A responsible approach to using AI in corporate reporting

Guidance for Boards and management on approach and disclosure

November 2023

Falcon Windsor

Contents

Introduction

Truth and accountability are the bedrock of corporate reporting, and thus the bedrock of our global system of capital markets. For markets to work, investors, regulators, governments and wider stakeholders must be able to trust the information on which they are based.

Early in 2023, when large language model artificial intelligence (AI) systems, or LLMs, really started to get traction with ChatGPT, I realised that the wholesale use of AI would raise issues for the accuracy of corporate reporting in general; and for directors' duties in terms of ensuring that reporting is fair, balanced and understandable. Then, in March 2023, the UK Government launched a broad consultation on regulating the use of AI. With the help of some other reporting experts, I submitted a proposal for regulating its use in reporting.

Over the summer and into the autumn, I ran a series of one-hour focus groups to gain feedback and input into the proposal from report preparers. These groups included FTSE 100, FTSE 250 and other UK companies, along with various interested parties such as auditors. Finally I held an investor focus group to gain feedback from the investor community.

In summary, while there were varying views on the merits of regulation, all felt that guidance on the use of AI in reporting is essential. Those who supported a regulatory approach felt that it might nonetheless be both challenging and too slow to materialise, given how quickly various forms of AI are being adopted by companies.

This guidance paper is the output. It is intended to help Boards and management think through the challenges and benefits of using various forms of AI in reporting, and how to disclose its use within the annual report. With the UK Government's intention to be global leaders in the regulation of AI, this paper also aims to help our regulators and inform their thinking when developing formal guidance. While this project has focused on the UK, the principles and approach should be useful for reporting in other jurisdictions too.

I am keen to hear feedback on this guidance from all stakeholders in reporting, and particularly from companies as they work on reporting over the coming year, so that I can update it and keep it current with the status of AI usage. So please do send me your thoughts and views.

Claire Bodanis

Founder and Director, Falcon Windsor
Author of *Trust me, I'm listed*

claire@falconwindsor.com

+44 7966 196808

Meaning of 'reporting': this paper uses reporting to mean any statements to the market on which investors and other stakeholders rely to make decisions, such as the annual and related reports, results statements, trading updates and so on.

About Claire Bodanis

Claire is one of the UK's leading authorities on corporate reporting, and the founder and director of Falcon Windsor (FW), a specialist reporting and communications advisory company, organised as a network of reporting and communications experts who come together to advise companies on how to report well. What sets FW apart is Claire's philosophy: it is only by understanding *why* we are required to report that we can navigate the complexities of regulation and do the difficult job of reporting well. This philosophy led to Claire being asked by the UK Chartered Governance Institute to write a guide to reporting for UK plc. With a foreword by Sir Donald Brydon, and contributions from experts across the reporting world, *Trust me, I'm listed – why the annual report matters and how to do it well*, was published (second edition) in October 2021.

How to approach using AI in reporting

The guidance on the following pages has been developed to address the views and concerns of the focus groups which are summarised on pages 6-7. The most important thing is to think through any use of AI, particularly large language model systems (LLMs), ahead of using it in reporting. The questions set out here should help Boards and management do just that. Following the questions are some practical steps to help bring the use of AI into the reporting process in a thoughtful way.

In this section:

Questions to help explore the use of AI in reporting

Status of AI within the company and how it affects reporting

Given how fast this agenda is moving, and that AI is currently unregulated, it's important that Boards and management stay up to date with what is going on inside the company, so that they can be sure how any use of AI will ultimately affect reporting. Consider these questions:

1. How is AI being used inside our company right now?
2. What is the plan for AI usage, and what's being developed?
3. Where does/would the use of AI touch our reporting process?

Benefits and drawbacks of using AI

There will be trade-offs in using different forms of AI in reporting, so it's important to explore these trade-offs before making a decision about if/how to use them. Consider these questions:

1. How would the introduction of AI support or detract from us achieving reporting's ultimate purpose, namely to build a relationship of trust with investors and other stakeholders through truthful, accurate, clear reporting that people believe because it tells an honest, engaging story?
   a. **Truth**: how will use of AI support or detract from the truthfulness and accuracy of our reporting?
   b. **Relationships**: how will use of AI support or detract from our ability to build a relationship of trust with investors and other stakeholders?
2. Which types of AI system could help us achieve reporting's purpose, and which could hinder that goal?
3. How will we verify information produced by AI?

