

Banking on AI

– The Next African Leapfrog Moment

Successful Adoption of AI in Kenyan Banks



Contents

1	The transformative potential of AI	03
2	A brief history of AI	04
3	Is AI intelligent?	05
4	Balancing opportunity and risk in AI	05
5	The AI opportunity and the risk of inaction	07
6	Steps to becoming AI ready	07
7	Regulation, Compliance and Legislation	10

Key Takeaways

- Like the Internet, AI is a transformational change for society
- AI is a tool like any other and the bulk of real-world applications are important but mundane – replacing tedious, repetitive tasks and streamlining operations
- The popular press focus on “Super Intelligence” and its risks or opportunities are far removed from the practical realities, at least in the medium term
- Like any new technology, banks need to be mindful of the potential issues, but not allow this to hinder the development of profitable and beneficial solutions
- The risk of non-adoption are arguably much greater than those of mindful adoption

The transformative potential of AI



Generative AI (the technology that underpins tools such as ChatGPT) has the potential to change society. For the first time, these “Large Language Models” are able to undertake tasks that traditional software has struggled to conquer but humans find easy – tasks like natural language, nuance, content generation and handling incomplete or conflicting data.

In the banking context, this means that AI can underpin call centres, summarise complex documents, automatically “understand” paper documents, cope with imprecise and unstructured data and create personalised content that engages customers as individuals.

It can be thought of as having an army of low cost assistants that are able to undertake a wide range of often mundane and repetitive tasks, freeing staff to focus on more important activities such as customer engagement. This transition represents a profound shift in the way of doing business that will transform the industry in the coming years.

A Brief History of AI

The subject of Artificial Intelligence (the ability for machines to learn and to act in an apparently similar way to humans) is not new. In fact, it has been studied since the 1950s. In that time, AI has gone through three distinct phases.

Phase 1: Symbolic AI. In this phase, it was thought that human knowledge could be captured as a series of “decision trees” and led rise to “expert systems”. These were marginally successful but it proved difficult to get the systems to work outside well-defined environments.

Phase 2: Artificial Neural Networks. This technology had its growth in the 1980s and was based on a totally different approach, drawing inspiration from the structure of the brain but implemented in software. It showed promising results in small systems for tasks such as image understanding and finding patterns in data (for example, predicting which customers will repay their loans from analysis of large quantities of historical data).

But implementing such solutions is difficult. It requires highly skilled engineers and a quantity of high quality, curated data. As a result, such systems found success in specialised applications such as self-driving cars, recommendation engines and loan quality management.

Phase 3: Generative AI and Large Language Models. Although originally developed by Google (“Attention is all you need”), this technology came to prominence in November 2022 with the launch of ChatGPT.

The key difference to previous AI was that the AI came ready-trained (on vast quantities of written text) and so could answer questions without any specialist knowledge. Anybody could ask ChatGPT for a recipe for Chicken Korma or to solve a maths homework problem.

This simplicity of use makes AI widely accessible to everyone. It means that AI projects that would otherwise have taken years can be achieved in weeks. And it’s being used in a wide variety of applications from customer chatbots to instant meeting summaries and computer programming.

With ChatGPT, the topic of generative AI took off. Companies and governments spent hundreds of billions of dollars (for context, equivalent to several times Kenya’s annual GDP) on research, technology and data centres. Companies extrapolated the leap to generative AI into the future and predicted “Artificial Super Intelligence” or “Artificial General Intelligence” as being just around the corner. Such development has some profound implications – will AI eliminate work except for a few elites? Will power be concentrated in the hands of just a

couple of companies? What happens if AI becomes self-aware and does not want humans any more?

At the time of writing (end 2025), the estimation or the rate of progress of AI has been tempered, with most now predicting that AI will develop more slowly than these initial expected. But that does not mean that AI is not already very useful. In fact, AI development could stop today and still it would take 5-10 years for industry to catch up to what it can do.

Is AI intelligent?

Generative AI in its current form can best be thought of as an “knowledgeable parrot”. It has been trained on a vast array of historical language and for most everyday problems, someone, somewhere has solved it in the past. So, generative AI is generally thought to mimic intelligence rather than necessarily being intelligent in its own right. Specifically, generative AI cannot distinguish between commands and data and is not thought to be self-aware (though it may be able to mimic this also). Imagine you, with only basic medical training, are at the end of a medical helpline. With access to a good enough search engine, you can probably provide a passable example of a doctor in most circumstances. The issue, of course, is you have no idea if what you are saying makes sense. And in the cases where the search engine gets it wrong, you will just authoritatively repeat the error.

