A community initiative, anchored by

INTERNATIONAL INSTITUTE OF INFORMATION TECHNOLOGY
HYDERABAD

IIIT Hyderabad

TechForward

DISPATCH

SEPTEMBER EDITION

TECHNOLOGY TRANSFORMING FINANCIAL SERVICES

Technology is reshaping financial services, making them more accessible and efficient. Right from digital banking apps that offer seamless transactions, to blockchain enhancing security and transparency, revolutionising payments and record-keeping. Artificial intelligence is driving personalised financial advice and optimising investment strategies for individuals. Fintech startups are challenging traditional banks, fostering innovation and competition. Contactless payments and digital wallets streamline everyday purchases, enhancing customer convenience. This issue delves into how technology is democratizing finance, empowering consumers and businesses alike.

IIITH's TechForward research seminar series is an academia-industry confluence around emerging technologies. The deep insights, directional talks and industry outlooks from accomplished thought leaders at the seminar are compiled monthly in the Tech Dispatch as a ready reckoner for technology directions.

From the Chair's Desk

Welcome to the latest edition of TechForward seminar series. This series is special because it serves as a bridge, facilitating conversations on the technologies shaping our future. I'm especially excited as we dive into a topic that is not only transforming industries globally but is also deeply reshaping the financial services landscape.

Our topic this month — 'Technology Transforming Financial Services'— could not be more timely. As we all know, the financial sector has always been at the forefront of adopting new technologies, from ATMs to online banking, but what we are seeing now is a paradigm shift. From banking apps to wallets and blockchains to crypto, emerging and constantly evolving technologies have not only changed how consumers interact with their money but has also created new opportunities for firms to be innovative and efficient. And today, we have AI – the impact of which has the potential to be immense on financial services. With the advent of machine learning algorithms, we're witnessing more personalized banking, faster fraud detection, and more efficient operations. But it's GenAI that offers unprecedented opportunities for automation, data analysis, and customer interaction, yet it also poses questions around governance, ethics, and trust that we must carefully navigate.

I joined Prof. Kamal Karlapalem, head of Data Science and Analytics Center at IIIT Hyderabad, to cover both depth and diverse perspectives on this topic. Our talks have been summarised in the Dispatch along with insights from other thought leaders in the field.

ALOK MADHUKAR
Head of Hyderabad Engineering, Goldman Sachs

CONTENTS OF THIS EDITION

The fourth edition of TechForward was hosted at Goldman Sachs campus.

Extraction And Analytics From Financial Documents

***Prof. Kamal Karlapalem** explores use cases of AI in the fintech world, particularly those applications that are being researched upon at the International Institute of Information Technology, Hyderabad. This article is a brief summary based on the talk he delivered at the TechForward Research Seminar Series on 'AI Transforming Financial Services'.*

In the financial world, AI leverages machine learning, natural language processing, and predictive analytics to enhance the quality of services in a myriad of ways - from automating everyday tasks adhering to regulations, to identifying threats and preventing fraud.

Interplay Between Regulators and Cos

At IIITH, active research is underway to determine the ways in which AI can ease the operations of financial institutions. For instance, Indian companies are required by law to adhere to the regulations drafted by the Securities and Exchange Board of India (SEBI), which is the regulatory authority or the 'watch dog' of the securities and commodities market. It essentially protects investors by laying down a code of conduct that prohibits unfair and fraudulent trade practices. These regulations which are documented need to be interpreted by the information technology departments and/or finance and legal departments of the companies. The companies' lawyers interact with their SEBI counterparts about case arguments and outcomes. The documents themselves are interrelated and extracting the relevant interpretation is a complex task. It is here that a well-formulated AI system can aid and improve these interpretations and interactions by processing and deriving insights from SEBI regulations, associated case files and other pertinent documents.