> "It wouldn't be advisable to use AI for forward-looking statements, but it could have a use in highly structured sections."
>
> Investor

A beneficial use of AI: smooth handling of enormous data sets

While LLMs carry significant risks in terms of writing report content, the types of AI that can smoothly handle enormous data sets could be very useful in finding patterns and summarising from source material which is then used in the annual report. Any such use should be disclosed as discussed on page 5.

How to approach using AI in reporting continued

Use of LLMs

Large language model systems (LLMs) of the ChatGPT type raise particular problems for reporting, given they are known to create false yet highly plausible narrative, including false sources. They also have their own in-built biases. If Boards and management are thinking about using LLMs, consider these questions:

1. Which sections might LLMs be able to help us with?
2. Which sections should be free from LLM usage?
3. How will we ensure that:
   - a. The questions we ask of an LLM are the right ones?
   - b. Any narrative written by an LLM is accurate?
   - c. Any sources cited by an LLM are genuine?
   - d. The output is not biased?
   - e. The people checking the output are qualified to do so?
   - f. The company brand and voice remains distinctive?
   - g. That directors can sign off the output as fair, balanced and understandable, or otherwise accurate?

Longer-term questions

If Boards and management decide to use LLMs in reporting, consider how to deal with the following longer-term issues:

1. **Critical thinking**: If the thinking is done by machine, how will we make sure that people still have the ability to exercise judgement and critical thinking?
2. **Ability to check output and confirm it is accurate**: Related to that, how will we make sure that people continue to have the experience and expertise to check the output and confirm it is accurate?
3. **Learning by doing**: How will juniors build up the knowledge and experience they need to produce reporting well, if the tasks they may have done previously as an essential part of their personal development and training are replaced by machines?
4. **Effect on reporting of AI learning on AI**: As more and more information is created by AI, how will reporting be affected, once AI systems start to train themselves on their own output?

Practical steps

1. **Get the use of AI in reporting onto the Board's/Audit Committee's agenda**: Use of AI should be part of the directors' discussion on reporting.
2. **Convene a discussion with the risk function**: The risk team should be involved in exploring the use of AI in reporting through the questions listed above.
3. **Convene a discussion with the reporting team**: The reporting team should also be involved.
4. **Develop an internal policy on the use of AI in reporting**.
5. **Communicate the policy** to the whole company so that everyone involved in providing the information on which reporting relies is aware of the policy, and can ensure it is adhered to.
6. **Document the use of AI** across the whole reporting process so that it can be properly managed and disclosed.

LLMs – red flags

- Do not put confidential information into public tools such as ChatGPT
- Do not use LLMs to draft anything that is a matter of opinion, since that opinion needs to be the opinion of the Board and/or management
- Do not use LLMs for forward-looking statements

Best use of LLMs – a summarising tool

- But don't forget that all output must be checked by a human being for accuracy

"I don't think people in our company have thought about the confidentiality issue, and there hasn't been any training on AI. It would be a good idea!"

FTSE 100 report preparer

How to approach disclosure

Investors in particular want to know if the companies they invest in or are considering investing in are using AI to create reporting, and if so, how. Here is a model for how to disclose the use of AI in the annual report. Note that the recommendations for the location of the disclosure statement and cross-referencing from other sections are based on the structure of UK annual reports, but the principles should be useful for reporting in other jurisdictions.

"I'm not sure any director could sign off a report as fair, balanced and understandable if it was generated using an AI model – and I'd certainly want transparency about its use in reporting."

Investor

A good 'use of AI in reporting' disclosure statement should explain:

1. The company's policy on the use of AI in reporting.
2. The company's overall process for creating reporting, and how use of AI fits in.
3. The types of AI used in the reporting process.
4. How each type has been used within the process, including but not limited to:
   a. The source material on which reporting relies
   b. Reporting itself
   c. The verification process
   d. Any other processes that are part of creating reporting.
5. Which section(s) of the report, if any, have been created using AI, whether in full or as a drafting aid; and, in each instance, which type of AI has been used.
6. How the company has ensured that the information touched by AI is accurate and truthful.
7. Where relevant, what additional steps the company has taken to ensure that, in light of the use of AI, the report is fair, balanced and understandable, or otherwise accurate.

