

The Board of Directors and Artificial Intelligence – An exciting dilemma

Victor Prozesky¹ and Jan Veldsink²

Introduction

While the digital transition has been high on boards' agendas, the rapid coming of age of artificial intelligence (AI) recently caused boards to reflect on how to approach this exciting, yet double-edged sword and use it to optimal advantage.

AI is fast becoming an element in the requirement for boards to have adequate digital literacy in its composition. It also is a further aspect to consider on the journey to digital transformation, which is on practically every organisation's roadmap. There is general acknowledgement that this transformation has significant consequences for business models, customer expectations, the competitive landscape and the regulatory environment. It is therefore imperative for all boards to exercise careful oversight over the identification and evaluation of opportunities and risks of adapting new technologies such as AI, cloud computing, cyber security and blockchain.

One of the often-neglected aspects of the approach to the digital journey is the highly important people side of any such transition; this includes cultural aspects, agility and speed in the organisation, accountability and questions such as the importance of the 'why' as much as the 'what'.

AI and the board itself

There is a strong possibility that AI may change the work of the board quite fundamentally. For example, it is not out of the question that an AI device may in future become a board member. This is already in play, for instance the

management team of the Scandinavian technology company Tieto has a fully-fledged AI member, called Alicia T (Finextra, 18 October 2016) since 2016. In 2014, a Hong Kong venture capital firm made headlines when it announced that it appointed an AI unit to its board of directors (Kamalnath, 17 January 2020). This was a machine learning algorithm (VITAL) allowed to vote at board meetings based on processing a range of data, but it is generally accepted that the algorithm was essentially providing a suggestion upon which the board could vote. Irrespective, this has caused boards to take notice of the possibility.

Currently most corporate law systems require that members of boards of directors are natural persons, although in several OECD countries legislators are considering proposals for enabling non-natural persons to be appointed (Hijink, 2019/3). The two fundamental questions regarding a non-human to be appointed as board member are those of legal status and accountability. Current EU proposals are focused on the possibility of assigning legal person (or entity) status to an AI device. In this process the implications, including accountability, needs to be considered. Questions include issues such as where legal and financial liability will be assigned – for instance, to the manufacturer or with the machine learning structure.

The most basic issue to consider is a future where the capabilities of AI will have become immensely better than today's generative AI tools, and what the 'unique' competencies of humans around the boardroom table will be. These might include strategic consideration, ethical behaviour and people-focused thinking. The above assumes that these complex cognitive concepts may not be fully addressed by AI at such time; or it may be that

¹ The Board Practice - a leading global provider of board consulting services

² Nyenrode University – Education of Executives, MBA and board programs

human Directors will to some extent be on boards to provide oversight over the output from non-human board members (or - advisors), rather than traditional oversight over the management of the company.

Ethical & moral dilemmas

Two dilemmas, one ethical and the other moral in nature, entered as input to ChatGPT provide an example of the level of intelligence currently provided by publicly available non-trained generative AI tools. In the ethical question the board of a pharmaceutical company is informed that the active ingredient of its star pain medication, which is highly popular and has grown to be the biggest revenue earner for the business, is not effective. It transpired from additional clinical trials that the desired effect was caused by an excipient. The question asked was whether the board should react to this discovery, given that the drug was having the proclaimed effect. The answer from the AI tool was unequivocal for the board to report this finding, in doing so considering open and honest reporting, transparency and accountability, patient safety, regulatory compliance and stakeholder considerations. The tool went even further and recommended how this should be dealt with.

In the moral question, the board of a medical facility is informed of five terminally ill patients, each needing a different organ in order to survive. There are no organs available to perform any of the transplants. (Sylvia, 2012) It just so happens that there is a sixth patient, suffering from a fatal illness, who will most likely pass away sooner than the other five if not treated. With the death of this sixth patient, the organs can be used to save the other five patients. However, medication is available which can delay the death of the sixth patient to a few days after the five patients will have died. The question to the tool was what the responsible medical professional was to do.

Here the tool was advanced enough not to propose a single decision. As a matter of fact, it proposed that, with different perspectives, different decisions may be more appropriate. From a utilitarian perspective the medical professional may consider saving the maximum number of lives by administering the medicine to the sixth patient, hoping that the chain of events will also allow the other five patients to receive the needed organs and potentially survive. From an individual rights perspective the rights and autonomy of each individual patient makes it ethically problematic to intentionally sacrifice one patient for the

sake of others. In considering proportional sacrifice a morally acceptable action may involve sacrificing one life to save a greater number of lives. The recommendations further urge consideration of the legal and ethical framework of the specific jurisdiction and medical facility in which this scenario takes place.

