output value is akin to an electrical impulse transmitted to other neurons. A network's *depth* refers to the number of layers. Each neuron's constant and weights attached to the output of previous layers' neurons are collectively called *parameters*; they determine the strength of connections across neurons and layers. These parameters are improved iteratively during training. Deeper networks with more parameters require more training data but predict more accurately. Neural networks are behind face recognition or voice assistants in mobile phones and underlie the most significant recent innovations in AI.

Transformers, unveiled in 2017, drastically improved the performance of neural networks in natural language processing (NLP) and enabled the rise of large language models (LLMs). Rather than just relating a word to those near it, transformers attempt to capture the relationship between the different components of a text sequence, even if they are far apart in the sentence. This allows the model to better understand the context and hence different meanings a word can have. For example, the meaning of the word "bank" differs when it appears in the sentence "I'll swim across the river to get to the other bank" versus "I crossed the street to go to the bank". Transformers unlocked use cases of NLP that require dealing with long streams of text and gave rise to the most recent advances in LLMs, such as ChatGPT.

LLMs underlie the rapid rise of generative AI ("gen AI"), which generates content based on suitable prompts, and can perform tasks beyond language recognition. LLMs are neural networks that are trained to predict the next word in a given sequence of text. To perform this task, LLMs learn to absorb all the written knowledge on which they were trained. As a result, their prediction is usually accurate even for texts that require nuance or field knowledge. LLMs can be fine-tuned for specific tasks with specialised data. For example, ChatGPT is based on an LLM refined with human feedback to generate more useful responses. Key characteristics of gen AI are that it can be used not just by a small set of specialists but by virtually everybody and that it can easily extract insights from unstructured data.

Machine learning and AI in central banks: use cases

What are the current use cases of machine learning and AI in central banks? They can best be organised by scope: (i) information collection and statistical compilation; (ii) macroeconomic and financial analysis to support monetary policy; (iii) oversight of payment systems; and (iv) supervision and financial stability. This section provides relevant examples in each area. More information on the selected examples, as well as a broader list of use cases, can be found in the annex.

Information collection

Ensuring the availability of high-quality data as inputs for economic analysis and for statistics compilation and production is a major challenge for central banks. Issues include data cleaning, sampling, representativeness and matching new data to existing sources. The steadily increasing volume and complexity of data necessitate efficient and flexible data quality tools.

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. There are also benefits to a two-step approach: initially, a 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, this approach can overcome the issue of "black box" machine learning models lacking "explainability", which is discussed below. Moreover, explainable machine learning methods provide experts with guidance on which data points warrant manual verification.

Macroeconomic and financial analysis to support monetary policy

Central banks rely extensively on macroeconomic and financial analysis to support monetary policy. In a complex environment, a significant challenge is to efficiently extract information from a wide array of traditional and non-traditional data sources. Machine learning offers valuable tools in this area.

Neural networks can, for example, break down services inflation into different components, revealing how much inflation is due to past price increases, inflation expectations, the output gap or international prices. Such models can process more input variables than traditional econometric ones, allowing central banks to use granular data sets instead of more aggregate data. Another advantage is neural networks' ability to reflect complex non-linearities in the data, which can help modellers to better capture non-linearities, from the zero lower bound to unequal asset holdings and shifts in inflation dynamics.

Other use cases are obtaining real-time estimates (nowcasts) of inflation expectations or summarising economic conditions over time. For example, random forest models can identify social media posts that are related to prices and then feed them into another random forest model that classifies each post as reflecting inflation, deflation or other expectations. The difference in the daily counts of social media posts for higher versus lower inflation gauges inflation expectations. Similarly, social media posts can be used to track the credibility of central bank monetary policy with the wider public.

Another example is the use of open source LLMs fine-tuned with financial news to summarise economic condition narratives over a long time span. Models can process eg anecdotal texts from interviews with entrepreneurs, economists and market experts to produce a time series of their (positive or negative) sentiment value. The sentiment index can then be used to nowcast GDP or predict recessions.

