



Services

Delivering Improved Experience *Through* *Generative AI*

Practical AI Applications in Global Payments





Executive Summary

This whitepaper explores potential applications of generative artificial intelligence (AI) in global payments, focusing on practical implementations of AI that could improve banking services for clients.

It begins by envisioning **seamless banking** through automated processes, enhanced client servicing and fewer manual touchpoints. Next, it looks at how AI can **unlock value** by providing **hyper-customised insights**. Finally, it touches on **client focus** – a core Citi strength – and considers AI's application for listening, understanding and co-creation.

Underlying these possibilities are considerations around **ethics**, model **risk management** and **compliance**, which are addressed in the final section, along with emerging legal frameworks.

Foreword

The possibilities unlocked by generative AI are astounding. At Citi, we continually strive to harness emerging technologies that can benefit our clients and drive innovation across industries. Equipping our staff with new foundational AI capabilities can potentially help streamline processes, bring us closer to clients, and improve customer experiences. Similarly, as data becomes an increasingly vital asset, AI enables us to extract greater insights and value from clients' information streams.

The path to realising AI's full potential requires new governance frameworks and responsible practices. This white paper provides guidance for embedding ethics, security, and trustworthy AI principles. Building a robust AI strategy is essential for our long-term success.

At Citi, we are committed to empowering our clients with the knowledge and resources to navigate the AI-driven future successfully. This paper represents our dedication to that mission. As well as informing future AI initiatives at Citi and providing insights to the wider banking community, we hope clients will benefit from exploring these powerful use cases and looking to a future where AI can propel your organisation to new heights.

Debopama Sen

Citi Global Head of Payments

Introduction

With advancing computing power and the availability of vast datasets, AI has made remarkable progress in the last decade. Applications based on transformers, a type of neural network architecture optimised for natural language processing, and especially large language models, have captured the public imagination since 2023. Today, AI has already been integrated with many aspects of business and is becoming a transformative force reshaping the financial industry.

Banking is among the industries that could see the biggest impact as a percentage of their revenues from generative AI.¹ It has both text-rich datasets, and substantial numbers of text-based customer interactions.

Generative AI can draw on banks' process and product knowledge bases to tailor its communication with clients, providing fast and accurate answers, which might cover anything from payment format setup to payment status information. Further enabled by accelerated document digitization, banks can offer faster, higher quality and more relevant services, such as onboarding and updating client mandates.

From a regulatory perspective, while AI can help enhance compliance processes in the banking sector, it is crucial to carefully manage AI capabilities through safeguards and address a full spectrum of potential technology risks.

As banks continue to modernise their applications, improvements made possible by generative AI will have a significant impact on speed and consistency. This progress, in turn, will open an opportunity for new product lines with a faster time to market.

Seamless Banking

01 | Onboarding

Onboarding for larger corporate clients is complicated by a heavy reliance on documents. This process requires an ability to efficiently validate, understand and store the content of documents to correctly implement client instructions. Documents often lack standardised formats and are not fully digitised.

From a client perspective, faster, more efficient onboarding would be very desirable. This is why with the help of generative AI (GenAI) any subsequent market expansions requiring

extra documents could focus only on the incremental or updated information not previously provided. An AI enhanced onboarding process could also provide step-by-step guidance on which documents are required and offer quick validation whether documents already provided are complete, up to date and acceptable.

Once digitised, these documents can be converted into vector databases which allow GenAI models to retrieve only the data specific to a query, helping to ensure that documents remain relevant and up to date.

