Introduction to AI Ethics and Governance

Reynold D'Silva[1] and Manik Bhandari[2]

[1]GO-JEK, Singapore
[2]Vulcan AI, Singapore

I. Introduction

Organisations that seek to develop or deploy AI solutions should include AI ethics in their corporate governance frameworks. That's because what may be considered as being suitable for an organisation may not always be good for its customers. It is, therefore, a good idea to set up governance structures with representatives from the business, legal, privacy, security and technology wings of the company, as well as include external subject matter experts (SMEs) as advisers.

External SMEs may offer unbiased views and provide checks and balances in AI solution development or deployment. In issues concerning ethics – generally, as well as related to AI – external or independent SMEs should be given weightage to ensure purely commercial considerations don't outweigh ethical ones.

A. Core Values

A company's core values are the basis on which its decisions get made. Top and senior management teams could be held accountable and liable for contravening the company's core values. That's why the core values could be a benchmark for companies to refer to, especially in case of doubt. AI solutions that are not well calibrated may have broad ramifications that could impact the company's reputation.

a. Corporate Mission

A company's corporate mission should be clear, precise and aim to inspire its staff to rally around a shared purpose and vision. Most companies will try to focus their resources and investments aligned with the company's purpose and mission. Similarly, it would be wise to ensure that its investments in AI, including AI ethics and governance, are also aligned with ethical guidelines and protocols.

o Example: Parkway Pantai Hospital Group's [1] stated mission is "to make a difference in people's lives through excellent patient care". One of their values is "excellence by striving for the finest clinical, service and operational outcomes". How does this manifest in reality? One issue patients used to face was not having a good idea of the cost of specific medical procedures. Bill estimates provided by the hospitals were generated using statistical methods, which were not accurate, resulting in "bill shock" for some patients.

o In December 2018 [2], Parkway Pantai reported that the Group would use AI to dynamically generate personalised, more accurate hospital bill estimates that vary from the actual bill by a low 18 per cent on average, a significant 60-percentage-point improvement over the current bill estimation system. That meant that patients would receive highly accurate bill estimates that fell within an 18 per cent margin from the final bill figure.

o "Using an advanced suite of AI and ML algorithms from UCARE.AI, a Singapore- based AI start-up, Mount Elizabeth, Mount Elizabeth Novena, Gleneagles and Parkway East hospitals will dynamically generate personalised bill estimates based on parameters such as the patient's medical condition and medical practices," the hospital reported [2]. "It also takes into account the patient's current age, revisits, and existing co-morbidities like high blood pressure or diabetes."

o The AI solution consists of a "Cost Predictor" and a "Price Guarantee Programme" for six surgical procedures – removal of piles, breast lumps, ovarian cysts, gallbladder, thyroid and tonsils. The AI program computes charges for these procedures against the Cost Predictor's price estimates.

o The hospital offered patients guarantees that they would be charged according to the initial estimate quoted by the Cost Predictor, even if treatments were added later. The hospital's financial counsellors worked with patients before or during admission to review their estimated bills. Patients received a precise estimate of the final bill and avoided the "bill shock syndrome".

b. *Balancing Act*

Organisations have to balance commercial objectives with the risk of deploying a new AI solution. Holding prior discussions and aligning issues with the company's core values can give corporate teams enough fodder to decide whether to proceed with the new AI solution or tweak it to comply with ethical guidelines and protocols.

o Example: In March 2020, Microsoft said it would no longer invest in third-party facial recognition companies, following controversy around its funding of Israeli start-up, AnyVision. Microsoft had previously turned down a request from law enforcement in California to use its facial recognition program in police body-cameras and cars, Reuters reported [3] on March 27, 2020.

o Speaking at an event at Stanford University, Microsoft President Brad Smith said the company was concerned that the technology would disproportionately affect some groups. It might lead to innocent women and/or minority groups being disproportionately held for questioning because the AI had been trained on mostly white and male photographs. "Anytime they pulled anyone over, they wanted to run a face scan against a database of suspects," Brad Smith said without naming

the agency. After thinking through the uneven impact, "we said this technology is not your answer." Microsoft was willing to turn away business when it conflicted with its core values of trustworthy computing.