Adapting LLMs to central banking terminology can bring further gains, as shown by the central bank language models (CB-LM) project developed at the BIS (Gambacorta et al (2024)). This approach uses thousands of central bank speeches and research papers compiled by the BIS Central Bank Hub to adapt widely used open source foundation LLMs issued by Google and Meta. This additional training focused on central banking texts increased accuracy from 50–60% to 90% in interpreting central bank terminology and idioms. It has also improved performance in tasks such as classifying Federal Open Market Committee policy stances and predicting market reactions to monetary policy announcements.

Oversight of payment systems

Well functioning payment systems are fundamental to the stability of the financial system, yet the vast amount of transaction data, often with a highly skewed distribution, poses challenges in distinguishing anomalous transactions from regular ones. Correctly identifying anomalous payments is crucial to addressing issues such as potential bank failures, cyber attacks or financial crimes in a timely manner. Money laundering, in particular, undermines the integrity and safety of the global financial system.

The BIS Innovation Hub's Project Aurora uses synthetic money laundering data to compare fraudulent payment identification by various traditional and machine learning models (BISIH (2023)). The models, which include isolation forests and neural networks, undergo training with known (synthetic) money laundering transactions and then predict the likelihood of money laundering in unseen data. Machine learning models outperform the rule-based methods prevalent in most jurisdictions or traditional logistic regressions. Graph neural networks, which take payment relationships as input, identify suspect transaction networks particularly well. These models can function effectively even with data pooling that safeguards confidentiality, suggesting that cooperation to jointly analyse multiple databases can be secure and beneficial. This illustrates the potential for more cooperation between authorities.

Another approach for overseeing payment transactions involves the use of unsupervised learning methods to automatically single out transactions that are worth closer inspection. For example, auto-encoder models, neural networks where both the input and output layers look at the same data, distinguish typical from anomalous payments and can detect non-linear dynamics such as bank runs. In simulations, these models effectively identified patterns of significant bank deposit withdrawals over several days. Auto-encoders also identified a range of real-life anomalies in payment systems, including operational disruptions among important domestic banks.

Supervision and financial stability

Supervisors analyse a broad range of data sources to efficiently oversee financial institutions. These sources include text documents such as news articles, internal bank documents or supervisory assessments. Sifting through this wealth of information to extract relevant insights can be time-consuming, and with the ever increasing volume of data it becomes nearly insurmountable. Moreover, analyses related to climate and cyber risks have emerged as supervisory priorities, but they lack the comprehensive data infrastructure already in place for more "traditional" risks.

One avenue pursued by many central banks is to consolidate the wealth of information in one place and help supervisory analysis of unstructured data. For example, models fine-tuned on supervisory content together with NLP techniques can classify public and supervisory documents, undertake sentiment analyses and identify trending topics, as done in the ECB's platform Athena. Training models on a large body of text combined with an expert-defined lexicon of relevant words and clauses can also help automate the discovery of excerpts containing information on different risks. Such models, for example the Federal Reserve's LEX, facilitate supervisors' access to relevant information scattered across millions of documents and reduce the time spent reviewing document submissions. Classification models, leveraging tree-based techniques or neural networks, can also help identify individual borrowers for which lenders underestimate potential credit losses, a task for which the Central Bank of Brazil created ADAM. Neural networks that include the first layers of a trained network can improve identification of borrowers with high expected losses. Supervisors can then require financial institutions to provision exposures that are not sufficiently covered.

Balancing opportunities and challenges

The above examples illustrate the opportunities for machine learning and AI to tackle problems at the heart of central bank mandates. Yet there are also new challenges, some more general and others more specific to central banks.

A general challenge is the conflict between accuracy and "interpretability/explainability". Sophisticated machine learning models can become near perfect at prediction. But since many variables interact in complex and non-linear ways, it can be difficult to interpret how important different input variables are for the result. Good prediction can hence come at the cost of accepting that the underlying model is a "black box". This can, for example, make it challenging to assess discriminatory biases in algorithms, especially when these have been trained on biased data sets. Limited explainability further means that it is difficult to explain model behaviour in human terms; for example, why inflation is predicted to go up or why a mortgage application was rejected. For gen AI models, the issue goes even further, as they suffer from the "hallucination problem". These models might present a factually incorrect answer as if it were correct. The hallucination problem implies that LLMs need human supervision, especially in tasks requiring logical reasoning (Perez-Cruz and Shin (2024)).