¹ <https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-AI-the-next-productivity-frontier#key-insights>

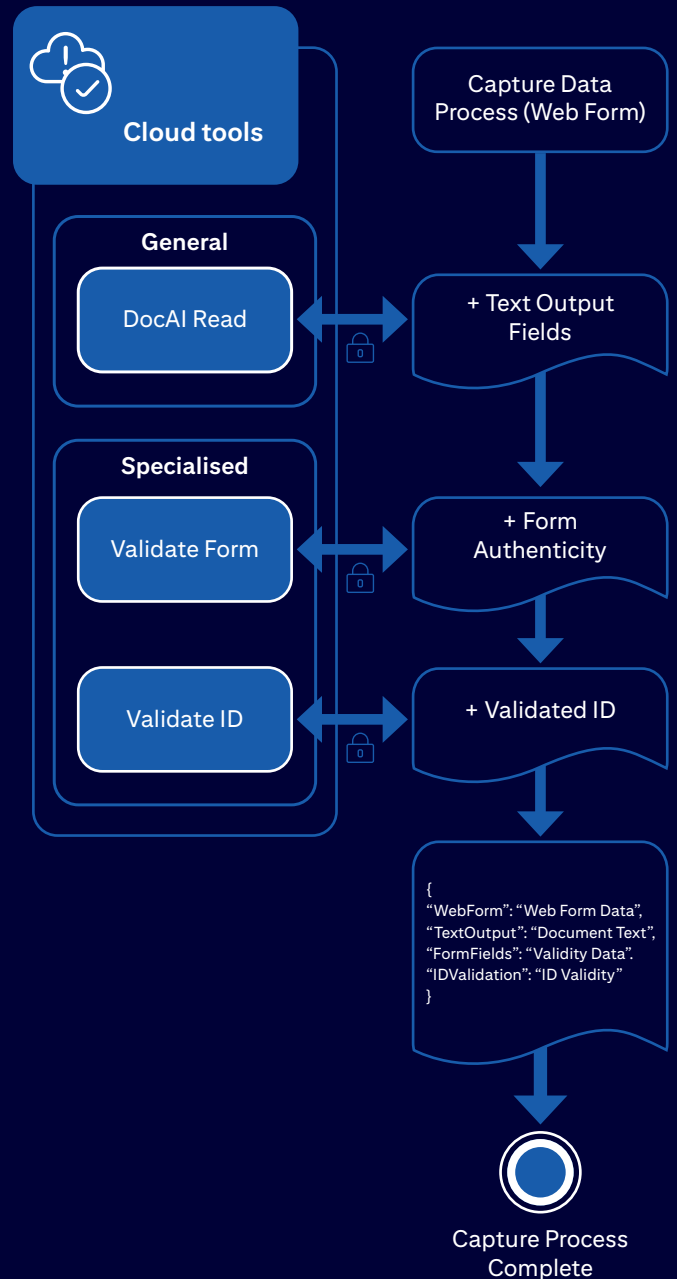
02 | Document Metadata and Context Extraction

Today, there are multiple advancements in document digitization and understanding. Many rely on a mix of optical character recognition (OCR) scanning, use of AI vision models (which process and interpret visual data) and accurate application of context to recognise a document with more accuracy. These improvements unlock the possibility of creating a graph knowledge database, which stores, queries, and manages data structured as a graph, with a digitized repository of all client documentation, indexed with contextualized embedding. These components could feed into an AI model to facilitate a meaningful retrieval augmentation generation (RAG) system to further automate instructions and documents by:

- Removing the need for the same documents to be provided multiple times – all relevant documents would be easily searchable because context would be extracted and indexed in the knowledge database.
- Accurately reading the text from a web form or scanned documents by considering context – this will further enhance the accuracy of document recognition.
- Understanding information in a form in the context of existing client information and requirements to validate intended instructions.
- Validating against official documents (ID, statement, certificate) to ensure consistency of name and address information to reduce the incidence of fraud.

We are exploring how such a system might benefit clients through improved onboarding and instruction processing. It might also allow clients to retrieve only the documentation relevant to their needs. In future, a context-enabled, fully authenticated client console could receive queries in various multimodal forms such as voice, text, or picture, anticipate the client need and suggest set up instructions.

We are also exploring the potential for efficiency and speed improvements in fraud prevention by building extensive graph knowledge bases allowing for contextualised analysis of payments networks and all interconnected nodes as part of customer onboarding (know your customer (KYC) and customer due diligence (CDD) processes), loan origination and compliance checks. These processes, which are more complicated and time-consuming in corporate than retail banking, often involve multiple touchpoints, paperwork, and approval and review cycles.