c. *Governance Structures*

The selection of SMEs with appropriate levels of experience or expertise is critical to ensure corporate governance and ethics bodies stay neutral and focused on doing what's right, and not what's merely profitable. On the other hand, SMEs with hidden agendas can undermine the credibility and effectiveness of the governance   structure.

o Example: On March 26, 2019, Google [4] set up a body to look at ethical challenges in AI development. The body, called Advanced Technology External Advisory Council (ATEAC), had eight members, including economists, philosophers, policymakers, technologists and others. It planned to discuss and debate issues on algorithmic bias, among others.

o The intent was excellent, but the outcome was not. Two of the members nominated to the panel were alleged to have an inherent bias: one had ties to the defence sector; another held controversial views on issues such as climate change, immigration and diversity rights. A third nominee declined to join the council. A week later, in response to a letter signed by more than 2,000 Google employees, the ATEAC was shelved, The Guardian [5] reported on April 5, 2019.

o "It's become clear that in the current environment, ATEAC can't function as we wanted. So we're ending the council and going back to the drawing board," Google's blog [4] on April 4, 2019, stated. "We will continue to be responsible in our work on the important issues that AI raises, and will find different ways of getting outside opinions on these topics."

This example highlights the challenges that companies face as they grapple with the complexities of aligning their AI ethics and governance with their corporate values and mission. In Google's case, employees had strong concerns about the company's approach in this specific situation.

B. Core Concepts

This chapter aims to guide organisations to develop appropriate internal governance structures to monitor how AI technologies are deployed in operations, products or services. These checks and balances are necessary to ensure ethical oversight over the use of AI.

An effective corporate governance structure is necessary for any business to meet its strategic goals. The governance structure should include controls and policies that help the company meet its objectives while satisfying ethical guidelines. There are three types of governance mechanisms in most medium and large companies:

a. Internal Governance Structures

All organisations have some internal structures or mechanisms to monitor their activities and take corrective action when required. These actions can be preemptive (such as during the process of development or deployment) or post-operational (when the product/service/solution is deployed in either the Alpha or Beta stages). Internal governance structures should involve management (senior, middle or junior), line and operations staff, and even suppliers or channel partners.

The objective is to ensure regulatory and ethical compliance. The operations should have clearly-defined reporting lines and performance measurement systems. Internal governance structures should include oversight of management, independent internal audits, and overall corporate ethical responsibilities by the board of directors.

b. External Governance Structures

External governance structures or control mechanisms are outside an organisation's management ambit. External control mechanisms are necessary to comply with laws, rules and regulations set by a country's regulators, government agencies, trade unions and financial institutions. These governance structures may deal with debt management, legal compliance, and ethical guidelines, with penalties for non-compliance.

Demands can also be placed on organisations by external bodies such as trade unions or industry associations. These may be in the form of guidelines or best practices. They may be optional, but deviations from usual industry norms could attract the attention of regulatory or fiduciary agencies.

c. Audit Structures

Audits serve an essential function in organisations by providing the necessary checks and balances. Audit governance structures can be either internal or external or both. An independent external audit of an organisation's financial statements is part of the overall corporate governance structure and process; it serves both its internal and external stakeholders.

An audited financial statement and the auditor's report helps investors, employees, shareholders and regulators assess the financial performance of the organisation. Most companies have multiple levels of audit, or oversight structures, including at the board level with "audit committee" members being part of the board of directors.

d. SME Structures

It's not just medium and large businesses that need to have formal corporate governance structures. It's all the more relevant for small and medium enterprises (SMEs). In Singapore, an SME is defined as an enterprise with an annual sales turnover of under $100 million, or with less than 200 staff.

In smaller SMEs, business owners make strategic decisions about how workers will do their jobs, and how their performance is monitored. This is part of its internal control mechanism or governance. It is essential to have such structures in place for SMEs because of external controls: If the business requires a loan from a bank, it must comply with the bank's terms and conditions. If the business is a partnership, a partner might demand an independent external audit if he feels something is amiss.

e. Alignment

Organisations can use three basic principles to align policies related to ethics and governance:

- Corporate communications, to share policies and governance plans with internal and

external stakeholders.

- Encouraging, listening and responding to feedback from internal and external stakeholders, especially customers.

- Alignment with the organisation's AI principles or policies. This can be reinforced with clear and consistent messaging, holding regular training sessions for staff, and getting staff to adhere to standard operating procedures.