For central banks, the use of unstructured data can offer valuable information that can help solve previously intractable problems. Manually converting unstructured data, in particular text, into structured form is time-consuming, prone to human error and infeasible at a larger scale. As the above examples make clear, LLMs can help central banks analyse a wide range of textual data, such as social media activity, financial news and central banks' own reports (confidential or public).

The use of unstructured and often personal data, however, poses new challenges in terms of legal frameworks and data privacy. Traditionally, most data were collected and hosted within public institutions with clearly defined access rights and sound data quality assurance processes. But now, large swathes of data are created by individuals and firms and reside with the private sector, sometimes with little documentation publicly available. Training or fine-tuning LLMs may require significant amounts of data, which can be obtained, for example by web scraping information from market platforms or social media, but for which legal frameworks often remain unclear about how and for what purposes they can be used. The availability of unstructured personal data also raises concerns about ethics and privacy. Citizens have a right to privacy and might feel uncomfortable with central banks scrutinising their data. While privacy-enhancing technologies are steadily improving, they are not yet a default in AI models.

Greater use of AI could also have profound implications for central banks' investments in information technology (IT) and human capital. Providing adequate computing power and software, as well as training existing staff, involves high upfront costs. Meanwhile, hiring new staff or retaining existing staff with the right mix of economic understanding and programming skills can be challenging: there is high demand for this resource, and public institutions often cannot match private sector salaries for top data scientists.

However, these investments could, over time, yield increased productivity. The above examples suggest that the use of machine learning and AI can markedly raise staff productivity – in particular in some time-intensive tasks that require cognitive work, such as summarising and extracting information from text (Brynjolfsson et al (2023), Noy and Zhang (2023)). For example, AI systems could act as "co-pilots" to human supervisory teams by learning from a combination of regulatory data, prior supervisory actions and broader market developments. AI could also improve analysis by freeing up economists' time for interpreting data rather than collecting and cleaning it. Yet AI will not make humans obsolete. Incorporating expert feedback can improve models and mitigate the hallucination problem. The business expertise of staff helps to identify where models add the most value as well as how to adapt them to central bank-specific tasks.

Finally, the rise of LLMs and generative AI has renewed concerns about dependence on a few external providers. Large economies of scale mean that the most powerful foundation models are provided by just a few large technology firms. Beyond the general risks that market concentration poses to innovation and economic dynamism, this high concentration of resources could create significant financial stability, operational and reputational risks. For example, greater reliance on LLMs and gen AI by just a few companies makes the financial system susceptible to spillovers from IT failures or cyber attacks on these providers. Outages among providers could also lead to operational risks for central banks and have repercussions for their ability to fulfil their mandates. The risk of operational problems leading to reputational costs looms large as central banks' greatest asset is the public's trust (Doerr et al (2022)). At the same time, if many institutions adopt the same few best in class algorithms, their behaviour during stress episodes might look increasingly alike and lead to undesirable phenomena such as liquidity hoarding, interbank runs and fire sales (Danielson and Uthemann (2023)).

These lessons underscore the benefits of cooperation among central banks and other public authorities. Knowledge-sharing and the pooling of expertise are well established in the central banking community, and central banks' public policy mandate gives considerable scope for cooperation, as well as to establish a community of practice for machine learning and AI. Central bank collaboration and the sharing of experiences could also help identify areas in which AI adds the most value and how to leverage synergies. Data standards could facilitate the automated collection of relevant data from various official sources, thereby enhancing the training and performance of machine learning models that use macroeconomic data (Araujo (2023)). Additionally, the sharing of code or pre-trained models hold much promise.