03 | Aligning Citi's Offering with Client Needs

Modern businesses grapple with the complexities of integrating financial operations. They often encounter issues such as convoluted banking interfaces, problems when setting up payment message formats, and difficulties in getting swift, consistent, and personalized responses to queries during periods of high transaction volumes.

Clients often need guidance to match their needs against Citi's product offering and find it difficult to ascertain which solutions will offer the greatest value and an improved banking experience. These challenges can lead to operational inefficiencies, increased expenses, and lower growth.

The use of generative AI presents an innovative solution. GenAI's strength is in rapidly analysing vast amounts of non-structured data, finding patterns, and making connections. At the same time, its natural language interface can deliver correct, consistent, and contextualised answers much faster, increasing customer satisfaction.

Tackling these problems with the large language models (LLMs) used in GenAI involves analysing and classifying customer queries based on their nature and urgency. When it comes to queries, two types in particular may see a noticeable improvement:

- (1) "How do I" queries relating to the use of Citi functionalities and products, and;
- (2) "What is the status of my transaction" queries.

Internal chatbots used by trained customer support personnel could improve the efficiency of customer support functions. Currently, agents rely on multiple tools and fragmented information sources, which often have imprecise search capabilities.

GenAI would enable them to leverage an integrated and contextualized knowledge base (outlined earlier in this paper). This removes the burden of needing to analyse pages of search results while helping a customer. An additional benefit is that agents would find it easier to update the knowledge base by feeding the latest information back to it, leveraging generative AI's ability to add content where fits best given existing information. Sentiment analysis of such interactions offers an additional strategic opportunity to differentiate through service quality.

This technology has exciting potential to make interactions productive and enjoyable, delivering multiple benefits. However, consideration must be given to Citi and client data security and the potential for GenAI hallucinations (explained later in this paper).

04 | Straight Through Processing (STP)

The elimination of manual touchpoints (MTPs) from banking processes is a key goal to deliver a seamless client experience and achieve near 100% straight through processing. Improvement of such processes involves leveraging advanced analytics and machine learning techniques, such as mapping client journeys. Also, a deep analysis of payment repairs to address sources of manual touchpoints is likely to be made significantly more efficient when using generative AI. Visualising the end-to-end client experience and overlaying this information with customer suggestions and preferences will add a qualitative dimension to the identification of MTPs.

For example, AI tools leveraging machine learning algorithms can reduce errors by prompting for details on manual entry forms, suggesting payment repairs and identifying and automating error correction in bulk payment submissions. They can also enrich payments data, help with reconciliations against invoices and purchase orders, map data from unstructured to structured formats like ISO 20022, and ensure that local remittance data for a specific geography is provided in the required format.

There are various efforts underway to remove friction by automating payment routing and decision making. Our corporate customers operating globally and engaged in international trade or with large volumes of transactions often face high costs associated with cross-border payments. This is partly due to the need to navigate a complex landscape and make manual decisions relating to payment products, fees, and cut-off times across geographies. We are exploring AI tools to improve routing of payments and decision making, consistent with the criteria set out by our clients, to reduce costs.

Insight-Driven Client Value

In an economy where data has become a cornerstone and a driver for innovation, economic growth, and competition in the market, it is imperative for financial institutions to enable customers to derive value from their data. Treasury managers at medium and large multinational corporations and governmental institutions rely on corporate banking products and services to manage their organisation's cash flow, financial assets, risks, and liabilities.

Large multinational corporates and their banks deploy significant resources to navigate regulatory frameworks, standards, compliance, and reporting requirements. This is both time consuming and costly. These activities depend on internal and external data and could potentially be vastly enhanced with the help of AI tools. Large language models connected to a reliable knowledge source, structured either as a RAG system or knowledge graph using graph databases, could deliver a significant increase in efficiency.