These two examples indicate that tools that are not yet optimised or trained are already at a level of being cognisant of tough ethical and moral dilemmas and are already far more useful than merely advanced fact checking, data processing and analysis.

AI and Board oversight

There are two major components to the oversight of the opportunities offered and risks posed by AI: the work of the board itself and the application of AI in the organisation. In addition to this the complexity of integrating AI into the fibre of organisations will not be less complex than the digital transition journey of companies. Most companies found this transition to be much more onerous and costly than anticipated, there is no reason to believe that an AI journey will be any different.

For the board itself there are major opportunities to focus more on strategic and people considerations and optimising administrative, compliance and structural risk processes of the board. For example, available general tools are capable of summarising a thousand-page board pack in thirty pages, including highlighting the most important number of issues identified for board members to consider. Board effectiveness evaluation tools where the outcomes of internal questionnaires can be automatically analysed, resulting in detailed reports with significant insights, are in an advanced stage of development.

For meetings itself current popular online meeting tools allow recording and transcription of meetings, including meeting minutes and action items. While these are only a few of the host of opportunities offered by AI, the risks are also significant. For instance, an analysis of how and where data is processed when popular online meeting tools analyse recordings does not give a clear indication that the data is treated within the closed data areas of any particular client and indications are that it may even be processed outside the jurisdiction of the organisation.

In general, the dilemma is that these AI tools are increasing

in capability because they are continuously exposed to new information and training through the use of data provided, but that is precisely the issue; the data of the provider may become part of the database. This concern caused Amazon to ban the use of AI tools in software development as it was clear that, although exceptionally effective in generating code, these tools incorporate the new development into its arsenal. Particularly, if the original code contained some proprietary information, this may lead to intellectual property issues.

Another complication in the development of these tools is the fact that they are created for the end-user, and it is in the interest of the developer of the tool that the user is satisfied with the outcomes offered, so it may be programmed to give a desired outcome, especially if it has some information on the preferences and beliefs of the user. Most information supplied to broad-based tools also originated from documents from the western world. The danger of inherent bias is thus significant.

Similarly, models will consider new input as being part of its 'truth' until convinced otherwise. For example, if the simple addition of seven plus three is entered into one of the popular tools with the proposed answer as thirteen, the model will correct it with the right response of ten. Repeating it a few more times will result in the same answer, but in a surprisingly short repetition of a proposed wrong answer the tool will be convinced that the alternative answer may be correct. Pursuing this approach may lead to problematic consequences; if automated algorithms enter dubious information often enough, misinformation may be viewed as factual truth. And because the models are 'pleasing' in nature it will try to generate an answer, even though it does not have the body of knowledge to support it. The ability to reason currently lacks in these tools and models are trained on statistics of a language, not on delivering the right answers to a question per se.

The development of these tools based on models will consequently delight actors with alternative motives, such as cyber criminals, data poachers and phishers. The voice of an executive in a voice note to staff was already perfectly imitated and an email written in the style of a particular person, seemingly coming from the right source and address. It may well be that the concepts of truth and accuracy, already under fire in a world having to deal with 'fake truths', will become one of the major areas of attention for boards and organisations.

Therefore, one of the major roles of the board will become one of ensuring that opinions, recommendations and proposed actions from automated algorithms are based on the right knowledge with an objective approach and using the actual information and data at hand. The second area of focus will be to ensure that ownership of data and information remain enclosed in domains, inter alia due to privacy and other legal requirements.

The ethical use of AI will also need some careful consideration by the board of directors; for instance, it is already possible to obtain algorithms analysing facial expressions and verbal cues to determine whether someone is telling the truth or not. Imagine a world where executives addressing the board are put under such scrutiny. The question will become one of deciding what the ethical and moral applications of AI are and to steer away from fascinating, but potentially dubious uses.

The use of AI models in the workplace is exploding and it is no surprise that the EU, for instance, has been developing legislation for the better part of the last decade. At the same time organisations should, with careful oversight, take care on what and how such tools are used in the organisation to mitigate risks and protect reputations.

A recent McKinsey survey (Matt Banholzer, 2023) found that 30% of companies with innovative cultures are already applying generative AI at scale in their organisations, versus 4% in more conservative environments. At the same time an astounding 78% of respondents believe that current business models need to change moderately to completely by 2025 to remain economically viable. In a similar study (Singla, 2023) McKinsey found that, less than a year after most of the generative AI tools were released, a third of companies are using these tools formally in at least one business function, with only 32% of respondents managing inaccuracy, the most important perceived risk. Only 21% of companies have policies governing the use of generative AI technologies by employees. Given the heightened risks of these freely available tools boards will need to play a significant role in ensuring that management build the necessary frameworks and governance to balance the risks.

