IoD

AI in the Boardroom
The essential questions for your next board meeting

IoD

This paper was produced by the Institute of Directors Science, Innovation and Technology Expert Advisory Group with the aim of helping members to understand AI.

With thanks to the members of the Expert Advisory Group:

- Michael Ambjorn, Managing Director PropTech Denmark
- Dr Phil Clare, CEO Queen Mary Innovation
- Paul Corcoran, Managing Director Interimconsult
- Dr Jonathan R. Everhart, CEO & Chief Investment Officer - Global ReEnergy Holdings
- Faisal Khan Chair, Institute of Directors South
- Rob Noble, Chair of Forestreet, The Webinar Vet, MBA-IT & Fidem Consulting
- Pauline Norstrom*, CEO Anekanta Consulting
- Yinka Olarinde, Principal Consultant, Strategy & Technology, Siep Global Consulting Ltd
- Professor Rameen Shakur, Founder and Chairman Cambridge Heartwear Ltd, Cambridge Science Park

Please note, this paper reflects the views and opinions of the members of the Expert Advisory Group, not the IoD itself.

Adapted from an original framework published by Anekanta Consulting 5 February 2020
Author: Pauline Norstrom

AI in the Boardroom
The essential questions for your next board meeting

Index

Executive summary

- An IoD members' survey revealed in 2022 that 80% of boards did not have a process in place to audit their use of AI. They said they did not know what questions to ask.

- Board governance of AI at the project inception is important because boards do not wish to unravel ethical issues later which may have a reputation and cost impact.

- Research reveals a gap between board governance and the use of AI in their businesses. Over 86% of businesses already use some form of AI without the board being aware of this.

- AI applications in business are expanding so AI opportunities and threats need to be understood.

- AI can amplify existing bias in human decisions and safeguards are needed to prevent this and its impact on perpetuating bias in existing culture

- AI governance may be rooted in the core ethical values of the business.

- AI risk and governance model requires a framework which boards can use as a blueprint.

A reflective checklist

This reflective checklist will help you get a board-level understanding of where you stand as an organisation when it comes to the ethical use of Artificial Intelligence (AI). It draws on a set of 12 principles first developed by Pauline Norstrom of Anekanta Consulting that can help guide the use of AI throughout an organisation.

It is clear that AI needs to be on the board agenda and considered seriously as part of the G in ESG (Environmental, Social and Governance) and the CSR (Corporate Social Responsibility) requirements.

AI should not be confined within the realms of IT, although the CIO may take responsibility for implementation and management.

The 12 principles are:

1. Monitor the evolving regulatory environment.

2. Continually audit and measure what AI is in use and what they are doing.

3. Undertake impact assessments which consider the business and the wider stakeholder community.

4. Establish board accountability.

5. Set high level goals for the business aligned with its values.

6. Empowering a diverse, cross functional ethics committee that has the power to veto.

7. Document and secure data sources.

8. Train people to get the best out of AI and to interpret the results.

9. Comply with privacy requirements

10. Comply with secure by design requirements.

11. Test and remove from use if bias and other impacts are discovered.

12. Review regularly.

Monitor the evolving regulatory environment

On 21 April 2021, The European Commission published their proposal for the regulation of Artificial Intelligence[1]. The new regulations which will automatically become law in all EU member states, set out a risk based approach to the implementation of AI horizontally across sectors, also vertically, providing sector specific guidance for AI in law enforcement, education, safety systems and so on. The impacts are going to be wide reaching. The responsibility and accountability is clearly defined through the entire supply chain. The developer, provider and user are all accountable for the safe implementation of AI.

Before AI is developed or deployed, a risk assessment must be carried out to determine what level of transparency and ex. ante, CE Marking and post market monitoring is required.

Furthermore, the EU AI Liability Directive complements and modernises the EU civil liability framework, introducing for the first time rules specific to damages caused by AI systems[2].

In the UK, a Government white paper on AI regulation pursuant to the policy paper[3] is due to be published in spring 2023 which is expected to set out a risk based approach based on the characteristics of the AI technology,

e.g., whether it can act autonomously, or learn and change in its operating environment. These are examples of regulatory developments that may affect your operations. Expect additional jurisdictions to catch up.

- What AI risk assessment tools do you use already?

- Is the business open and engaged with key stakeholders around ethics and technology?

- Where the business is regulated, to what extent is it engaged in discussions with regulators about any changing requirements?

- Is there recognition of the potential for the role for regulatory sandboxing e.g. via UK Financial Conduct Authority and Information Commissioner's Office?