Where to include the 'use of AI in reporting' disclosure statement

The disclosure statement should be included in the governance report, wherever makes the most sense in terms of the type of company and its relevant disclosure requirements. Here are some suggestions for UK companies, based on where directors are most likely to discuss their responsibilities for accuracy and integrity of information. Non-UK reporters could apply this principle to their own disclosure framework.

For companies reporting under the UK Corporate Governance Code: within the fair, balanced and understandable assessment.

For companies reporting under the QCA Code: within Principle 4, which covers risk management.

For companies reporting under the Wates Principles: within Principle 3, director responsibilities, which covers accountability and integrity of information.

For companies that don't use any of these: it may make sense to include a section in the directors' report.

Cross-referencing to the 'use of AI in reporting' disclosure statement

Wherever the report discusses processes and procedures relating to the accuracy of reporting, and directors' duties in this regard, the use of AI should be mentioned, with a cross-reference to the disclosure statement. Sections likely to be affected include, but are not limited to:

Risk report (strategic report)

Risk management and controls (governance report)

Financial statements

Statement of directors' responsibilities.

Views and concerns

The views and concerns of companies and investors at a high level were very similar, while companies had additional points relating to the process of producing corporate reporting. The overall view was that, like it or not, AI will have an impact on reporting in some way; that it will happen very soon; and that guidance is urgently needed to ensure that the risks it presents are managed effectively, so that any usage is beneficial to the purpose of reporting. Many of the issues raised are of course not specific to the use of AI in reporting, but its potential impact is so significant that existing processes for dealing with these issues should be reviewed and strengthened.

> The views expressed in this report are from those who took part in the research, which included five investors, representatives from 18 companies (including 11 FTSE 100s), plus others from audit, governance/advisory, legal and academia. In total, 43 people took part.

"I'm sure it'll become possible to produce a report with AI, but what value would stakeholders see in a report written by a bot?"

FTSE 250 report preparer

Main concerns – use of LLMs

The main concerns came from the potential use of large language model systems (LLMs) like ChatGPT for drafting narrative. These centred around the accuracy of information; whether machine-generated information can genuinely represent the views of the directors; and the even greater importance of verification and assurance. There was also a general feeling that the use of AI in reporting is not on most Boards' agendas – but that it should be. Nonetheless, some felt that LLMs could have a place in making the drafting/editing process quicker and easier, if carefully controlled and restricted to certain sections.

All agreed that the use of AI in reporting should be disclosed in the annual report – a particular concern for investors.

Summary of key points raised

The rest of this section summarises the key points raised, grouped under the following headings:

- Status, knowledge and regulation of AI
- Risks of using large language model systems (LLMs)
- Potential benefits of using internal LLMs
- Views on disclosure
- Other general points.

These points informed the development of the guidance set out on pages 3-5. For all contributors, the focus of the discussion was LLMs. This is because they are the type of AI most likely to be used in reporting, given their perceived efficiency benefits, but which also present the greatest risks.

Status, knowledge and regulation of AI

Little knowledge about how AI is being used within companies: Many company representatives weren't aware of how AI is being used within their companies and therefore how it might affect reporting.

A very small number of companies are beginning to experiment with LLMs in reporting: While no companies that took part in the research have used LLMs yet in their reporting, some are experimenting with internal LLMs to see how they might help edit text, although at this stage their usage is minimal and restricted to editing/summarising information that has already been published.

Differentiating between types of AI: It's essential to be specific and consistent in the terminology we use to describe different types of AI. Different types have different benefits and risks.

Guidance is urgently needed: But views differ on what form it should take. Some investors felt that guidance would not be sufficient and regulation is needed; but most in the company focus groups felt that regulation would be both impractical and too slow, while some felt that regulation was in any case a step too far.

"The volume of sources [used by AI] might pose a real issue in unpicking what is fact and what isn't."

FTSE 100 report preparer

Views and concerns continued

Risks of using large language model systems (LLMs)

All agreed on these risks, although some were more concerned about them than others.