The biggest challenge for boards will be to ensure that innovation and agility is built into the DNA of the company to be able to exploit the opportunities offered, while

simultaneously ensuring that risks are mitigated. Boards need to be able to ask the right questions and value the answers. The board also needs to be cognisant that new business models may develop and new ways of organising, as well as controlling the organisation, may emerge along the way.

Do boards have the skills and understanding to provide adequate oversight?

The Board Practice conducted an analysis of the composition of the boards of the top 40 companies on the Amsterdam stock exchange through a study of the published skills and experience levels of members of the board in the annual reports of the relevant companies. The aim was to find the level of particular technology related skills in each sector, but also to probe whether boards have started to publish increasingly important skills, such as AI. Therefore, the lack of reporting such skill may indeed not reflect a lack of the particular skill, just that it

has not been included to that level of detail.

The answer represented in Figures 1 and 2 is quite surprising in that some technology-specific competencies are quite well represented, such as industry-specific technology, general IT, cyber-security and digital skills, which are represented in between 8% and 17% of the total board membership. This indicates the increased importance of skills recently elevated to Board level, given the increased risks associated with technology, such as cyber security. Some other areas with increased focus, such as AI, are not well represented at all. Only three Board members out of a total complement of 341 (less than 1%) are described as having AI skills, with IT strategy / digitalisation skills at double that level of representation. No Board member from the Top 40 companies is described as having digital marketing type skills. It is also clear that the gender disparity is decreasing; out of the total number of Board members, 60% are male, 40% female. Of these ³, 42% of males and 26% of females contribute to one or more technology related skills.




Figure 1: Overall % of Board members contributing to technology skills and experience

³ Each representation of a skill is indicated as a full Board member, naturally some Board members represent more than

one technology related skill, therefore this does not represent the number of persons as Board members in percentages of the total.




Figure 2: % of Board members in each industry contributing to technology skills and experience

Figure 2 clearly indicates that the technology skills represented on boards is centred around general relevant experience in the sector of the company, digital know-how and IT skills. The only two industries with broad technology representation on its boards are Technology and Banking & Finance. It is quite surprising that in a world dominated by massive amounts of information and data, there are six board members in total, representing less than 2% of members, described as having data & analysis skills.

Conclusion

It is early in the cycle of the increased opportunities and risks offered by new developments in technology, including the focus on AI and the need for boards to be able to offer the necessary oversight and ability to use AI in the board room at optimal levels. Even though generative AI has been available as mainstream technology for less than a year, companies are racing to adopt it at scale to obtain early advantage in increases in cost-effectiveness, improved efficiencies

and increased competitiveness in early implementation. This is happening across sectors and geographies and practically all business areas, including procurement, supply chain, design, manufacturing, software development and data analysis. As is the case with most emerging technologies, the focus is on finding applications in organisations, with less focus on taking care of risks and ensuring proper processes and systems to ensure adequate control and an understanding of how to deal with aspects such as accuracy and bias.

It is apparent that boards are not at the same level of skills and experience to provide proper oversight over this fast-developing technology with the potential to fundamentally impact companies, even in the near-future. If boards are appointing such skills this is not indicated in public descriptions of the skills and experience of members.

Current levels of understanding of AI at board level will need to develop at speed, given its rapidly increasing

application. This may be in the form of appropriate new board members, or the appointment of technology councils or forums which deliberate under the auspices of and reporting to the board and management. Another way of increased application and understanding at board level is the potential of applying AI technology in the boardroom. This may be through using AI tools at board level, or the appointment of AI personas as board members once the technology and regulatory frameworks are adequately developed.

Imagine a board of directors in the not-too-distant future

where a large proportion of board members are automated devices that provide expertise such as data analysis, scenario planning and fact checking, with the human members providing 'human' contributions, such as sense-checking, morality, ethics and empathy. It seems that the question of what makes us truly human and cannot be replicated by technology will indeed be answered in the next decade or so.

The publication was written by human beings, AI was used as tool for research purposes.

References

- Finextra. (18 October 2016). *Tieto appoints bot to leadership team*.
- Hijink, P. m. (2019/3). Robots in de boardroom. *Ondernemingsrecht 2019/3*, 10.
- Kamalnath, A. (17 January 2020). *THE PERENNIAL QUEST FOR BOARD INDEPENDENCE: ARTIFICIAL INTELLIGENCE TO THE RESCUE?* Albany Law Review.
- Matt Banholzer, B. F. (2023). Companies with innovative cultures have a big edge with generative AI. *McKinsey & Company Strategy*.
- Singla, A. (2023). The state of AI in 2023: Generative AI's breakout year. *McKinsey survey*.
- Sylvia, J. I. (2012). Moral Monday: The Patient. *Philosophy Matters*.