Generative AI is similar. It can “hallucinate”, filling in gaps where it does not have specific information to respond to a query. As long as the user is an expert, the AI is still useful but they need to be able to spot where it goes wrong. This becomes more apparent when the AI is asked to reason. Imagine an AI is 90% right. It does not need many steps before it will go off track. So, a simplistic approach of “just get the AI to do it” is rarely successful. But building AI into a structured solution to a problem can be highly effective and deliver products that would be tremendously difficult to do in software alone.

Balancing opportunity and risk in AI

Like any technology, AI has benefits and some risks. The size of each depends on how the AI is used. One might ask “is a car a benefit or a risk?”. It is of course both and the magnitude of each depends how you use it. But with sensible guidelines, the value far outweighs the risk. AI in banking is similar.

AI risks fall into two categories: the risk of using AI yourself and the risk of others using AI against you. Many of the well-publicised risks of AI fall into the latter category. These include

the potential for “deep fakes”, cyber-bots, surveillance and infringement of civil liberties, harms due to mis-diagnosis for example in a medical context etc. While these risks are absolutely real, they are external to most applications in banking. In other words, they will happen whether or not a bank implements AI solutions and so are not the focus of this document. These risks need to be planned for whether or not the bank chooses to adopt AI.

The risks associated with AI applications a bank might implement include issues such as hallucination (providing incorrect information), data privacy and potential leakage of privately identifiable information and solutions that impinge on civil liberties.

Fortunately, there already exists a good regulatory framework around such issues. The GDPR act is clear about data privacy and this applies equally to AI and non-AI solutions. Issues relating to civil liberties are covered well in the EU AI Act, which is being closely mirrored by the regulatory framework in Kenya.

Therefore, the risk of AI (like the risk of a car) is in the way you use it. There is a wide range of valuable applications for AI that have minimal risk:

- A customer care chatbot that provides generic product information
- A tool to accelerate customer loan provision by automatically checking supporting documents (but leaving the ultimate decision to a human loan officer)
- A back-end solution that reduces the number of payments that cannot be reconciled to a given account
- A tool to personalise customer marketing
- A WhatsApp agent or voice agent that assists customers in opening a bank account



Like any software, the solutions need to be implemented with due care and consideration for topics such as cybersecurity, data privacy and correctness. But with straightforward (yet tremendously useful) services such as these, the incremental risk from the AI is modest indeed and straightforwardly managed.

The AI opportunity and the risk of inaction

According to research by Akili AI in 2025, over 80% of senior executives in Kenyan banks believed that AI would be transformational to their business. However, only around 20% felt “prepared” or “very prepared” for the transition.

Digital banks have very different economics when compared to traditional “bricks and mortar” organisations. The following table is representative for banks in Kenya:

Metric	Traditional Kenyan Bank	Digital Fintech Bank	Ratio
Cost-to-Income Ratio	50–70% (typical in SSA)	~30–50%	
Operating Expense per Customer	KSH 13,000 – 39,000/year	< KSH 4,000 /year	3-10x
Customer Acquisition Cost	KSH 6,500 – 20,000	< KSH 1,300	5-15x

The disparity between operating costs of digital and non-digital banks means there is a substantial first-mover advantage for organisations that adopt AI and, for those that choose not to embrace the technology, the risk of being severely commercially disadvantaged in the future.

There are specific areas where AI is already having an impact in the banking sector. These include:


Account opening and customer onboarding – replacing manual document checks, written forms and multiple stages of decision-making

incremental risk from the AI is modest indeed and straightforwardly managed.

Steps to becoming AI ready

The future use of AI is likely to be so widespread that it’s sometimes difficult to know where to start. AI will become broadly embedded in commercial products and a wide variety of AI platforms will emerge. But none of these are useful until they are integrated into the bank’s systems and data.

The advantage of generative AI over earlier technologies is that it can be applied to a wide range of tasks in a period of weeks rather than years. Therefore banks can take a step-wise approach to AI adoption, selecting individual profitable solutions one at a time and evolving towards a business model where AI plays in integral part in almost all business operations.