05 | Querying Own Payments Data

Clients might be able to interact with their data through a generative AI interface without needing data manipulation proficiency or being restricted to predefined data analysis features within an application. Instead, they would have the flexibility to ask questions in natural language, such as "Who are my top suppliers and customers, based on payment value and volume, segmented by geography and payment methods?". The responses go beyond mere data; they are intuitive insights in natural language, accompanied by visualisations and the capability to explore 'what if' hypothetical scenarios. The system provides text summaries that dynamically update as the data evolves, offering a seamless and accessible experience.

06 | Open Banking Integration

The effectiveness of AI-based analysis depends heavily on the quality and completeness of the data it uses. In terms of data completeness, it is important to highlight that the introduction of Open Banking and Open Finance in various regions allows multi-banked corporations to gain a comprehensive view of their finances. These regimes enable a company's primary banking partner to access data from other financial institutions, making it easier to incorporate this information into AI-based analysis than in the past.

07 | Compliance Insights

Graph LLM-powered tools can help monitor and stay up to date with local transaction requirements using natural language processing algorithms. These tools identify relevant updates, which could be used by the bank to adapt its processes and policies, and by clients as they incorporate changes into their business. Banks' transactional processing flows would also benefit from continuous data validation against requirements, which would flag non-compliance promptly, allowing both the corporate and the bank to respond in sufficient time to avoid penalties.

Client Focus

08 | Co-Creation with Clients

By sharing analytics and AI insights with clients, Citi can establish a common goal and collaborate with them to co-create new products and solutions. This partnership would empower clients to make informed decisions, take forward-looking actions, and work in tandem with Citi to reach their goals. It transforms the traditional banking relationship into a more dynamic and interactive partnership that benefits both Citi and our clients.

The true voice of corporate clients is often obscured by limited feedback channels, rigid surveys, and generic market research. In the near future, generative AI will be able to unlock novel approaches to express feedback more naturally and move away from static forms. This would allow AI to generate insights that illuminate hidden needs and preferences, fuelling a virtuous cycle where client feedback directly shapes transformative, highly tailored products. For example, a client seeking to transfer funds would receive guidance on how to best complete the action to meet certain criteria. At the end of the process the chatbot could ask, “How could we make this process easier next time?” The client’s natural language feedback would provide far more actionable insights than a 5-star rating.

In the future, generative AI could provide summaries of insights derived from social media posts, enabling a quick and tailored reaction to emerging trends, client ‘chatter’ about various financial products, and unmet needs. For example, if AI identifies a spike in mentions of supply chain finance limitations among exporters, it could signal a need to develop tailored solutions. In addition, identifying emerging trends could help prepare for potential market shifts, ensuring products and services remain ahead of the curve.

Another area of exploration is hyper-personalization within client segments, industries, and even individual company profiles, which can uncover niche needs and enable customization at scale. Feedback from generative AI channels could determine feature prioritisation, product design, and service improvements that are aligned with genuine needs to produce bespoke offerings. Deeper understanding of client pain points would become a foundation on which Citi could further build trust and gain competitive advantage.