IIITH's AI Framework

One of the ways the institute has enabled analytics on SEBI documents is by developing a multi-layer Applied Semantics Extraction and Analytics (ASEA) framework for document processing. Work in this area has been supported by the JP Morgan AI Faculty Research Award. The lowest layer of the framework deals with document pre-processing that includes entity extraction and entity linking. The middle layer deals with semantic analytics where classification, language modelling and so on is performed. The lowest and middle layer can be supported by GenAI solutions. The top layer requires user domain inputs and understanding to deploy the solutions. This layer is the applied semantics layer wherein the extracted semantics and analytics that is performed is used for various user relevant tasks. Thus the entire pipeline of the framework supports a range of use cases such as extractive question answering, provenance of the documents, regulation violation prediction, legal case file segmentation (through sentence classification), regulation simplification, and case file summarization.

Semantic Segmentation of Case Files

In this case, first a unique dataset was created with annotated adjudication orders pertaining to regulations. Adjudication rules are essentially rulings passed by SEBI against alleged violators of the SEBI Act, rules or regulations. So here, the IIITH team experimented with a number of machine learning models to train a sentence classifier that can help in semantic segmentation of case files which in turn helps in document retrieval.

Regulation Violation Detection

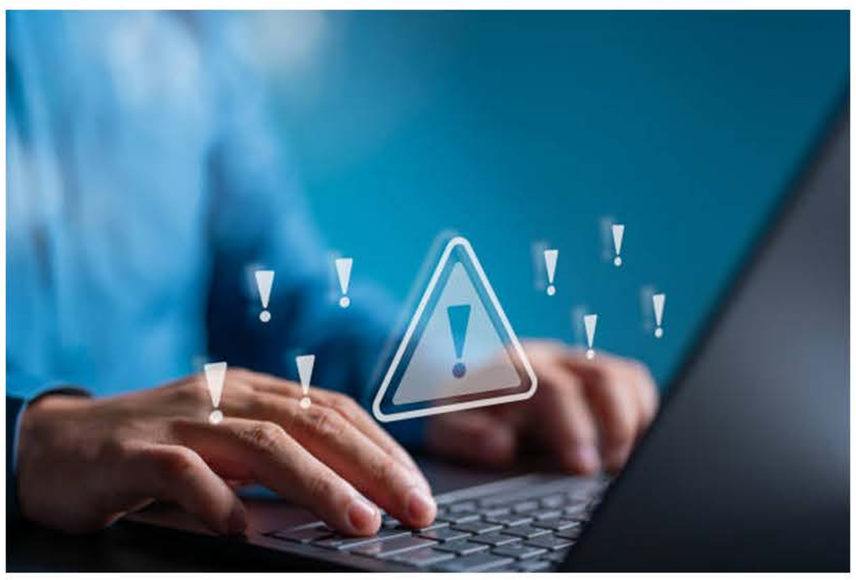

Typically a single case file has a multitude of sections. In order to zero in on the section which deals with violation of regulations, the team has developed a semantic segmentation engine that separates out the different sections of the case file. Transformer-based machine learning models for NLP tasks have become extremely popular and hence the IIITH researchers too built a transformer-based multi-label classifier. Traditionally such models perform poorly on domain-specific tasks as those found in the legal, medical or scientific domain, but fine-tuning them greatly improves their performance.

Similarly, in this case too, the ML model was fine tuned specifically for the SEBI domain on a dataset comprising SEBI regulations, case files, as well as SEBI-related news articles.

Regulation Biography

SEBI's regulatory documents are a mine of information which include their amended versions and additional supporting documents related to the domain of banking. With the help of AI, the IIITH team has demonstrated how this information can not only be analysed, extracted and tagged with the help of NLP methods but also how you can visually identify the changes made from one version of a document to another. Besides, these methods can also provide additional information extracted from annual reports and concept papers which will help understand the rationale behind the amendments itself. It also provides tags to categorise the type of amendments and identifies references to the regulations in news articles.

