

Ethical Issues in Technological Era

The developments in technological area have accelerated tremendously in the past decade, throwing up both great convenience and learning and also new challenges in the world of finance. We still lack sufficient material to fully understand the impact of the machine learning systems on the financial architecture and the shape of the financial institutions of the future. The increasing use of technology in traditional areas of banking, insurance and capital markets has both expanded the options available to the customer and automated some actions, significantly reducing costs and nudging the financial system towards fairness and democracy. This aspirational attribute of the uses of technology drives us to many ethical questions which need to be answered today. The pace of developments in technology leaves no time to think, putting considerable pressure on the ethical frameworks to keep abreast.



CA. M. Ravindra Vikram

The author is a member of the Institute. He can be reached at mrvikram57@gmail.com and eboard@icai.in



Ethics and the Development of Technology

Simply stated, Ethics is about doing “good” and not causing “unnecessary harm”. The alignment of the value systems of the humans with that of the technology being developed is essential if the machines should also do good, as per the ethical values we believe in.

Theoretical developments in sciences would not have been possible if there were limits placed on the imagination of the scientists and technocrats working with the problems at hand. The developments in the world of science in the twentieth century not only transformed our lives

unalterably, but also made it possible for some countries to develop nuclear bombs. Could society have had the power and reason to control these developments if limits were placed on the imagination of the scientists and harness the same for only ethically appropriate technologies?

There are two opposite views on the ethical impact of the development of new technologies. For the consequentialist approach, science and technology are ethically neutral and it is only the application of technology that may lead to unethical ends. On the other hand, for the non-consequentialist or

A rules-based AI system expects rationality in the decision making process. Rationality need not be always the optimum for creativity, limiting the frontiers of the development of the human mind.

the deontological approach, it may be argued that science and technology has to be always seen through the prism of ethical values at every stage of action in the development. Consequently, all involved in the development of these technologies including scientists, technocrats, entrepreneurs, institutions and regulators, government and even citizens have to find their ethical responsibilities while participating in the development of technologies. Extending this argument further, the ethics of technology can be examined in the context of social or communal harmony, rather than the benefits to the individual. Typically, Professional Ethical frameworks are designed to satisfy community needs while determining the code of ethics for the individual practitioner.

The ethical development of technologies should stand the three way test. Firstly, what should we do with the Technology. If the technology can cause unnecessary harm, however remote those chances may be, should it be developed? In the long run, technologies must be judged not only from

the point of who can benefit from it but also from who is likely to be affected by it. Secondly, what can we do with the technology. If under the present laws and regulations, it is prima facie illegal or blasphemous, thought must be given before the technologies are experimented with. Thirdly, what can technology do? This question must be answered in the context of whether there are frontiers of what we can explore and judge for ourselves at each stage the value of breaking with tradition and societal norms.

Human values also vary over time. It has taken many centuries for societies to accept women's right to vote in a democracy. If the design of the AI system had happened in the previous century, such abhorrent practices would have been part of the ethical frameworks of the designers. It is time we understood that the design of an ethical framework of the AI systems must be enduring not only to time tested human values but also be flexible enough to recognise that the value systems change over a period of time.

A rules-based AI system expects rationality in the decision making process. Rationality need not be always the optimum for creativity, limiting the frontiers of the development of the human mind. Inefficiencies, mistakes, playfulness and even serendipity are human values. Killing them in exchange of rationality and increasingly

relying on technologies could prevent humans from discovering something unusual. In one sense, it is these human values which will endure and add to the progress of technology.

The notions of right and wrong travel a little beyond the explicit words captured in law. But who determines what is good? Are the values of being good universal? Can elected politicians, dictators, influential businessmen, university professors be trusted to make informed decisions on these ethics?