- The IMDA Model AI Governance Framework [6] offers good guidelines for companies to follow:

Figure 1: The IMDA Model AI Governance Framework

II. Setting up Structures

Many start-ups – and a few mature companies – in AI development, products and services, are SMEs. The use of AI solutions will help double employee productivity and the pace of innovation in Singapore by 2021, according to a study by Microsoft and IDC Asia-Pacific [7]. "Singapore's business leaders and workers hold positive viewpoints about the AI's impact on the future of jobs. The majority (62 per cent of business leaders and 71 per cent of workers) believe that AI will either help to do their existing jobs better or reduce repetitive tasks," the study [7] reported. "When it comes to creating or replacing jobs, 29 per cent of business leaders believe that AI will produce new jobs, whereas 9 per cent feel that AI will replace jobs. Interestingly, workers are more optimistic, with only 4 per cent expecting AI to replace jobs, and 10 per cent anticipating AI to create new ones."

The study also found that workers were more willing to reskill than business leaders believed – 20 per cent of business leaders said it might be too difficult for workers to develop new skills. In contrast, only 12 per cent of workers felt that it was a challenge.

The organisation's existing internal governance structures can be adapted – or new structures set up – to include AI ethics and governance. Alternatively, risks associated with the use of AI

can be managed within the enterprise risk management structure; ethical issues can be handled through ethics review boards.

Which features should be included in an internal governance structure? Should you opt for a top-down, centralised governance structure? Or a loosely-structured decentralised model? What are the pros and cons of each? Which kinds of organisations would benefit from which type of structure? The bottom line is clear: Whatever the structure, it is essential to have top management's – and the board of directors' – support, sponsorship and approval for the process to be effective.

A. *Centralised vs Decentralised*

Ethics committees were first created as a solution to deal with clinical dilemmas. They helped in clinical decision-making (through consultation) and helped develop institutional ethics policies. A University of Pennsylvania study shows that the power and activity of ethics committees depends upon the composition of the committee and the relationship of members to the parent organisation.

a. Ethics Committees

Many organisations have some form of ethics governance policies or committees. That's because products or services with embedded AI might at times behave in unintended ways. With the widespread adoption of AI technologies, more firms may need to have such bodies. Research by Fitzgerald and Phillips [8] suggests that there are at least three types of systems – fully centralised, dual, and decentralised/multi-committee. Each can have up of two interrelated components – the administrative process, and the ethics review.

b. Centralised Structures

Centralisation of the governance process can help an organisation gain visibility across different departments. Centralised structures can set policies that generally apply across

the enterprise and help tackle the most common issues. Centralisation of the ethics review can help the company deal with a multitude of related issues and evolve compliance policies or guidelines. The centralised approach may work well when the ethics committee and the board of directors adhere to a robust ethical culture. It may fail when the profit motive sways this group.

o Example: Boeing has a centralised Office of Internal Governance and Administration under a Chief Ethics and Compliance Officer who reports directly to Boeing's Board of Directors. This structure was meant to ensure that employees in different locations and divisions would not have to make different decisions on the same ethical issue.

o According to The New York Times [9] report on March 23, 2019, many Boeing employees had privately discussed problems with the design and decision-making process on the 737 Max aeroplane. The paper reported cases when managers dismissed engineers' recommendations or just prioritised profits. This shows the limitations of the centralised approach.

c. Decentralised Structures

Decentralisation of the ethics review allows for more empowered and rapid decision-making by teams that are most knowledgeable about the issues as well as the fallout from non-compliance. The governance process can also be decentralised or democratised, allowing for a wider net of potential ethical issues to be identified.

o Example: In 1990, Unilever set up the Safety and Environmental Assurance Centre (SEAC) [10], which in 2019 consisted of 200 scientists who independently carried out safety risk and environmental impact assessments of new products and processes. The SEAC made decisions independently of business and profit considerations, keeping sustainability, and safety of employees, consumers and society in view. "This means new products and processes are always designed to ensure that Unilever supplies products that

are safe for our consumers to use, for our workers to make, and for the environment," the company reported.

o That is a great example of independent and decentralised oversight. Most companies would not need 200 scientists to do something similar. However, by removing profit and revenues as being the motive, companies can reassure customers and regulators that it pays adequate attention to environmental, social and ethical values.

d. Pros & Cons

A centralised structure makes it easier to have a consistent and standardised process to evaluate the risk and impact of an AI solution. It also helps to assess that it aligns with an organisation's mission and core values. There is also greater control over the review process.