Challenges and Future Direction

The Applied Semantic Extraction and Analytics framework can be applied to many other domains like health, scientific literature, manuals, etc. The framework can judiciously use GenAI for some of the lower level ML-NLP tasks. The challenge is to deploy GenAI at the right places to come up with a production quality solution. Extracting and managing high level semantics for documents is a challenge IIITH is keen to take up. ASEA on documents like SEBI regulations and SEBI case files requires tailormade solutions and the hope is to see further research in this direction in the future.

PROF. KAMAL KARLAPALEM

IIIT HYDERABAD

is a Professor and Applied Computer Scientist at the Data Science and Analytics Centre and Agents and Applied Robotics Group, IIITH. His areas of research interest include Multi-Agent Systems Simulations, Multi-Robotic Systems, Workflow Management Systems, Visual Data Analytics, Data Mining, Text Analytics, and Database Systems. He also guides institutions for academics and curriculum development.

A Sneak Peek into a High Throughput Trading Database

***Alok Madhukar** provided insights into the design and architecture of a high throughput trading platform used in the firm that powers Goldman Sachs pricing and risk management discipline.*

This enterprise scale platform provides a reliable and resilient trade booking system. It also provides a hyper productive environment of code and data, which is used daily by thousands of quants and engineers to create cutting edge analysis and flows.This groundbreaking technology has been an inspiration for several copies in the market. However, none of these copies have reached the scale and power of this platform. Its ability to navigate difficult times is well known in the industry.

The platform was built in response to automation needs nearly three decades ago and was the first in the industry to do so. The platform was built for providing a persistence layer for all risk artifacts across the firm, real-time risk for a business across all locations, quick reproducibility of historical results, efficient what-if analysis, identical pricing across all divisions, and minimal time to market.

He further explained how these requirements were met using specific design principles. He concluded the discussion with remarks on modernizing the plant, including modern language integration, the use of open-source technologies, enabling a data-driven and a cloud-native approach, and AI/ML-based fine-tuning.

ALOK MADHUKAR

HEAD OF HYDERABAD ENGINEERING, GOLDMAN SACHS

is the Head of Hyderabad Engineering at Goldman Sachs. Prior to joining the firm, Alok held various academic positions at the Indian Institute of Technology Bombay, Innovation Alberta and the University of Calgary. Alok earned an MS in Computer Science from the University of Calgary

Efficiency And Security For Financial Applications with Federated Learning

***Prof. Sujit Gujar** discusses the emergence of Federated Learning (FL) while enumerating the ways in which the Machine Learning Lab at IIITH has been working on FL for the financial sector.*

AI and machine learning (ML) have significantly impacted our daily lives, including financial services. For example, they are used in fraud detection, predicting loan defaulters, portfolio management, stock price predictions for algorithmic trading, personalised customer support through chatbots, and more.

Consider a specific scenario: a company wants to create a module for predicting loan defaulters and offer it as a service to different banks. The company aims to enhance customer experiences and refine the underlying machine learning models. To achieve this, the company needs to gather data from client banks to improve the models. However, due to legal and privacy concerns, banks may be hesitant to share their data. Such situations have led to the concept of federated learning (FL).

Challenges with FL in Fintech

Such a decentralised approach in the financial sector is the way to advance. However, it poses challenges. Though the clients are not sharing data, there is the possibility of model inversion attacks, which may predict the presence of specific individuals. It may be problematic if the data is about loan applicants, and the applicant would not be the financial institution's customer who leaked its data. Another challenge is: Is such a model fair across demographics? From a bank's perspective, it may prefer to free-ride on other banks' contributions unless offered incentives.

IIITH's Contributions

At the machine learning laboratory in IIITH, our group has looked at all three challenges in adapting AI through FL in financial services and proposed innovative solutions combining machine learning and game theory. We have developed techniques that optimally learn in FL with minimal compromise on privacy. We have used incentive engineering concepts to ensure that all the participants receive fair compensation for their contributions. We proposed heuristics on combining different client models to ensure that the aggregated model is fair across demographics. Though we derive motivation from financial domains, the techniques we developed are general and can be adapted to multiple other domains as well.