---

1   https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A52021PC0206

2   https://ec.europa.eu/commission/presscorner/detail/en/QANDA_22_5793

3   https://www.gov.uk/government/publications/establishing-a-pro-innovation-approach-to-regulating-ai/establishing-a-pro-innovation-approach-to-regulating-ai-policy-statement

Continually audit and measure what AIs are in use and what they are doing

The ethical principles must be auditable and measurable; they should be embodied in the ISO 9001:2015 quality system (or equivalent suitable system for example ISO/IEC 42001 when ratified) to ensure a consistent approach to the evaluation and use of AI by the organisation.

- What is your risk appetite when it comes to AI? And are AI application risks built into your Risk Register?

- Does your board audit committee (or equivalent committee) have an active role already?

- Are your ethical principles human readable, and understandable at board level? Plain language helps.

- Are the principles machine readable so they can be read by the code that needs to abide by them?

- Are the principles built into your ISO 9001:2015 or equivalent quality system?

- Will you want to report on AI in your annual report?

- If yes, consider your timeline and whether external advisory assistance is needed.

Undertake impact assessments which consider the business and the wider stakeholder community

Impact assessments must be undertaken which consider the possible negative effects and outcomes for employees who interact with the AI or whose jobs may be affected. The approach to the change of any job's purpose must be in line with the organisation's policy and compliant with the relevant employment, privacy, and equality laws.

The employees must be given sufficient information about the potential use of the AI and the opportunity to provide feedback to the board via the ethics committee. The employees must be advised so they are aware of what processes are being undertaken by AI even if they do not directly interact with it e.g., they must be informed and aware when they are interacting with an HR chatbot.

Similarly, impact assessments must be undertaken for all stakeholder groups including customers, suppliers, partners and shareholders.

- Is all AI use appropriately labelled throughout the organisation?

- Is there a structured, continuous method of evaluating impact – perhaps through an independent third party?

- What level of explainability is available for AI-led decisions?

Establish board accountability

The board is accountable both legally and ethically for the positive use of AI within the organisation including third party products which may embed AI technologies. Ethical principles, policies checks and measures must be in place and rigorously monitored. This is to ensure that the AI is safe, transparent, does not have a negative effect on the well-being of employees and stakeholders, and prevents the introduction of bias into automated processes and decisions. The board holds the final veto on the implementation and use of AI in the organisation.

- Does the board have the right skills and knowledge to consider the risks and issues?

- Does it understand how data, algorithms, and other technologies are being used in the business, including third party products which embed AI technologies, especially to make key decisions or redirections?

- How is ethics around technology included in board governance?

- How is the board communicating the importance of an ethical approach to technology adoption?

- Who is accountable at board level for these issues?

- Are you aware if your organisation is using Machine Learning (ML) or Artificial Intelligence (AI)?

- If the answer is yes, then this list is for you. Carry on.

- If the answer is you are not sure, then now is the time to check. Ignorance is said to be bliss, but not in the boardroom.

- And if you are not using it directly, you might be relying on AI indirectly through clients, suppliers or perhaps sub-contractors. Going through this document will help give you a better understanding if you should start digging.

Set high level goals for AI in the business aligned with its value

High level goals for the use of AI in the organisation must be created in line with its vision, mission and values. The AI must be to the benefit of the organisation, its employees, customers and wider stakeholders.

Examples of the goals of AI include:

1. Augmenting human tasks

2. Enabling better, consistent, faster human decisions

3. Preventing bias

4. Not causing harm to any employee or stakeholder

- Are the goals clearly written in plain language and understandable from the board to the frontline?

- And where relevant, are they machine readable as well, for any code that relies on them?

- Are they aligned with the organisation's vision, mission and values?

- Is the benefit measurable?

Empowering a diverse, cross functional ethics committee that has the power to veto

The ethics committee which includes employee representatives should oversee AI proposals and implementations. Proposals should include a risk assessment, impact assessment and evaluation of the transparency of the AI's decision-making process.

The committee must recommend to the board whether the AI implementation may have a beneficial effect and understand potential negative impacts. The ethics committee must examine the risk of bias becoming embodied in the AI and propose measures to prevent and remove bias.

- Does the ethics committee have clear Terms of Reference?

- Is the committee made up of a diverse group of employees?

- Have the members of the committee received adequate training?

- Does the ethics committee have the power to veto harmful use pending investigation?