However, having a significant centralised function requires massive investment. In MNCs with presence in multiple markets and countries, it may be challenging to incorporate local moral, ethical and legal norms in a centralised governance structure. If every AI process or model needs to go to the HQ for review and approval, it will slow down the decision-making process, impacting the business at various levels. A decentralised structure might be ideal for MNCs spread across multiple markets.

The question is not which model is better, but which functions need to be handled in a decentralised way, and which functions can be handled in a centralised structure. Organisations can also use the Risk-Impact Assessment Matrix [11] to help them decide the probability and severity of impact, and if the AI solution needs to be escalated to a centralised governance team for review.

Some organisations have found that a dual structure, which combines a centralised ethics review with a decentralised administrative process, is better suited for their corporate

culture. There is no one superior approach; organisations must consider which would be the best fit for their culture.

Here are some real-world examples from Business Ethics: Ethical Decision Making & Cases [12], by John Fraedrich and OC Ferrell:

Company	Culture	Characteristics
Nike	Decentralised	Creativity, freedom, informality
Southwest Airlines	Decentralised	Fun, teamwork, loyalty
General Motors	Centralised	Unions, adherence to assignments, structured
Microsoft	Decentralised	Creative, investigative, fast paced
Procter & Gamble	Centralised	Experienced, dependable, proven historically

B. Role of the Board

The board of directors need to get more involved in the corporate AI discussion. Senior managers and AI team members should provide specialised training to the board on issues such as the AI solution and its features, and the potential risks and implications. The discussions need to be frank and forthright so that the board has all aspects of the situation to exercise oversight across the enterprise.

"Board members and corporate executives of all companies are responsible for stewarding their companies through the current period of unprecedented technological change and its attendant societal impacts," according to the World Economic Forum [13]. "A practical set of tools can empower them in asking the right questions, understanding the key trade-offs and meeting the needs of diverse stakeholders, as well as how to consider and optimise approaches when overseeing and operationalising AI, to ensure ethical and responsible implementation."

In January 2020, the WEF published an AI Toolkit for Boards of Directors [13] in collaboration with key stakeholders from diverse companies and industries and academia. The Toolkit has 12 modules designed to help company directors understand the impacts and potential of AI in their company strategy, which may include customers, brand,

cybersecurity, employees, operations, and competition, as well as the main areas of control, such as audit, risk, and governance.

The WEF project team is currently co-creating a new toolkit focusing on how to operationalise AI for C-suite executives. It will address topics that are specific to conventional C-suite positions (CEOs, COOs, CFOs) as well as emerging positions (Chief Data Officer, Chief Digital Officer, and others).

a. CEO, CAEO, CECO?

The CEO or Chief Executive Officer is a universally accepted abbreviation. But then, what do we call the Chief Ethics Officer in an organisation? Should she/he be called Chief AI Officer (CAIO)? Or Chief AI Ethics Officer (CAEO)? Or Chief Ethics & Compliance Officer (CECO)?

b. Brief History

The US Congress passed the role of a Chief Ethics Officer in 1991 under the Federal Guidelines for Organisations [14]. This role had existed in the financial services and the healthcare sectors even before 1991 but was not quite as useful. In October 2001, the Enron scandal [15] erupted and eventually led to the bankruptcy of the US energy giant, as well as the de facto dissolution of Arthur Andersen, which was then one of the five largest audit and accountancy partnerships in the world. Apart from being the most prominent bankruptcy reorganisation in US history, Enron was cited as the biggest audit failure.

That led to the role of the Chief Ethics Officer becoming paramount. The role's importance rose after the passage of the Sarbanes–Oxley Act [16] in 2002, and amendments made to the US Federal Sentencing Guidelines. In 2002, the US Congress passed the Dodd-Frank Wall Street Reform Act and the Consumer Protection Act [17],

putting huge responsibilities on the Chief Ethics Officer. Now, the position is responsible for developing and ensuring codes of ethics, developing training programmes for staff, and monitoring and auditing compliance with government regulations.

c. Do You Need a CAEO?