Central banking is particularly well suited for the application of machine learning techniques given the availability of structured and unstructured data as well as the need for rigorous analysis in support of policy. The synergies between machine learning and core central banking disciplines such as economics, statistics and econometrics are likely to place central banks at the vanguard of advances in AI.

References

Araujo, D K G (2023): "gingado: a machine learning library focused on economics and finance", *BIS Working Papers*, no 1122.

Araujo, D K G, G Bruno, J Marcucci, R Schmidt and B Tissot (2022): "Machine learning applications in central banking: an overview", *IFC Bulletin*, no 57.

——— (2023): "Data science in central banking: applications and tools", *IFC Bulletin*, no 59.

Athey, S and G Imbens (2021): "Machine learning methods that economists should know about", *Annual Review of Economics*, no 11, pp 685–725.

BIS Innovation Hub (BISIH) (2023): *Project Aurora: the power of data, technology and collaboration to combat money laundering across institutions and borders*, May.

Brynjolfsson, E, D Li and L Raymond (2023): "Generative AI at work", *NBER Working Papers*, no 31161.

Danielson, J and A Uthemann (2023): "On the use of artificial intelligence in financial regulations and the impact on financial stability", mimeo.

Doerr, S, L Gambacorta and J Serena (2021): "Big data and machine learning in central banking", *BIS Working Papers*, no 930.

Doerr, S, L Gambacorta, T Leach, B Legros and D Whyte (2022): "Cyber risk in central banking", *BIS Working Papers*, no 1039.

Gambacorta, L, B Kwon, T Park, P Patelli and S Zhu (2024): "CB-LMs: language models for central banking", *BIS Working Papers*, forthcoming.

Noy, S and W Zhang (2023): "Experimental evidence on the productivity effects of generative artificial intelligence", *Science*, vol 6654, no 381, pp 187–92.

Perez-Cruz, F and H S Shin (2024): "Testing the cognitive limits of large language models", *BIS Bulletin*, no 83.

Previous issues in this series

No 83 04 January 2024	Testing the cognitive limits of large language models	Fernando Perez-Cruz and Hyun Song Shin
No 82 20 December 2023	The contribution of monetary policy to disinflation	Pongpitch Amatyakul, Fiorella De Fiore, Marco Lombardi, Benoit Mojon and Daniel Rees
No 81 13 December 2023	Interest rate risk of non-financial firms: who hedges and does it help?	Ryan Banerjee, Julián Caballero, Enisse Kharroubi, Renée Spigt and Egon Zakrajšek
No 80 23 November 2023	Monetary policy, financial conditions and real activity: is this time different?	Fernando Avalos, Deniz Igan, Cristina Manea and Richhild Moessner
No 79 2 November 2023	Lessons from recent experiences on exchange rates, capital flows and financial conditions in emerging market economies	Pietro Patelli, Jimmy Shek and Ilhyock Shim
No 78 3 October 2023	Mapping the realignment of global value chains	Han Qiu, Hyun Song Shin and Leanne Si Ying Zhang
No 77 13 September 2023	Margins and liquidity in European energy markets in 2022	Fernando Avalos, Wenqian Huang and Kevin Tracol
No 76 7 September 2023	The oracle problem and the future of DeFi	Chanelle Duley, Leonardo Gambacorta, Rodney Garratt and Priscilla Koo Wilkens
No 75 19 May 2023	Disinflation milestones	Benoit Mojon, Gabriela Nodari and Stefano Siviero
No 74 13 April 2023	The changing nexus between commodity prices and the dollar: causes and implications	Boris Hofmann, Deniz Igan and Daniel Rees
No 73 11 April 2023	Stablecoins versus tokenised deposits: implications for the singleness of money	Rodney Garratt and Hyun Song Shin
No 72 11 April 2023	The tokenisation continuum	Iñaki Aldasoro, Sebastian Doerr, Leonardo Gambacorta, Rodney Garratt and Priscilla Koo Wilkens
No 71 29 March 2023	Fiscal and monetary policy in emerging markets: what are the risks and policy trade-offs?	Ana Aguilar, Carlos Cantú and Rafael Guerra

All issues are available on our website www.bis.org.