Document and secure data sources
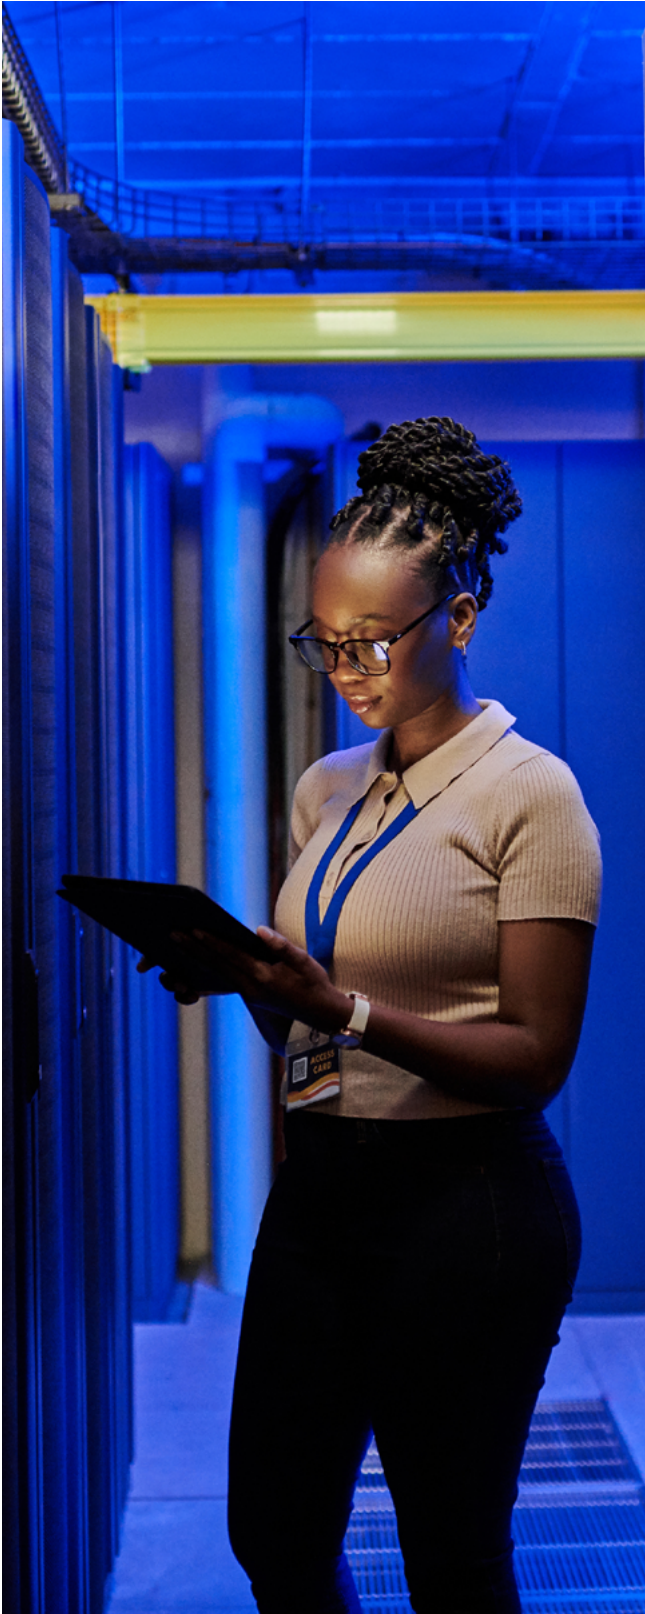
In the definition of the purpose of the specific AI implementation, the sources of data must be identified and documented. There must be a method of detecting and reporting bias. If bias is discovered, action must be taken to identify the source, discover why it is there and remove it from the AI. The data sources must be preserved and any adjustments documented for traceability. KPIs (Key Performance Indicators) must be implemented to keep the bias out of the organisation and out of its data sets.

- Is the decision logic understandable by human beings, including board members? If not, are you in black box territory which means unquantified risk?

- Are tripwires incorporated in case the AI goes off track?

- Are logs secured and regularly backed up?

- Is ethics considered when reviewing or signing off new AI projects or use cases?

# Train people to get the best out of AI and to interpret the results

Training needs of employees must be taken into account and programmes must be implemented in the company in order to prevent bias and potential harmful outcomes. The employees must be trained on the system used to monitor and report bias.

- Has adequate, accessible and audience-appropriate training been designed and made available?

- Have staff been given adequate time and incentives to undertake and maintain training?

- How are new staff taught about the ethical values of the business?

- Are staff aware of issues surrounding AI applications and hidden bias?

Comply with privacy by design requirements

The AI must be designed for privacy and audited to ensure compliance with data privacy legislation such as GDPR (General Data Protection Regulation) and organisation policies. The technical teams responsible for implementation must be trained on the ethics framework and to challenge the AI developers to ensure that the transparency of the AI decision-making is understood and whether any human final decision-making is needed or not. The AI technical teams must report their findings to the ethics committee as part of the AI proposal process.

- Are there mechanisms for employees to raise concerns about ethical questions, such as whistle-blowing processes?