09 | Feedback Collection at Scale

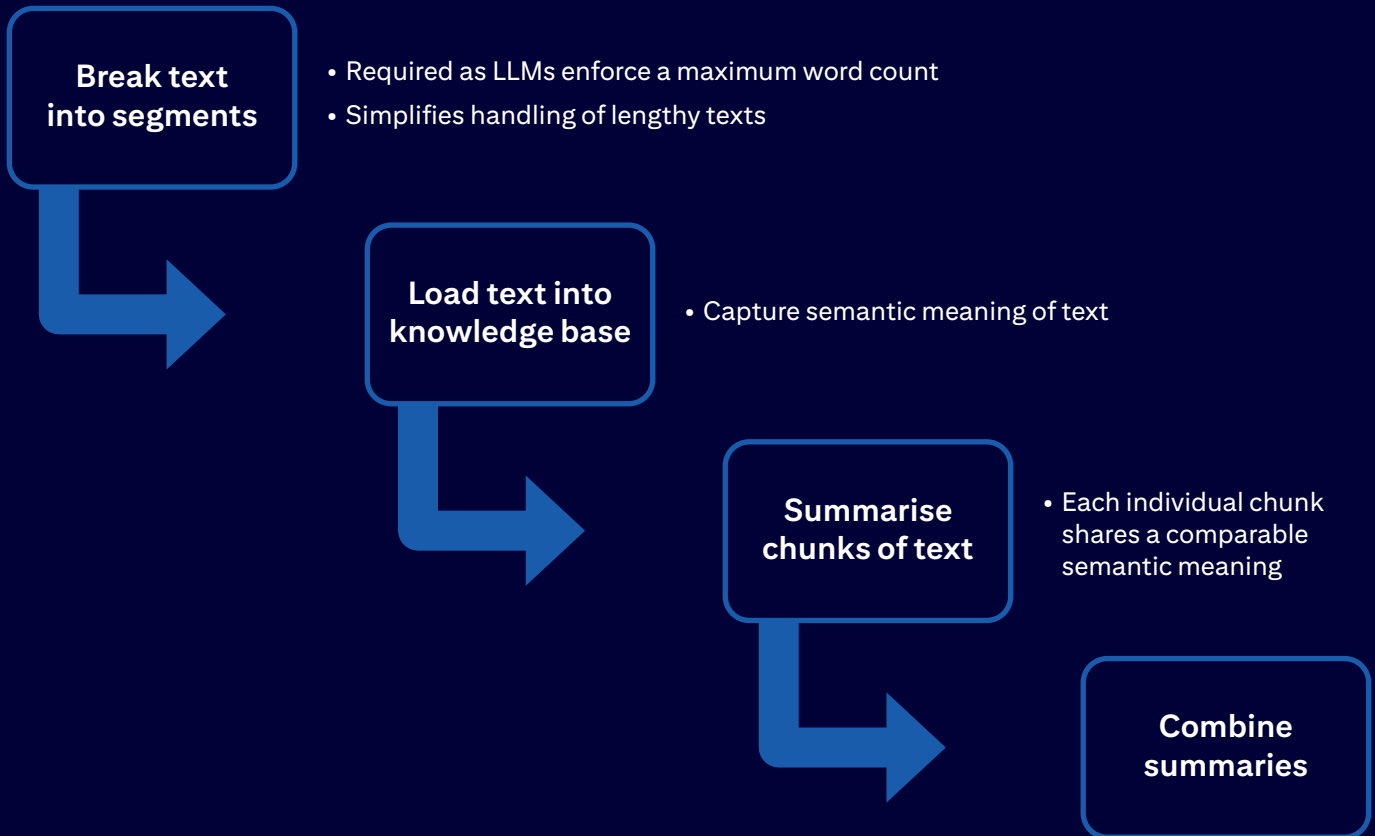
How do we ensure that emerging solutions satisfy the needs of a majority of our clients in a given segment? Collecting information at scale is a challenge for any large organisation. To avoid the high costs associated with manually collating responses for a large customer base, organisations adopt one of the following approaches:

- Collect samples from a subset of the customer base or;
- Rely on simplified multiple-choice questionnaires that can be quantified.

The problem with these approaches is that key information is likely to be omitted.

Since generative AI can summarise text, a powerful use case would be to allow all, rather than a few, customers to provide feedback in natural language. The combined responses could then be summarised, providing insights from responses provided.

The algorithm to summarise the text performs the following steps:



After collecting client feedback, an automated process can then summarise the responses. In addition, tools could be developed to allow analysts to ask questions relating to the feedback using natural language instead of technical query language or tools. For example, the analyst could simply ask, “What are the five main pain points for customers?”. Reactions to new products could be assessed in the context of the knowledge base to determine how fully the original client needs were met. The response could be provided either as natural text, or if the analyst prefers, as visual graphs.

Challenges and Ethical Considerations

Large language models are the technology that underpins generative AI, enabling communication using natural language. To achieve this, the models need to be pre-trained using vast amounts of existing text. One challenge is that most of the companies that develop the models do not disclose which data was used to train models, making it impossible for users to verify the training data. While vendors claim data is responsibly sourced, this lack of transparency forces users of generative AI to blindly trust their claims.

