

Banking on Artificial Intelligence and Machine Learning (AI/ML)¹

Prompted an AI app to give a quote on artificial intelligence, it returned “Machine intelligence is the last invention that humanity will ever need to make” – Nick Bostrom in 2015. That seemed to be one of the most fraught quotes one can encounter. Most Indian banks, like their global peers, have been using traditional AI for decades, if not on the bleeding edge. However, the expression ‘AI’ seems to have suddenly come of age when it was named as the most notable word of the year 2023 by dictionary publishers Collins, last month and CAFRAL conducts its first programme on AI today – that is a long maturing process since the word was first coined in 1956 by John McCarthy. The term ‘machine learning’ was coined by Arthur Samuel in 1959 driving home the point that computer could be programmed to outplay the programmer. Now it is counted in the league of transformative technological innovations turned the world sideways, such as personal computers, Internet, and mobile phones. If somebody were to say that AI/ML is rocket science, it indeed could literally be so, if one followed the role of AI/ML in landing of Chandryan-3 on the south pole of moon. At the same time, one can go easy knowing that each one present in this virtual room have also been personally using AI/ML for a long time. The objective of today’s programme is more around strategic disruption of AI/ML in banking rather than technical insides of AI/ML. One would also add another objective to this programme i.e to ring a bell for banks who have not yet put in place an AI strategy for their business. AI strategy cuts both ways – it can help realize the business strategy you already have in a more efficient way, or it can even help in developing a business strategy you do not have yet. New technologies and particularly AI always have high scope for adding new lines of banking business. My remarks will be an attempt to draw a rough outline of a few dimensions of AI relevant at this point in time, with banks in the foreground.

Understanding AI/ML

2. AI is not a monolith but a selection of different technologies and sub-types. There is no regulatory definition of AI and even for regulatory purposes, principle-based or risk-based approaches to the definition of AI/ML with a focus on its specific characteristics

¹ Keynote address by Jayant Kumar Dash, Executive Director, Reserve Bank of India, at CAFRAL’s Virtual Program on Artificial Intelligence and Machine Learning on November 10, 2023

and risks posed or amplified by it have been advocated to be more appropriate. Distinction between AI-ML and conventional software lies in how a task is performed by a computer. A conventional software is 'programmed' to perform a defined task; AI is 'programmed' to learn to perform the task. Einstein had similar definition for 'Education' i.e 'the training of mind to think' rather than just learning of facts. Unlike in conventional software, despite being the building block, code is not the dominant artefact in AI which has data at its 'back and centre'. That explains why it took a while for AI to be mainstreamed as it needed oceans of data to gather before setting sails.

3. If one were to select a single word that AI deals with, it will most probably be 'prediction' - the process of filling in missing information. The process takes the information (i.e data) that is available and generates information that isn't. The prediction needs to be actionable. Decision-taking for most financial activities by a bank can be reframed to a 'prediction' task and hence, be assigned it to AI for solving. To give a better context, the 'prediction' need not always connote future (e.g a loan repayment probability) but also present (e.g a fraudulent credit card transaction, real time). When we talk of AI/ML today, it is the Generative AI that has created the excitement which has 4 essential ingredients, namely algorithm, compute, data and business process. Simplistically speaking, while the first two belong to the realm of technology, the last two is where business i.e banks must contribute.

Generative AI (GenAI)

4. Generative AI (GenAI) is a type of AI technology that can produce various types of content, including text, imagery, audio and synthetic data (useful in fraud detection or privacy maintenance) in a humanlike manner. The recent popularity around GenAI has been driven by the simplicity of user interfaces for creating high-quality contents instantly. The technology, evolving over decades, got a shot in the arm with development of an algorithm called generative adversarial networks (GANs) in 2014 (it is same ML used in generating deepfakes). The other advances that played a booster role are transformers and large language models (LLMs). Transformers enable training of large models without pre-labeling of data as required earlier and have capability to track connections across large a spectrum of data population. LLMs are models with humongous number of parameters enabling GenAI to create contents in multiple formats. That capability makes it a suitable candidate to be used, say, to

summarize content, answer questions in a chat format, and edit or draft new content in different formats. As a corollary, GenAI in banking could rapidly and cheaply generate hyper-personalized products and services, or automate software engineering, IT migration, and upgradation of programs. There are talks about the hyper-personalization becoming the next big thing in banking.