Accuracy and reliability of output – highly plausible yet false narrative/attributions: Unfettered use of LLMs – which are 'truth agnostic' – is likely to lead to errors in reporting.

Compromising confidentiality: Publicly available LLMs like ChatGPT should not be used for confidential information (a universal view); although internal LLMs, once created, could be useful.

Lack of authenticity of the company/individual director's voice: Many felt that if narrative were created by LLMs, it could not genuinely be considered the view of the company, or, in the case of authored statements, the individual director.

Heightened accountability of directors: The ultimate accountability of directors is even more critical when LLMs are involved. While many Boards are discussing how AI is used in products and services, it's less clear if directors are aware of the issues around reporting.

Potential implications for directors' insurance: There could be implications for directors' liability insurance if AI were used.

Reduction in critical thinking around the reporting process: When done properly, the reporting process requires senior management to come together to think, discuss and debate, which in itself is a valuable process that can lead to broader and better outcomes. There is a danger that this critical part of the reporting process becomes compromised or lost if reports were to be produced by LLMs.

Heightened importance of the human checker: All agreed that all reporting, whether produced by a human being or by some form of AI, should be checked by human beings. The particular challenge is that the more AI is used, the harder checking becomes, because a) AI will start learning from itself, making checking sources more difficult, and b) human beings will be less experienced at creating reporting themselves, and therefore less able to confirm the output is accurate.

Risks vary depending on how LLMs are used: While generally people had no issue with using an internal LLM to summarise information, they felt that this use was different in nature from using it to analyse information and/or draft narrative, which is far riskier.

Reporting could become pointless: Investors in particular were concerned about reporting becoming pointless if written by LLMs, since insight into the minds of management and the Board would be lost, and all reporting could end up sounding the same.

Potential benefits of using internal LLMs

While all agreed that external LLMs should not be used at all, some felt that even internal LLMs should not be used either, since the risks of introducing them were too great to make the potential benefits worth considering. The benefits discussed are listed here.

Time-saving: This was by far the main potential benefit, in that it could make the laborious process of drafting/editing easier, and automate tasks hitherto done by juniors, for example low-level data analysis.

Summarising information: The use of LLMs could help cut down the time/effort required to summarise information before drafting, although all agreed that even in this instance, the output must be checked by human beings.

Views on disclosure

Use of AI must be disclosed: Reporting should include a disclosure statement about the use of AI, setting out whether AI has been used and, if so, what types, how it has been used and in which sections.

Other general points

'That ship has sailed' – guidance/regulation must assume AI/LLMs will be in use inside companies: Despite the risks, many felt that LLMs may well come to be used anyway, so, rather than trying to prohibit their use, guidance/regulation should get companies to consider how LLMs are used and how to ensure that any usage is responsible.

Some parts of reporting should never be created by LLMs: All investors and most company preparers felt that forward-looking statements and opinion-based sections should not be created by LLMs.

Heightened importance of the human checker: All felt that any output must be checked and confirmed by human beings (i.e. AI should never be allowed to check and confirm reporting on humans' behalf).

Difficulties in checking LLM-generated content: Some company preparers felt that fact-checking LLM-generated content would be much more difficult since it would be unclear where the information came from.

Longer-term challenge for the human checker: Many agreed that in the longer term, there could be an issue for human beings' ability to review and confirm the information, since they may not have the experience or skills to judge what is and is not correct.

Longer-term challenge for training juniors: Despite the time-saving benefits of automating information, some raised concerns about how juniors would be able to build up the knowledge and experience they need to work in reporting, given the value of learning by doing, and the potential of AI to hinder people's ability to develop the capacity for critical thought.

New systems, processes and verification: Preparers felt that the use of AI (particularly LLMs) would require its own processes, systems and checks/balances, which themselves could be assessed, verified and reported on.