Human Machine Interaction

Algorithms shape the way machines are learning about decisions and actions. Machine learning is designed to use data from its experience to generate new data or even a new set of rules based on the inferences from the data it acquires. Humans are intuitively geared to take decisions even when inadequate data is available. Machine learning systems will have to take actions based on the set of rules embedded in its system, as designed by the technocrat. However, we still do not know enough about what the machine will end up learning. The consequences of decisions taken by machines can also be judged differently if the same had been taken by humans. The ethics of the use of drones in military warfare is a case in point. In spite of advanced image processing and face recognition techniques, the drones have frequently ended up killing civilians.

As machines become more sophisticated, ethics demands that transparency must be a priority. Customers should be informed that their solutions are guided by machines and not humans, as this makes a difference to many people.

Society is only now beginning to grapple with the ethical issues in such situations.

Machine learning systems are learning very fast, though. While it is still in the realm of science fiction, the day may not be far behind when machines could be super intelligent, simulate human minds and in instances transcend them. This idea, referred to as a 'singularity', refers to a point of technological progress when machines begin to control humans and at some point, may be uncontrollable or irreversible. Should we begin to worry about the ethics of such possibilities?

The increasing use of AI in financial services results in frequent human interaction with machines. AI chatbots are becoming better and better at modelling human conversations, but it must be admitted that they are still not human. As machines become more sophisticated, ethics demands that transparency must be a priority. Customers should be informed that their solutions are guided by machines and not humans, as

this makes a difference to many people. It is generally believed that humans have a level of empathy which machines are incapable of. The gap is narrowing, though.

Oftentimes, the machine chatbots have to convey a message that it is "sorry" for the delay or difficulty in fulfilling a transaction. Feeling sorry is a very human emotion which is imposed on the machine. When a human feels sorry, there are

When a human feels sorry, there are consequences of the regret expressed. A machine learning system could then be trained to alter the course of its action consequent to expressing regret, very much as a human would do.

consequences of the regret expressed. A machine learning system could then be trained to alter the course of its action consequent to expressing regret, very much as a human would do. The difficulty would be to parameterise emotions of regret and consequences to the expression of regret. Is empathy a necessary ethical value to be protected in the development of technology?

The Bias of the Developer

Financial markets are being expanded to develop AI systems through the development of a series of rules or algorithms. The algorithm,

written by a developer, can direct the customer towards exercising choices in a subtle way offering some choices and not others. The algorithms may be written based on customer profiles, behaviour, risk taking abilities, previous decisions of the customer, etc. These AI systems are created by humans who can be biased and judgemental. It is logical to assume that the developer will include some of his biases, whether deliberately or unconsciously, on the development of the AI system.

For example, many financial websites now run chat-bots, which are driven by AI. They may act in a political way by steering a discussion towards mis-selling products, or profiling and consequently giving access to systems which are available to some whose abilities match the profile the system recognises and denying others who it refuses to. In the alternate, the system may be designed to exclude or limit the use of certain financial services by people belonging to a certain pin code, as the programmer knows fully well that the people in that geography are largely under privileged. This meets the goal of the designer of the financial services, because he does not have to deploy capital and resources on what appears to be a risky transaction. But this perception is that of the programmer or the person controlling the process of choice. These systems work on developing a pattern of users and it is possible that the

Can we afford even a single error? If the system were to make an error, who is to blame? Is it the developer of the software or the owner of the software or the regulator who has missed the possibility of detecting the error or is nobody to blame?

pattern does not include users who were out of the sample. Is it fair to exclude such persons? Who has given the system, the right to exclude? What will the system do when such an excluded person chooses to access the system?

One of the important areas of AI development is in voice activated services. If natural language processing systems are used for information access by the customer, it raises issues of awareness and accessibility. Imagine a spoken dialog system that is easy to use for a young financial professional user with a urban middle class background's pronunciation, but that may barely work for an elderly lady from rural India. As automated information systems are becoming more pervasive, they may eventually substitute human information kiosks for cost reasons, and then out-of-sample user groups could be excluded and left behind without an alternative. Is it ethical to deny access to benefits of AI technology, succumbing to the bias of the developer?

What happens if the AI system makes mistakes?