The critical question your organisation needs to ask is whether you need a Chief AI Ethics Officer. In 2018, Deloitte surveyed 1,400 US executives with knowledge of AI and found that 32 per cent ranked ethical issues as one of the top three risks of AI.

While there are myriad ways for companies to access ready-made AI or develop their own, many also seek outside expertise. A total of 53 per cent of respondents said they co-developed cognitive technologies with partners, and nearly 40 per cent used crowdsourcing communities such as GitHub.

"Through cloud services and enterprise software, companies can try cognitive technologies and even deploy them widely, with low initial cost and minimal risk. The growing number of cloud-based options may explain the spike in pilots and implementations between 2017 and 2018. About 55 per cent of executives say their companies have launched six or more pilots (up from 35 per cent in 2017), and nearly the same percentage (58 per cent) claim that they have undertaken six or more full implementations (up from 32 per cent)," the Deloitte study reported [18].

In much of the tech world, the concept of a Chief Ethics Officer remains a hard sell, in large part because of pressures to stay competitive. "Being first in ethics rarely matters as much as being first in revenues," Timothy Casey, a Professor in Residence at California Western School of Law and a member of the Ethics Committee of the San Diego Bar Association, was quoted in a Forbes [19] article published on March 27, 2019. The other big issue is simply a matter of history. Professor Casey noted that while certain professions have ethics baked in from the start, computer programming decidedly does

not. "In medicine and law, you have an organisation that can revoke your license if you violate the rules, so the impetus to behave ethically is very high," Forbes [19] quoted him. "AI developers have nothing like that."

o Example: Salesforce.com appointed a Chief Technology Ethics Officer and an AI Ethicist to work with the AI production team. Other firms are exploring the creation of Ethics Advisory Boards. Microsoft has expanded its institution of Champs to include an AI Champ, a role that's responsible for all AI issues. Microsoft's framework [20] (image below) provides a practical checklist for AI issues that these executives need to consider and oversee.
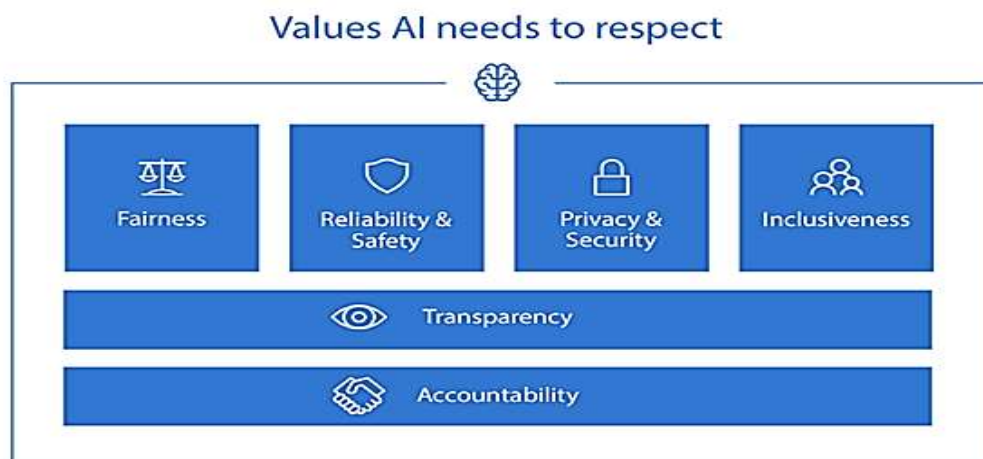

Figure 2: Values AI needs to respect [20]

d. *Values-based Approach*

Which approach is right for your organisation? It depends on factors such as the countries of operations, the market environment, the vertical segment, and the nature of your AI solution. It is essential to recognise the need to assign oversight responsibilities to the board of directors and get them trained on AI issues and outcomes. It would be wise to appoint a senior officer to oversee the day-to-day work being done by different teams and to identify and manage risks systematically.