After pre-training, models undergo supervised fine tuning, during which they are trained – using a large set of commands and responses – to help them avoid unethical responses. Human operators participate in this process, providing feedback on the model's outputs, so the quality of the model's behaviour depends on how these operators rate its outputs.

One major concern with AI is that if training data contains bias, the model will learn and reproduce that bias. For generative AI, bias can come from the training data, or from the humans providing feedback. This bias can lead to content being created that reinforces cultural, racist or gender-based stereotypes. If the model has inherited biases during its training and tuning, its outputs could discriminate against certain groups, or lack diversity and inclusion. This presents a risk as an organisation using generative AI may generate output at odds with its own cultural beliefs and policies.

While LLMs generate text that is fluent and coherent, they can sometimes produce text that is inaccurate or illogical, a phenomenon called hallucinations. These vary from minor inconsistencies to total fabrications. To prevent hallucinations, it is necessary to understand their causes, the most important of which are:

- **Data accuracy:** If a model is trained on data that contains factual errors or bias, the model may repeat the inaccuracies. Data usually includes pretraining data (from its supplier) and domain-specific knowledge provided by an organisation.
- **Data completeness:** If the dataset lacks the necessary information to answer a question, the model might generate incorrect responses.
- **Input context:** If a query is unclear or lacks necessary contextual information, the response is more likely to be unreliable.

To minimise hallucinations, data has to be accurate and complete. Pretraining data needs to be evaluated for accuracy and bias in the specific application domain. Domain-specific data used by an organisation to train a model must be accurate, complete, free from bias and not infringe intellectual property rights.

Handling of personally identifiable data by the LLM requires strong safeguards to prevent leakage, and ensure privacy and security. Extreme caution must be exercised and users should be educated before it is deployed.

The EU AI Act specifies that in most cases, customers should be informed if they are interacting with AI systems. Even for organisations not directly in scope of the EU AI Act, it is good practice to follow its provisions.

Generative AI's ability to generate high-quality output, whether it is text, video or audio, creates the risk of deepfakes, and misinformation. It complicates the use of biometrics, such as voice recognition for a telephone channel, in fraud detection.

Conclusion

Generative AI marks a seismic shift in how banks can engage with corporate clients. By leveraging this technology in feedback mechanisms, financial institutions can move from a reactive to a predictive approach, and from generalisation to customisation. This transformation is key to building enduring client relationships based on co-creation and delivering exceptional value in the competitive corporate banking landscape.

This technology needs to be implemented responsibly. Generative AI training data must be complete, accurate and free from bias. Systems have to be thoroughly evaluated to ensure they do not generate inaccurate or unethical responses. The ability for systems to communicate using natural language necessitates more extensive testing than past systems might have required.

As AI systems become more integrated into banking operations, Citi will proactively manage AI risk to ensure transparency, fairness, and privacy. To maintain trust with clients and regulators, Citi continues to evolve its AI risk management framework enabling safe adoption through an appropriate set of guardrails to ensure alignment with our high ethics and governance standards.

Citi invites clients to join us on this transformative journey. Together, we can enhance our competitive edge and create more efficient and personalised banking services. By embracing AI, we can deliver innovative solutions and shape the future of the banking industry. While there will be challenges on this journey, the potential rewards make it a worthwhile endeavour.

The Use Cases described are not necessarily ones that Citi will adopt, or that Citi will adopt in the near term.

Authors



Jiten Mehta
Citi



Santiago Dinelli
Citi



Damian Kowalski
Citi



Tamsin Crossland
Principal AI
Architect,
Icon Solutions



Eligijus Kubilinskas
Principal Payments
Architect,
Icon Solutions



Dean Ogden
Principal Solutions
Architect,
Icon Solutions



Services
citigroup.com/global/businesses/services

© 2024 Citigroup Inc. Citi, Citi and Arc Design and other marks used herein are service marks of Citigroup Inc. or its affiliates, used and registered throughout the world.

2306177 11/24