Serious risks to financial systems may arise when we recognise that AI systems are not infallible. These risks include risks of contagion or causing instability in the market which may run out of control, however brief the event horizon is. High-frequency trading, for example, is a classic case of AI based algorithms achieving what humans cannot physically accomplish in the given time. Nearly 80% of the trades in the US stock market are now done by computers without any human intervention. The machines dominance of the markets is sure to extend further. AI based computers move in and out of short term positions at high volumes and speeds aiming to capture the minutest of profits per trade but relying on large volumes. The speed at which information is absorbed makes it possible for the price to be as close to the underlying value.

The firms using AI cannot be blamed for implementing these trading strategies in spite of not understanding the models fully due to its inbuilt complexity. Machine learning systems may also be adding its own rules to the algorithms adding complexities not known to the developer either. Consequently, when AI machines interact with each other, there may be unintended, unanticipated consequences in the models resulting in chaos in the market. Can we afford even a single error? If the system

were to make an error, who is to blame? Is it the developer of the software or the owner of the software or the regulator who has missed the possibility of detecting the error or is nobody to blame?

Three incidents come to mind in this context: the 2010 Flash Crash, Knight Capital's \$140 million loss due to a trading error in 2012 and more recently, an innocent login glitch at the Tokyo Stock Exchange. In the case of Knight Capital there was no malintent, it is said, on the computer's part, yet its actions wiped out 75 per cent of the firm's equity value within days. There is a responsibility that lies with the developer that the AI system acts as planned and there is sufficient learning built into the software to "learn" from its mistakes. Otherwise, a single error may cause a cascade of actions leading to a risk of contagion. In today's globally connected capital markets, we can least afford to get to this stage. Is it ethical to cede control of decision making to machines and only be prepared to face the consequences of mistakes? But on the other hand, undermining a technology just because it is difficult to find a person to blame in the event of failure is counterproductive to progress.

The question will remain that when the market crashes and the situation is unfamiliar to the system, will the AI system itself crash, or will it make decisions to maximise profit, heedless of the damage caused to the ecosystem in which it exists.

Humans who rely for a long period on AI systems run the risk of losing experience in dealing with situations when the AI system fails. Aircraft pilots relying on AI systems for too long for flying may not even know or draw from experience in the event of failure of the system. Designers of systems may even “forget” that some possibilities of mistakes exist. When such black swan events happen, humans must be capable of reacting to the event according to the ethical values built in the individuals reacting. This will be a challenge, though.

AI, Financial Markets and Trust

Financial markets are essentially trust based systems. There are rules of ethics between the constituents of the markets (Stock exchanges, Banks, Insurance companies, Investment bankers, etc.), which are designed to keep the markets stable, especially at times of volatility. These rules of ethics go a little beyond law as they are also designed to guarantee a level of credibility, professionalism, quality of services. More importantly, they instil confidence in the users of the services. These attributes appear to be personal and linked to human relationships and the way individuals would react in given situations. Developing AI systems, especially in the world of finance, need to build sufficient trust mechanisms to protect the integrity of the systems. Trust-based mechanisms have inbuilt safeguards to protect the

system by giving it the primacy of survival over individual constituents also be socially responsible.

The working of the trust mechanism is a process which is actually not fully understood by the customer (and need not be understood either), but perceived by her customer to exist in a rational form. This trust can be expressed as a confidence measure, just as we are confident that the brakes of a car would work when applied and need not know how it works. The institutional mechanism of protection of trust has Regulators, Government, Courts working towards providing a fair and transparent mechanism to the customers and protecting them at times of vulnerabilities. Over a period of time, as the AI system keeps learning, these rules will embed within themselves fair practices which improve the trust of the customer in the system. Can we always trust the AI driven system to be fair, transparent and trustworthy, without human intervention?

To transfer this trust to an AI system will need a new set of ethics of accountability and autonomy. It will be necessary to depend on people to build ethics into the objectives for the systems through accountable governance. It is time that Regulators catch up with this.