Appendices

Appendix 1 – Methodology

Process for developing the guidance

All dates are 2023

April/May – initial proposal developed

With the help of two reporting experts, I developed my initial proposal for regulating the use of AI in reporting with a view to submitting it to the Government's consultation, *A pro-innovation approach to AI regulation*

June – with advice from an MP, proposal submitted to the Government consultation

I sought advice from a sympathetic MP, who encouraged me to submit my proposal to the consultation

I revised my proposal and submitted it to the consultation on 20 June

July/August/October – six company focus groups held and feedback received

Outside academic research, focus groups are one of the most widely used qualitative research methods, and so I convened a series of focus groups with company executives directly involved in producing reporting and statements to the market, along with other interested parties including non-executive directors, auditors, the legal profession and academia

Six focus groups were held on 28 and 30 June, 4 and 25 July, 2 August and 3 October

In total, there were 37 participants from corporates, audit, governance/advisory, legal and academia. The 18 companies involved included 14 FTSE companies (11 FTSE 100s), two small listed companies, one large private company, and one private equity firm

I compiled a feedback document which was circulated to all participants for comment/approval

I sent the feedback to the team running the Government consultation

October – investor focus group held and feedback received

On 16 October I chaired a focus group with investors

The group included six investors: four from large UK investment managers, one independent investment manager and a representative from an investor body

I compiled a feedback document which was circulated to all participants for comment/approval

I sent the feedback to the team running the Government consultation

November – guidance for Boards/management developed and published

With input from various focus group contributors, I drafted this guidance document

I published this guidance on 22 November

My thanks and acknowledgments

In total, 43 people took part in the focus groups. These included five investors, along with representatives from 18 companies (including 11 FTSE 100s) that produce reporting, plus others from audit, governance/advisory, legal and academia. They were invited from my network of those involved in or with an interest in reporting.

I am enormously grateful to all those who gave their time and thoughts to this research. Most wished to remain anonymous, but those happy to be mentioned by name include Damian Carnell, Carolyn Clarke (Brave Within LLP), Janice Denancourt, Paul Lee (Redington), Luciano Rubén Lilloy Fedele, Emma Scott-Smith and Teresa Watkins (PCC Wealth).

Appendix 1 – Methodology continued

How focus groups were run

All sessions were held online (on Teams) and lasted one hour

Pre-session briefing: Ahead of each session, I sent my original proposal to participants, along with various blogs I have written on the subject of AI; a link to the Government consultation; and some questions to prompt discussion

First 15-20 minutes: I presented my proposal to the group, along with my views of the benefits and risks to reporting of introducing AI, particularly LLMs

Remainder of the session: Open discussion amongst participants under the Chatham House anonymity rule, moderated by me

Contributions: All participants contributed

Consistency: All focus groups received the same proposal and presentation, and all company focus group participants received the same set of briefing questions

Recording: All sessions were recorded and transcribed; the output was seen and used solely by me and my assistant to write up the feedback documents and as reference for developing this guidance

Briefing questions sent to company focus group participants:

- Do you currently use any form of large language model AI (LLMs) in your company? And in reporting specifically? What's your experience if so?
- Are there any plans to do so? Is your company discussing the use of LLMs in any aspect of the business? What are the points/issues being raised?
- With reporting in mind, what might the benefits of using an LLM be? And the drawbacks? Please consider the specific context of reporting.
- Are you involved/interested in being involved in any other discussions on this subject?
- Have you/your company been involved in the Government consultation on regulating AI?
- What do you personally feel about the use of AI/LLMs in general? (We don't have to cover this if you'd prefer not to; it would just be useful insight.)
- Any other questions/thoughts/subjects you have in mind.

Briefing questions sent to investors:

- As investors, would you view the information in an annual report, or any statement made by a company to the market, differently if you knew it had been written – or parts of it had been written – using an LLM?
- Do you have any concerns about companies using AI, particularly LLMs, in creating reports/statements to the market?
- Do you think reporting could be improved by companies using AI, particularly LLMs, to write their reports?
- Do you think the use of AI/LLMs in reporting should be regulated? Or guidance provided?
- Do you think companies should disclose in the annual report if/how they have used AI/LLMs in its creation?
- Does your view on the use of AI/LLMs in creating reporting vary by type of information included? I.e. would it be OK for some sections and not others?
- Do you think the use of AI could potentially compromise directors' accountability for reporting?
- Do you currently use any form of LLMs in your company? Do you use AI at all for summarising/digesting the information you receive via annual reports?
- What do you personally feel about the use of AI/LLMs in general? (We don't have to cover this if you'd prefer not to; it would just be useful insight.)
- Any other questions/thoughts/subjects you have in mind.