- Is the reporting process clearly understood by everybody?
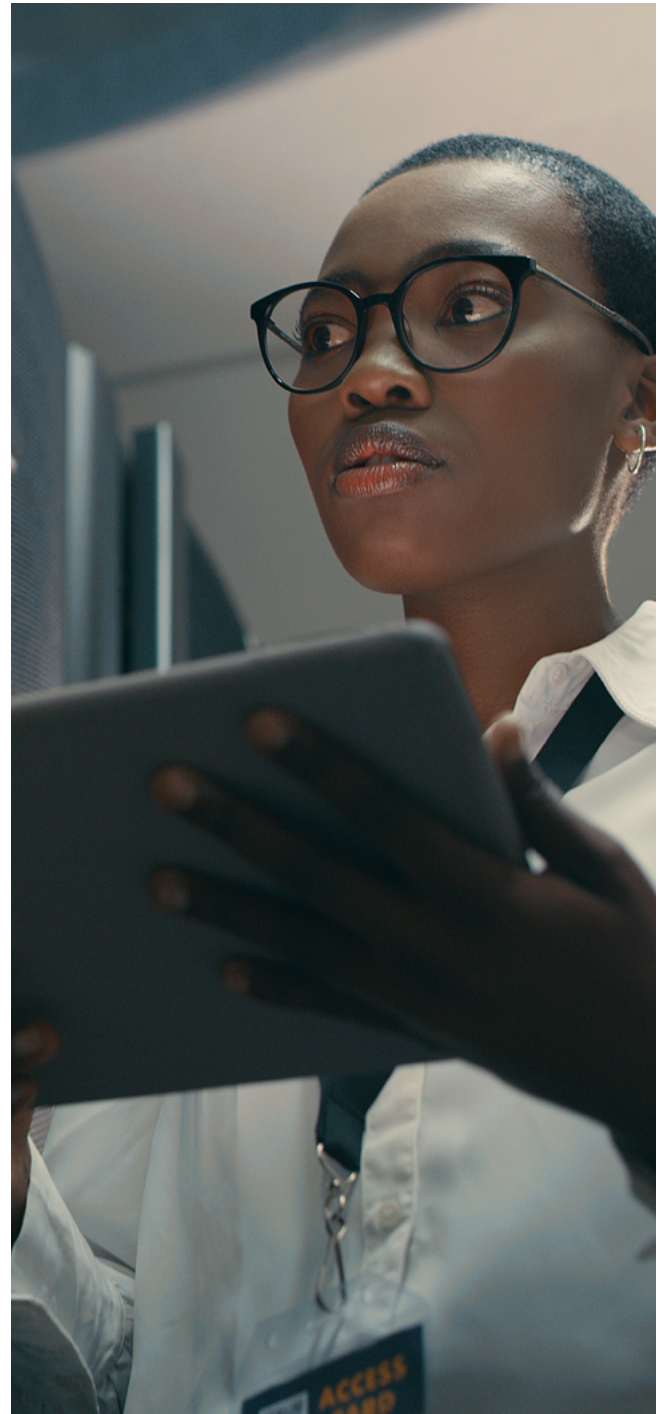
Comply with secure by design requirements

The AI must be secure by design and stand the scrutiny of external test and certification processes such as Cyber Essentials Plus. This may include penetration testing which may be used to ensure that the data sets used in the AI cannot be breached.

- To what extent should individual AI system designers and engineers within the business be explicitly required to declare their adherence to a set of ethical standards? This is particularly relevant to AI developers.

- If you have recently run a penetration test, what did you learn and what action will you take?

- Are the implications of future AI system learning and reconfiguration clear?

Test and remove from use if bias and other impacts are discovered

The AI must be tested prior to any organisation-wide implementation to ensure compliance with the ethics framework and the scope of the implementation. The ethics committee makes recommendations to the board to proceed with full implementation, adjust the scope or to veto if conditions are not met.

The decision to implement rests with the board as does accountability for ongoing safe and consistent performance which does not deviate from the intended purpose.

- If AI solutions are externally sourced, are there ethical requirements engrained in procurement processes?

Review regularly

Regular reviews must be undertaken to monitor the decisions of the AI and audit them against the intended use and ethical framework.

If the AI deviates from the purpose and ethics in any way, the deviations must be documented, reported to the ethics committee and corrective actions implemented within a reasonable period since discovery depending on the severity of the issue based on the risk and impact assessment.

- Where there is automated decision-making, to what extent have controls been reviewed to ensure there is a human in the loop?

- Is the board satisfied it has the necessary diversity and inclusion in the AI expert groups with different perspectives when developing technology which enables them to spot problems of bias in training data and decision-making?

- Has the board, in seeking assurance on the standards for training, testing and operation of AI solutions, considered relevant tools such as:

  - Algorithmic impact assessment

  - Ethical audit mechanisms

  - Consumer assurance through Kitemarking or

  - Standards for ethically aligned design

About the Science, Innovation and Technology EAG

The Science, Innovation and Technology Expert Advisory Group is one of a number of Expert Advisory Groups that the IoD Governance and Policy Unit has established to help tap into the expertise of IoD members on the key issues for UK directors, providing insight from those who have substantial front-line experience.