PROF. SUJIT GUJAR
IIIT HYDERABAD

is an Associate Professor and Computer Associates Technology Chair Faculty at IIITH. His research interests lie at the intersection of machine learning, game theory and blockchain technology.

Debunking The Myth Around Post-Quantum Cryptography

Prof. Kannan Srinathan explains why the current frenzy about the post-quantum cryptography apocalypse is unwarranted.

There has been a great deal of buzz about post-quantum cryptography. This is because quantum computers are expected to be ready in the immediate future with the purported ability to break current encryption methods. The goal of post-quantum cryptography (also called quantum-resistant cryptography) is therefore to develop cryptographic systems that are secure against both quantum and classical computers, and can interoperate with existing communications protocols and networks.

Even as India nears completion of its first small-scale quantum computer at the Tata Institute of Fundamental Research (TIFR), there is a parallel discourse about the vulnerability of current systems and upcoming associated risks to security. One can argue that the paranoia associated with the latter is not justified simply because of a unique facet of quantum computing. Consider this: if a 2n-qubit quantum computer is required to break a cryptosystem and the maximum you have built is an n-qubit quantum computer; *how many such n-qubit computers are required to break the cryptosystem?* It is mistakenly thought that one just needs to juxtapose two such machines. But one really needs 2^n such quantum computers which are never going to be available!

Uniqueness Of Quantum Computers

Unlike in the classical world, where in order to have the computational power equivalent of a 64-bit machine, but the prevalent technology is that of a 32-bit one, all you have to do is to add the two 32 bits together and simulate a 64-bit machine. The quantum world does not work that way. Adding every qubit potentially doubles its power. So having two 32-qubit machines is only as powerful as one 33-qubit machine and not 64.

As a direct consequence of the above fact, if n-qubit machines are the maximum available, the 'key' (pun unintended) to protecting an algorithm from being broken is to use a large enough key so that say only >2n-qubits machines can break it, which is exponentially beyond the reach of an n-qubit machine.

Currently, quantum computers that break the prevailing cryptosystems exist only on paper. They don't exist in the industry or in reality. At the risk of sounding controversial, it is not even ruled out that they may never ever exist. The shift to post-quantum cryptography could actually be never necessary, because the human race may never be able to build a quantum computer that will break the current cryptosystems. And there's a reason for that. As already pointed out, we need a technology that can design n-qubit quantum computers, for any given value of n, without a pre-set bound! This is beyond just technology, as there is a distinct possibility that Nature may not even permit it beyond a limit!

Old Wine In New Bottle

There is another perspective with regards to post-quantum cryptography. That it has nothing to do with either quantum mechanics or with quantum cryptography. In fact it refers to those classical, mathematical algorithms (some of) that were discovered long before Feynman suggested quantum computing in early 1980s. It is just that their security is based on NP-hard problems, which are conjectured to be beyond the reach of efficient quantum algorithms. One can only speculate about their non-usage then despite their discovery and a probable reason could be due to the large size of the public key. Advances in this respect are happening at a rapid clip, and are certainly welcome, with (or without) their post-quantum nature.

Conclusion

It is not our intention to either argue for or against research in post-quantum cryptography. All we are saying is that the ability of quantum computers to crack existing cryptosystems is moot at the moment. And the likelihood that it will happen in the immediate future is also unclear. Hence, I don't see the need for banking or other financial institutions to be in a rush to overhaul or change anything that is working perfectly fine right now.

PROF. KANNAN SRINATHAN
IIIT HYDERABAD

is an Assistant Professor at the Center for Security, Theory, and Algorithmic Research, IIITH. His current research interests include cryptography, quantum algorithms, and all aspects of theoretical Computer Science.

A community initiative, anchored by

INTERNATIONAL INSTITUTE OF INFORMATION TECHNOLOGY
HYDERABAD

IIIT Hyderabad

TechForward
DISPATCH