5. The potential of GenAI's to cater to banking sector, more than any other sector, comes from its ability to understand natural language, among others. The impact of GenAI on different segments of bank operations may vary but broadly the benefits could be overall value accretive in terms of cost reductions and efficiency gains. AI-facilitated automation and prediction are not new to banks' digital transformation. However, tools driven by newer, GenAI powered systems in banking are rather in its infancy; premature to enumerate its possibilities or the risks. This new wave of AI at the same time offers opportunities for the large and growing network of FinTechs, nimbler than banks. The resulting capabilities, unless absorbed by banks in time despite certain symbiotic relationship already built, could create some disruptions to the banking sector. Industry experts opine that testing of GenAI solutions will accelerate over the next two to five years and the benefits are likely to accrue incrementally rather than immediately. Sustained testing of GenAI and ongoing investments during this period will be required before scaling up, customer facing deployment and engaging in more transformative projects. Focus on incremental innovation resulting in efficiency gains or product / service improvements as per identified business needs are likely to be short-term preferences of banks. This stage may also remain human supervised, known as human-in-the-loop (HITL), to ensure expected results in terms of acceptable precision, accuracy, and compliance. Banks are likely to apply foundation model (gigantic systems trained on enormous corpuses of data) of GenAI to other running or legacy applications with a view to improving their efficiency, in addition to newer use cases. Example: using the digital data trail arising from digitalization and automation of customer-facing processes that can be processed by GenAI to tailor both the service and its internal processes. Digital data trails could also be used to improve risk management, data collection, reporting, and monitoring, in turn delivering further digitalization, leading to hyper-scale customization

Navigating the Hype

6. “Banking is necessary, but banks are not” – words of a younger Bill Gates nearly three decades ago has been quite a buzzer for bankers during last few years, without being anywhere near cracking. One of the consultancy firms a few years ago propagated its vision of an ‘invisible bank’ where “enlightened virtual assistants” replace people at all points of customer interaction. Now, it is early stretch of AI/ML where investment into it is done out of either faith or fear (of missing out). Gen AI is posited on the ‘Peak of Inflated Expectations’ on the Gartner Hype Cycle for Emerging Technologies, 2023, projected to reach transformational benefit within two to five years. Presently, AI companies have found it hard to deploy the technology to production, as the number of commands to get a near-precise outcome is still very huge. As per the Hype Cycle graph, there has to be dip in the hype cycle (i.e ‘trough of disillusionment’) followed by a ‘slope of enlightenment’ when it reaches the ‘plateau of productivity’. That would indicate the dynamic strategy horizon for banks. The Initial fervor for a hyped technology slowly gives way to a rigorous analysis of risks and implementation challenges.

7. Not all AI is created equal, or for that matter, AI at all and the word more often tends to be used as a marketing buzzword. If one excludes the pure-hype AI products, many others simply apply existing AI to a given category of banking need, often licensing big AI systems. In order to qualify as AI, a system must demonstrate certain level of learning and adapting. Many products which promise to be ML or AI may actually be well-dressed, sophisticated or impressive algorithms. Technology that operates without human interaction is not necessarily AI and most of the time are sophisticated algorithm. IoT and cloud computing are extensively used in AI, but it’s not same as AI. More and more AI products are pitched for different propositions of value to make autonomous decisions at speed of a machine for financial services. A close look at the full AI system including the infrastructure, learning process, data, and technology and resources used to create by cutting through the surface specifications or marketing hype may be helpful before committing big money. In the absence of ML, all situations or scenarios anticipated at the time development of the app can show its smartness, but it cannot discover new facts outside the programmed universe. Banks need to be observant to the fact that smart is not same as intelligent.

Crafting AI Business Strategy for Banks

8. That AI powered strategies will shape the future of banking, often hits the pitch of a slogan of late. For banks, an AI strategy assume criticality owing to the momentum of obsolescence of traditional applications as all applications tend to have some form of AI. Adoption of an effective AI strategy and governance are expected to use to gain (i) an edge on franchise value that boosts customer acquisition and retention; (ii) improved margins through efficiency gains and profit maximization; and (iii) enhanced risk management outcomes for credit, market, and operational risks. A Gartner's web poll in March-April 2023 may show the signs of time for AI. About 45% reported that the publicity of ChatGPT has prompted increased in AI investments. Further, 70% of executives said that their organization is in investigation and exploration mode with generative AI, while 19% are in pilot or production mode. Primary focus of such investments was in the areas of customer experience / retention, revenue growth, cost optimization, business continuity. The poll found that 68% of executives believe that the benefits of generative AI outweigh the risks, compared with just 5% that feel the risks outweigh the benefits. However, executives may begin to shift their perspective as investments deepen.