While it would be possible to go all-in on artificial intelligence (and some challenger banks have), for most established banks the adoption of AI is a more stepwise approach. This is in part because many banks with legacy operations have data stored in multiple siloes with no unified means of access. While this needs ultimately to be addressed, it limits the extent to which AI can easily be rolled out across the organisation as a single layer. In addition, AI technology is moving at such a pace, it is prudent for banks not to become too aligned to one specific technology that may be superseded in short order.



The journey toward adopting generative AI should be strategic and grounded in practical steps that align with resources, the competitive landscape and customer needs. The first step is to identify specific business problems or processes where generative AI can add value. Banks should focus on high-impact, low-risk areas such as customer service (e.g., AI chatbots), content generation (e.g., automated customer emails or reports), marketing personalization, document summarization, and compliance support. Selecting well-defined, narrow use cases helps ensure early wins that build internal confidence and stakeholder support.

Generative AI is a new field that requires both technical and domain knowledge to adopt in the most cost-effective way. Specifically, it is key for the bank to create the appropriate foundations of reporting and governance to act as a basis for future development work. Banks can partner with technology providers, fintechs, or AI consultants to access needed expertise without having to hire large in-house teams. At the same time, banks should consider upskilling a small internal team through online courses and vendor training programs to ensure that the management of technology providers and consultants is undertaken in the most effective manner. A designated AI lead or project manager can coordinate these efforts.



The process can be broken into three phases: The preparation (Engage) phase sets up the organisation with the core capabilities for success without substantial investment. The next phase (Evaluate) involves creating a small number of rapid proof-of-concept applications to demonstrate the value of the technology and to align the organisation to the changes in process, management and governance that will be needed with the new solutions. Only once the foundations are in place and the business alignment has been demonstrated does it make sense to invest in larger-scale rollout of technology.

While this process is comprehensive, it need not be time-consuming or expensive. With proper planning, most organisations will be able to undertake the preparatory Engage and Envision phases within an elapsed period of 3 months.

Regulation, Compliance and Legislation

As generative AI becomes increasingly embedded in the financial services landscape, banks are exploring how to harness its potential responsibly. However, due to the regulated nature of banking, adoption must be carefully aligned with regulatory requirements, compliance frameworks, legislative obligations, and robust internal operations. Unlike larger multinational banks with deep legal and risk teams, many African banks must be especially strategic in how they address these issues in order to balance resource and impact.

“AI compliance is like tax returns. It is not especially difficult but it requires dedicated attention and is not something to be ignored”

Before deploying generative AI tools at scale, banks should take into account the relevant legal and regulatory frameworks. This includes data protection laws (such as GDPR, Kenya's Data Protection Act and similar laws), financial sector regulations (e.g., central bank guidelines), consumer protection standards, and sector-specific AI guidance. Banks should also take into consideration the Kenya AI Strategy 2025 and ensure that operations are aligned with the overall direction of their recommendations.

However, for the bulk of early applications of Generative AI, the additional considerations regarding compliance are modest. While the regulatory environment for AI is still in its infancy, there are a number of well-defined AI policy frameworks including NIST (USA), the EU AI Act and Singapore's Model AI Governance Framework. All the frameworks have similar underlying principles and approaches. The majority of use-cases for banks fall into the “low” or “minimal” risk categories, meaning that the processes and procedures currently used for mainstream software development can be reused, almost without change.



There are cases where additional measures need to be taken into account. These include where personally identifiable information is shared outside the organisation and in cases where the AI is making decisions autonomously, for example for loans. But both of these are generally handled without difficulty in the same way as current banking systems work, by encrypting or anonymising data and by ensuring there is always a human in the loop that is making the ultimate decision.

It is also important that employees understand how to work with AI responsibly, particularly in relation to data privacy responsibilities and ethical considerations. “Shadow AI”, where employees use tools such as ChatGPT to support their work, is widespread in the industry and is a potential source of data leakage or incorrect information. Employees should be made aware of the risks of using AI in this way, particularly if it involves personally identifiable information and banks should have appropriate policies in place to cover what, in most cases, is already happening.

Adopting generative AI can bring powerful capabilities to banks, enabling them to operate more efficiently and serve customers more effectively. By choosing initial applications of AI that are inherently low risk but still highly valuable, and taking appropriate proactive steps across legislation, reporting, and operations, banks can confidently harness AI’s benefits while minimizing risk and maintaining regulatory trust.