III. Roles & Responsibilities:

A. By Function/Functionary

Function/ Functionary	Roles & Responsibilities
Board of Directors	Understand AI ethics issues. Ensure oversight. Ask probing questions. Make informed decisions.
CEO	Overall management and execution. Strategic decisions on where and how to use AI. Recommendations to the Board of Directors.
AI Ethics Committee	Independent oversight of AI ethics solutions/deployment. Draft and uphold AI ethics principles/Code of conduct. Iteratively learn from internal and external examples. Take required action to prevent ethics violations.
Chief AI Ethics Officer	Manage the administrative processes. Conduct periodic ethics review of all AI-related projects. Take required action to prevent ethics violations.
Chief Data Officer	Ensure consent for the collection and use of customer data. Ensure outcomes are always beneficial to the customer. Ensure customer data is protected and secure. Ensure customers cannot be identified by AI technology.
Chief Risk Officer	Provide the CEO and board of directors with a probabilistic assessment of any adverse outcomes from AI projects.
Chief Legal Office / Counsel	Assess potential legal liabilities. Prepare mitigation strategy. (Read Chapter 3.4).
Chief Information Officer	Audit risk and security of databases used by AI systems. Audit risk of AI solutions procured or used on a SaaS model.
Chief Technology Officer	Audit risk and security of AI customer interfaces. Audit risk of misuse of AI systems.
Chief HR Officer	Manage ethical issues in hiring, assessment and promotions. Manage impact of AI technology on jobs and careers.

B. By Stage of Tech Development

	Stage of Development				
	Data Provenance	**Model Verification**	**Model Testing**	**Interpret-ation**	**In-use Safety**
Assessments	- Data lineage - Data diversity - Flaws/Biases - Regulatory Compliance	- Algorithm soundness. - Extent & impact of false negatives or false positives	- Ethical goals & outcomes - Relationship between input/output	- Explainability of outcomes - Visualisation of deep neural networks	- Adversarial testing - Attack & defence algorithm testing

a. Data Provenance

Data lineage	Metadata. When, where and by whom the data was created.
Data diversity	Inclusion of data from multiple sources. Structured & unstructured data. Goal is to strengthen the ability of the data to model or represent an outcome.
Flaws & biases	Flaws include errors in data collection, errors in meta-tagging, and data dredging. Best explained as the patterns that emerge from correlation rather than causation. Biases include datasets that do not represent the entire population, or datasets that represent historical biases and prejudices.

b. Model Verification

Algorithm soundness	Soundness is the ability of an algorithm to return a true answer every time.
False negative	Output of a model that wrongly predicts the absence of an attribute/feature.
False positive	Output of a model that wrongly predicts the presence of an attribute/feature.

c. Model Testing

Relationship between inputs and outputs	Increasing the level or intensity of an input increases or decreases the volume or likelihood of an output.

d. Interpretation

Explainability of outcomes	A transparent report on why and how the model arrived at an outcome.
Visualisation of deep neural networks (DNN)	Visual representation of data, features and relationships in the DNN.

e. In-use Safety

Adversarial testing	Identification of weaknesses in the system through a cyclic process of attack surfaces and vectors, including datasets, input features, data labels, metadata.
Attack and defence algorithm testing	Testing system defences against the modification or deletion of datasets, input features, data labels or metadata.

C. Checklist

a. Does your organisation have a published AI Code of Ethics?

b. Is it aligned with your mission and culture?

c. Has it been communicated to key stakeholders?

d. Has feedback been taken and incorporated?

e. Have the board and senior leaders trained in the core concepts of AI and AI-related ethical issues?

f. Have you set up an internal or external governance structure?

g. Does your organisation's culture fit best with centralised, decentralised or dual-system ethics review?

h. Does your organisation have or need a Chief AI Ethics Officer?

i. Does your organisation have clearly defined AI ethics review processes?

j. Do key executives know their roles and responsibilities related to AI ethics?

k. Do you have appropriate assessment processes for each stage of the AI development process?

l. Do you have an anonymous whistleblowing and escalation process for ethical issues?

m. Will AI solutions cause job losses in your organisation?

n. Do you have a plan to manage outcomes for people affected by decisions made by AI solutions?

o. Are diverse points of view and datasets included in developing of your AI models?

p. Do you have safeguards to protect people from unexpected, unjust, unfair or dangerous outcomes?

References

[1]   Parkway Pantai, "Parkway Pantai - One of the largest private healthcare provider in Asia," [Online]. Available: https://www.parkwaypantai.com.