9. A condition precedent will be apt to preface any AI strategy talk. An AI strategy for a bank cannot be delinked from a robust data strategy where the data architecture is granularly aligned to business goals and data governance, with high priority on data quality. The AI technologies deployment has to be 'at scale' and require a holistic transformation spanning multiple layers of the banking and IT organization. Further, the AI strategy for small and medium sized banks would be different from that of a large or very large bank. Now AI model's size alone may no longer define its utility, or even its capability for that matter. Rather, integrability of banking apps with these models or compressibility of LLMs to run on mobile devices, for example, are more critical for increasing their efficacy. The availability of GenAI, drawing on all data stores of an enterprise in productivity applications as well as other data stores, seeks to increase its progressive adoption. For an effective strategy formulation, clear understanding of what price has changed for AI now and how that price change will cascade throughout the broader economy will help appreciate the investment value.

10. Every company spends significant amount of money on technology and banks do more. Somebody was heard saying that when you add up what banks and other financial services companies spend on technology, it could be more than what technology companies do themselves. While AI investments require upfront lay out, they are expected to allow banks to become more efficient over time and, at the very least, could lead to a swifter and better experience for customers. Multifarious surveys which rank the uses of AI/ML in banks broadly list them in the following order: Fraud Detection, Optimizing IT operations, Digital marketing, Risk Assessment, Personalization of customer experience, credit scoring, optimizing product designing, sales and marketing optimization, personalizing investments / wealth management, and portfolio optimization. If properly deployed in the next five to seven years, the ballpark estimates of experts indicate that AI could increase banks' revenues by as much as 30% and potentially reduce their costs by 25% or more. To date, however, banks have typically been scattershot in their strategy to deploy AI, with models ranging from customer-support chatbots to analytics on price-elasticity. But an à la carte approach, without a comprehensive strategy, could prove to be sub-optimal given the number of potential use cases AI keeps developing.

Implementing Strategy / Becoming an AI-First Bank

11. With exabytes of data, banks are essentially into information business and hence, provide a lush ground for AI which lives on data. Supervised learning is a ML process of providing labeled input data as well as correct output data to find a mapping function to map the input variable with the output variable. This advance in ML made banks, with abundant information on inputs, decision and outcomes, a natural fit for switching to an AI-first business model, rather than being stuck in a revolving door of AI picks. It is an oft-heard pitch that banks have little option than to become and AI-First bank. The AI-first business model is one where AI becomes central to bank's core strategy and technology-led operations. For banks, this would assume some imminency on the face of elevated customer expectations, steady uptick in use of advance AI technology by leading financial service providers, progressive disintermediation of financial services through digital ecosystem and creeping entry of technology behemoths in financial service as a contiguity to their core business models, without having to depend on it for revenue.

12. Talking about AI is the easiest part, but the hard work involved in its implementation is very time demanding. Parts like figuring out the right architecture, getting the data ML-ready, having the standards, the processes, multi-cloud environment management needs good investment in resources. Further, transforming the derived insights in a way to impress/ benefit the customer is a further long-drawn process requiring adequate resource budgeting. Unless each step is measured as a part of the whole journey and meticulously signed off, any mistake may potentially put one back at starting line again. Three pillars of successful AI implementation in banking are said to be comprehensive and 'Clear' strategy (emphasis on 'clear'), ongoing Investment in talent and infrastructure, bank-wide culture of innovation and entrepreneurial approach. A recent study by AI Talent Report found that the 60 largest North American and European banks and payment companies employ at least 46,000 people in the "AI and data core", including AI development, data engineering and governance and ethics roles, with as many as 100,000 global banking roles, including implementation, quants and model risk involved in bringing AI to market. Some of the global banks who have successfully implemented AI had invested in a lot of efforts in crafting a charter on AI ethics before starting off, in order to prevent mistakes that could be harmful to the customer upending the whole implementation exercise.

13. The components of AI-first paradigm banks would depend upon their size, business model and where they stand today in the data value chain. However, the following typical aspects in the AI strategy and implementation may be contextual.

(a) Preparing Data Assets: Good data being the staple feedstocks of AI, data transformation programme focusing on data quality is a key step in implementing. For banks with data scattered across systems, a method to build up the data assets enabling instant access to internal and external sources, both in terms of history and real-time data, becomes a necessity. The tradeoff between proactive collation of different data sets and an on-demand search for data when the business seeks to deploy a specific AI application needs careful balancing. The principle of capturing and stockpiling data that is relevant to AI models to be built or deployed has to be tempered by the data-minimization principles of Digital Personal Data Protection Act, 2023.