Appendix 2 – Types of AI

Artificial intelligence (AI) is a broad discipline with roots in the 1950s, focused on creating machines capable of mimicking human intelligence. Companies like IBM, with its Deep Blue and Watson systems, were pioneers in this field. AI encompasses a vast range of technologies, including machine learning, generative AI, and large language models (LLMs), among others. For reporting we are mainly concerned with the potential impact of the use of LLMs in generating and writing content, although other forms of AI may be used in the development of source material.

Machine learning

Machine learning (ML), a subset of AI, was developed in the 1980s. Its main aim is to enable machines to learn from data, improve their performance, and make decisions without explicit programming. Google's search algorithm is a prime example of ML application, using past data to refine search results.

Generative AI

Generative AI (GAI) evolved from ML in the early 21st century, and represents a class of algorithms capable of generating new data. These algorithms construct data that resembles the input, making them essential in fields like content creation and data augmentation.

An important subset of GAI is the Generative Adversarial Network (GAN), introduced by Ian Goodfellow in 2014. GANs consist of two neural networks: a generator that produces synthetic data, and a discriminator that purports to distinguish this data from real instances. GANs have gained popularity in image synthesis and modification.

Large language models

Large language models (LLMs) also arose from the GAI subset. In simple terms, LLMs generate human-like text by predicting the likelihood of a word given the previous words used in the text. They are the core technology behind many voice assistants and chatbots. OpenAI's GPT model is a well-known example.

GPT, or 'Generative Pretrained Transformer', is a specific type of LLM developed by OpenAI. GPT models are trained on vast amounts of text data and can generate coherent, contextually relevant sentences. Introduced with GPT-1 in 2018, it evolved to GPT-2 in 2019, and GPT-3 in 2020, each generation producing more coherent, relevant results.

ChatGPT, a derivative of the GPT family, is an AI conversational model. It generates responses to text input, meaning it can be used to draft emails, write code, create written content, or carry out what appear to be engaging conversations. OpenAI has been at the forefront of this technology, making strides in its development and application.

> These explanations have been extracted from an article on AI, *So What's The Difference Between AI, GAI, ML, LLM, GANs, and GPTs?*, published by Nicholas Cropp, Global Digital Design Director at Jo Malone London, on LinkedIn on 20 May 2023.

Appendix 3 – About Falcon Windsor

Founded in 2004 by Claire Bodanis, Falcon Windsor is a team of independent experts committed to delivering thoughtful, creative and meticulously accurate corporate reports, for companies small and large, private and listed. We bring together critical thinkers, strategic planners, writers, designers, and production and project managers with impressive credentials in corporate reporting.

Between us, we've delivered hundreds of annual reports and thousands of other communications projects. Many of us have worked client-side too. Reporting's in our bones.

That's why we love sharing our expertise: through our book, webinars, conference appearances; and through working with regulators, and people from every aspect of company life.

But let's park the monumental erudition for a minute. Because that's not what gets you through the midwinter marathons and springtime sprints of the reporting year. That takes a mixture of resilience, integrity, the instinct to help, and a genuine love of the job. It's that indefinable something – character? – that helps us become partners to our clients.

Trust me, I'm listed

In July 2019, the Chartered Governance Institute commissioned Claire to write a book on how to do corporate reporting well. With a foreword by Sir Donald Brydon, and contributions from experts across the reporting world, *Trust me, I'm listed – why the annual report matters and how to do it well*, was published in June 2020. The second edition, with updates on the ESG reporting landscape and the future of digital reporting, was published in October 2021.

Contact

Claire Bodanis

claire@falconwindsor.com

+44 7966 196808

falconwindsor.com

trustmeimlisted.com

Views? Thoughts? All welcome!

Please get in touch with me:

Claire Bodanis
claire@falconwindsor.com
+44 7966 196808

Falcon Windsor

AI disclosure statement

I hereby confirm that:

- This guidance document was conceived, written, designed, reviewed, checked and published by human beings, and
- To the best of my knowledge, no AI system of any kind has been used in its development, creation and publication.

Claire Bodanis