The Issue of Data Stewardship

The 2009 “Galway Paper” that established the Essential

Elements of Data protection Accountability says:

“Accountability is the obligation to act as a responsible steward of the personal information of others, to take responsibility for the protection and appropriate use of that information beyond mere legal requirements, and to be accountable for any misuse of that information.”

Many ethical issues arise on the use of data responsibly and to protect people from issues of privacy and data breaches. In a landmark move, the European Union has codified the General Data Protection Regulation (GDPR), emphasising on the way data is handled vis-à-vis citizen’s rights. Data Ethics focus both on the ethical rights of an individual on data generated by individual and the ethical use of data by entities which need data to develop algorithms. AI takes data stewardship to a new level, building ethics into the systems through accountable governance. Data Stewardship has to resolve issues on ownership of data, clearly laying down the rules of sharing the data for analytical purposes. Consent must be obtained from the persons whose data is being shared. While this may prove difficult in practice, firms should not be allowed to hide behind legalese to obtain data by stealth. At all times, users of financial systems must be made alert to the possibilities of abuse of the data to the detriment of citizens. Data Security too assumes importance and customers must be made aware

AI takes data stewardship to a new level, building ethics into the systems through accountable governance. Data Stewardship has to resolve issues on ownership of data, clearly laying down the rules of sharing the data for analytical purposes.

of the systems which holders of data employ to protect the data from theft or abuse.

AI and the Regulations in the Financial World

Neither businesses nor regulators may fully understand the consequences of the software being developed. The unbridled growth in blockchain technologies and cryptocurrency is a case in point. Central banks around the world are concerned about the potential harm posed by current usage of these often poorly understood crypto assets. Surely, the software programmers pioneering the currency have not clearly understood the ethical issues in introducing the technologies in the market and cannot be expected to do so. Increasingly, there are third party dependencies on developments of technology outside the regulatory purview, risking the ethical bases of developments in AI.

Before adopting new technologies, regulators and the government need to

understand their current data, processes and infrastructure in order to keep up. Poor quality data is a burden no one wants to own. To conduct a successful digital transformation, without adequate preparation and understanding of the underlying data is unethical.

There is a serious risk of concentration of market power by those controlling data. Regulators across the globe must develop capacities to develop norms for data flow across borders and also limit regulatory arbitrage in data usage.

The way forward

Like the proverbial sword, disruptive technology, if used incorrectly, can act as a weapon in the wrong hands. Fortunately, with the right behaviour and conduct, it can form a protective force to herald the move of centuries-old financial institutions into the future. There is no doubt that compliance cannot compete with technological innovation, but it can certainly adopt it.

Society, including Governments, Regulators, and multilateral bodies have to agree on a set of standards for an Ethics review board on introduction of new technologies, which is not very different from the Bank for International Settlements. The Board must have space for all stakeholder views including technocrats and under privileged groups in society like poor and racial and ethnic groups. The Board approves

at every stage of Research, Product Development and Deployment to test it on the platform of ethical standards.

There is a need for a multi-stakeholder collaboration to understand the immense benefits the new technologies of finance offers as also to mitigate the ethical risks that these technologies need to address.

While the revolutions in the technological era brings in fundamental changes in the financial world, the time has come that we ask the right questions on the ethical issues these changes bring up. As a profession which deals with the core of changes in the financial world, the Institute of Chartered Accountants of India is best placed to take up such matters. India, with its multicultural and diverse ethnic groups offers a wonderful platform to lead the world to develop these ethical standards acceptable to all. ■

References :

- Artificial Intelligence: potential benefits and Ethical Considerations:* A European Parliament briefing
- Artificial intelligence and machine learning in financial services: Market developments and financial stability implications* - Financial Stability Board, November 2017
- Artificial Intelligence, Ethics and Enhanced Data Stewardship:* Information Accountability Foundation, September 2017
- Ethical Use of AI in Finance:* The Finance innovation Lab, February 2018
- The New physics of Financial Services: How AI is transforming the financial ecosystem:* World Economic Forum, August 2018