(b) Deployment 'At Scale': Key market standard metrics / KPIs to track performance of an app typically include lead time to change, deployment frequency, mean-time to

recovery and change failure rate. Rated on a scale of 1-4, the highest rated deployment KPIs (say, on-demand deployment with less than one hour lead time to change / time to recover with change failure rate below 15 per cent) is considered as deployment 'at scale'. In a more translated sense, it is about facing the possibility of thousands of models running at any point in time, some serving millions of users or reviewing millions of transactions while many of these models getting simultaneously recalibrated. While deploying a stand-alone AI model may present itself as easy, applying it to a million customers poses a challenge given the unpredictable nature of user contexts. Managing multiple models raises the degree of difficulty multiple notches. Banks are often advised to develop infrastructure with single clear deployment platform across the organization and a method for managing models.

(c) AI organizational culture: AI has to be seen as a timely opportunity bank-wide, without creating any sense of threat to job security. AI has wide application in both support role for existing tasks as well as fully automating some processes. Employees can be encouraged to act as vanguard to identify potent areas for its deployment. Separating of clear process driven, repetitive tasks for AI deployment, banks can showcase pathways for newer and more contemporary business models by freeing the employees' time for higher-level fulfilments.

14. As with any new technology having world-changing potential, there are always races for supremacy which creates inevitable tension. Inherently, there are two sets of counter-balancing goals in AI adoption i.e agility & flexibility traits of technology vs. managing scale, security standards and regulatory compliance, which often speed-break AI adoption. The common most obstacles quoted as challenges for AI implementation are lack of a 'clear' and comprehensive strategy, unsupportive core technology / data assets and a non-coherent operating model. The barriers to adopting / incorporating AI tech, which will be different for small/ medium / large banks, are generally seen to include Privacy and Security Concerns, Access to Quality data, Lack of consumer trust around AI, Regulatory compliance, complexity and uncertainty, limitations of technology infrastructure /lack of computing capacity, lack of internal skills and talents, high cost / perceived risk of Return on Investment, risk-averse corporate mindset, inability to aggregate data from multiple systems / sources, lack of third party suppliers / partners. Given the benefits in terms of cost savings and new

customer acquisitions, the potential of AI is too material to be ignored on the face of threat from agile startups looking to reselling the banks their customers. With a clear AI strategy and governance plan, banks can build relatively simple models that provide a short-term impact which can then pay for the platformisation for more sophisticated models down the line.

Understanding the AI concerns and risks

15. The adoption of Gen AI is accompanied by its share of costs, risks, and concerns, given the pace of the technology evolution, deployment capabilities, and the potential breadth of its use case. The intended large scale deployment AI will have likely impact around “risk, compliance and enterprise” — with enterprise being a catch-all phrase for HR, legal and other business support functions. More conspicuous among the concerns are ethical questions, security privacy and control risks, nebulous AI regulatory environment, workforce risk, investment requirements for integration of new AI with legacy infrastructure and environmental cost. The key risks that are generally associated to AI/ML relate to security and privacy breaches, failure of AI systems, legal liabilities from AI decisions, workforce and labour displacement, customer trust loss, ethical risks and risks of regulatory non-compliance. However, AI we do not know is actually more worrisome than the AI we know.

16. Weaponization of new technologies is not a new phenomenon and AI with large scope for international spill over is no exception. The new AI/ML models, particularly, Gen AI, with electrifying capability and popularity, have raised safety concerns in its management and governance across stakeholders such as public, business, policy makers and even important political leaders. The AI-related ethical concerns generally involve questions like ability to explain generated content or selection biases embedded in data. Another generic issue is environmental concerns due to the high levels of energy consumed by AI models, particularly LLMs.

Trends in Global AI/ML Regulation

17. The chair of SEC, USA put the likelihood of an AI-driven financial crisis within a decade as “nearly unavoidable”, without regulatory intervention. The immediate risk is more of a new financial crash than a robo takeover. Regulators’ challenges for any