[2]   Parkway Pantai, "PARKWAY PANTAI AND UCARE.AI LAUNCH AI-POWERED PREDICTIVE HOSPITAL BILL ESTIMATION SYSTEM," [Online]. Available: https://www.parkwaypantai.com/docs/librariesprovider6/media-releases/media-release---parkway-pantai-and-ucare-ai-launch-ai-powered-predictive-hospital-bills-estimation-system.pdf?sfvrsn=9653871e_2.

[3]   J. Dastin, "Microsoft to divest AnyVision stake, end face recognition investing," [Online]. Available: https://www.reuters.com/article/us-microsoft-anyvision/microsoft-to-divest-anyvision-stake-end-face-recognition-investing-idUSKBN21E3BA.

[4]   K. Walker, "An external advisory council to help advance the responsible development of AI," 26 March 2019. [Online]. Available: https://www.blog.google/technology/ai/external-advisory-council-help-advance-responsible-development-ai.

[5]   S. Levin, "Google scraps AI ethics council after backlash: 'Back to the drawing board'," 5 April 2019. [Online]. Available: https://www.theguardian.com/technology/2019/apr/04/google-ai-ethics-council-backlash.

[6]   Infocomm Media Development Authority, "Artificial Intelligence - Infocomm Media Development Authority," [Online]. Available: https://www.imda.gov.sg/infocomm-media-landscape/SGDigital/tech-pillars/Artificial-Intelligence.

[7]   Microsoft, "Microsoft – IDC Study: Artificial Intelligence to nearly double the rate of innovation in Singapore by 2021," [Online]. Available: https://news.microsoft.com/en-sg/2019/02/20/microsoft-idc-study-artificial-intelligence-to-nearly-double-the-rate-of-innovation-in-singapore-by-2021/.

[8]   M. H. Fitzgerald and P. A. Phillips, "Centralized and non-centralized ethics review: a five nation study," 2006.

[9]   D. Gelles, N. Kitroeff, J. Nicas and R. R. Ruiz, "Boeing Was 'Go, Go, Go' to Beat Airbus With the 737 Max," 23 March 2019. [Online]. Available: https://www.nytimes.com/2019/03/23/business/boeing-737-max-crash.html.

[10]   Unilever, "Leading safety and environmental sustainability sciences," [Online]. Available: https://www.unilever.com/about/innovation/safety-and-environment/leading-safety-and-environmental-sustainability-sciences/.

[11]   Task, "How to Use the Risk Assessment Matrix in Project Management?," [Online]. Available: https://www.ntaskmanager.com/blog/risk-assessment-matrix/.

[12]   O. C. Ferrell and J. Fraedrich, Business Ethics: Ethical Decision Making & Cases, 2014.

[13]   "Empowering AI Leadership," [Online]. Available: https://www.weforum.org/projects/ai-board-leadership-toolkit.

[14] Wikipedia, "Compliance and ethics program," [Online]. Available:

https://en.wikipedia.org/wiki/Compliance_and_ethics_program#:~:text=On_May_1,_1

991,_as,organizations_convicted_of_federal_crimes..

[15] T. SEGAL, "Enron Scandal: The Fall of a Wall Street Darling," 22 September 2020.

[Online]. Available: https://www.investopedia.com/updates/enron-scandal-summary/.

[16] W. Kenton, "Sarbanes-Oxley (SOX) Act of 2002," [Online]. Available:

https://www.investopedia.com/terms/s/sarbanesoxleyact.asp.

[17] A. Hayes, "Dodd-Frank Wall Street Reform and Consumer Protection Act," [Online].

Available: https://www.investopedia.com/terms/d/dodd-frank-financial-regulatory-

reform-bill.asp.

[18] B. Ammanath, D. Jarvis and S. Hupfer, "Thriving in the era of pervasive AI,"

[Online]. Available: https://www2.deloitte.com/insights/stateofai.

[19] Insights Team, "Rise Of The Chief Ethics Officer," 27 March 2019. [Online].

Available: https://www.forbes.com/sites/insights-intelai/2019/03/27/rise-of-the-chief-

ethics-officer/?sh=5ef753e95aba.

[20] T. O'Brien, S. Sweetman, N. Crampton and V. Veeraraghavan, "A model for ethical

AI | World Economic Forum," 14 Jan 2020. [Online]. Available:

https://www.weforum.org/agenda/2020/01/tech-companies-ethics-responsible-ai-

microsoft/.