emerging technology primarily emanate from unforeseeable nature of business models and data security, privacy and ownership issues. Hardly a day passes by these days without another new idea on AI regulation springing up. However, the road ahead for framing regulation appears to be long at this stage, despite all discussion on how 'something' should be done with little clarity about who should be doing the regulating. A proposed legislation- the Artificial Intelligence (AI) Act, of European Union is likely to be the first regulatory framework for AI in the world. In approval stage, it primarily focuses on strengthening rules around data quality, transparency, human oversight, and accountability. Transparency in the inner working of AI is not merely ideal but an essential to successful AI accountability. It also seeks to address ethical questions and implementation challenges across sectors ranging from healthcare and education to finance and energy. Expecting AI transparency or explainability to be better than that of corporate transparency may sound aspirational but proprietary AI are notoriously unintelligible. Neither regulators nor consumers can be left completely unaware of the risks linked to such technologies making banks plan only for responsible AI. On 30th October last, the US president issued an Executive Order on Safe, Secure, and Trustworthy Artificial Intelligence containing certain sweeping and hard-coded actions ever taken to protect citizens from potential risks of AI systems. UK initiated Bletchley Declaration earlier this month, with participation of 28 countries to protect against the potential for AI to cause 'catastrophic harm', had three major outcomes (i) a multilateral agreement by tech companies to collaborate with government for testing advanced AI models; (ii) an international declaration addressing risks associated with AI and (iii) UN confirming its support for creation of an expert AI panel akin to the Intergovernmental Panel on Climate Change (IPCC). Further, 29 countries have formed Global Partnership on Artificial Intelligence (GPAI) to guide the responsible development and use in line with OECD recommendations on AI. The next summit of GPAI under Indian Chair is scheduled for December 2023.

18. The Bank of England has adopted a wider programme of supervisory work related to AI by setting up an 'AI Public Private Forum' (AIPPF). Some of the convergent views relating to regulatory frameworks include regulators' capability for 'live' guidance keeping pace with rapid changes in AI capabilities, ongoing industry engagement, inter-regulator coordination at both domestic and international level to smoothen the complex and fragmented regulatory requirements including that involving data

regulation. The key focus should be on consumer outcomes with regards to ensuring fairness and other ethical dimension. With increased use of third-party models and data by AI, more nuanced regulatory guidelines were warranted. Banks will do well to track such developments towards regulatory rails for AI to factor them in for possible buffer zones while designing the AI architecture of their respective institutions. At the same time, applicability of existing regulations to the outcomes generated by AI should not be lost sight of.

Future of AI

19. Somebody coined a catchy headline that future of AI is the future of work when AI prediction capability cannot yet predict its own future given the size of dynamics and the variables involved. Despite the great strides that AI has made in linear task-oriented or domain specific realms (called weak or narrow AI), it is miles away from an artificial general intelligence (AGI) capability or a machine that can truly feel, emote or have a moral compass or, so to say, pass the Lovelock Test. Deep learning, the technology behind advanced AI systems, is not yet advance enough to anticipate all situations, particularly, those not amenable to backpropagation. However, certain developments such as advancements in quantum computing technology, algorithmic shifts to reduce the volume of training data requirement for ML etc. augur a spurt in its adoption of AI by many industries. According to International Data Corp, global expenses on artificial intelligence is expected to reach \$166 billion in 2023 which will rise to about \$450 billion by 2027. Banking is one of the largest contributors by industry at about 13% to above projections. As per the 'India AI 2023' report published by Expert Group under MEITY in October 2023, "AI will be the kinetic enabler of India's digital economy and make governance smarter and more data led. AI is expected to add USD 967 Bn to the Indian economy by 2035 and USD 450–500 billion to India's GDP by 2025, accounting for 10 % of the country's USD 5 trillion GDP target."

20. Not all banks will move at the same speed with or at AI. The competitive environment in banking, regulatory developments, banks' investment capacity and willingness, and customer preferences will all play a role in determining the extent to which AI usage proves to be conservative or more transformative. Hence, a multi-speed adoption is likely to be seen. Retail and transaction banking are likely to go into

AI-driven hyper automation with enhanced customer experience as target which will be facilitated by better capability on anomaly detection. There will be need for developing APIs for integrated services to foster co-innovation and build ecosystem for that would be required. While instances of banks' interest in investing in FinTech/ AI companies in India have been seen, the intensity of bank-FinTech collaboration will also be a deciding factor in realizing improved banking solutions using AI.

Conclusion

21. The genie is out of the bottle and as they say, it is everywhere, at the front and centre of it all. There's little doubt that technology is disrupting financial services, but it is in a positive sense of simplification and it's the customers that are prompting it. The accelerated transition to a digital economy in India by extensively digitalizing the national infrastructure while fostering the DNA ecosystem (i.e data, network, and artificial intelligence) is in full motion. Hence, AI is bound to force rethink banking. The way that customers communicate, work, and consume today means it's not possible for banks to ignore the old song: 'The answer, my friend is blowing in the wind'. In such a juncture, a right first step is always AI appreciation programmes covering the concepts, methods, use cases, risks and challenges of AI for board members and senior management – just as thought of by CAFRAL.

My best wishes for the curated sessions ahead today.

I will sign off with a riddle thrown in – what are the odds that the contents of above remarks were sourced from a GenAI application?

Thank you.