

Using AI Ethically in Business Intelligence

This material is produced by the Ethics and Compliance Team at the BHBIA. It is meant as general, not specific, advice on the use of AI. The AI space is evolving rapidly and it is an AI deployer's responsibility to stay current on technical capabilities, limitations, and administrative risks

This guidance focuses on the **ethical issues specific to AI**, rather than the broader legal, compliance, privacy and governance requirements that must be followed. Those requirements remain essential, but they do not fully address the ethical questions AI raises.

The distinctive ethical challenge of AI in healthcare business intelligence is not that AI outputs can be wrong (humans are frequently wrong too), or that it needs governance. It is that it can generate persuasive, decision-shaping outputs that alter what organisations believe, but in a different way to the traditional business intelligence process.

AI changes the character of insights work. It can improve the quality and speed of analysis, but also changes the way knowledge is formed, and whose perspectives become visible. That matters in healthcare analytics because these outputs can shape strategy, evidence priorities, campaign development, forecasting, etc. long before any decision is made. In roles built around guiding strategic decisions using market research and analytics, this ethical layer is intrinsic to our day-to-day work – meaning we need to think about how to maintain high ethical standards when using AI tools.


Seven Principles for Using AI Ethically

The AI-specific ethical questions in healthcare business intelligence can be framed around the following principles:

1. Fluency bias: do not mistake AI polish for proof

AI can make an answer sound clear and confident even when the evidence behind it is thin. AI can make errors – which is fine, although frustrating, when we can recognise them – but the real danger is that we falsely believe an incorrect output because it sounds fluent and plausible. This could lead to a business intelligence professional sharing a misleading insights summary with a client or stakeholder, having trusted the AI output because it appeared complete and well articulated.

This also extends to building our own tools. A synthetic respondent powered by AI that is built on thin source material will answer a question as confidently as one built on a rich evidence base – and even a well grounded synthetic respondent may answer confidently in areas where it has little source data to work with.



AI tools are valuable, but outputs should not be accepted at face value: they need validation against the evidence, and we should ensure that our reports and insights are correct, whether AI-generated or not.

2. Be transparent with how AI has been used

Be clear about how AI has shaped your output. Using AI to transcribe interviews is not the same as using AI to interpret evidence, or to draft a full summary report. Transparency is always required, but the level of explanation and/or validation should increase when AI materially shapes the narrative of an output. The speed benefit of creating a first pass fully AI-generated report could be valuable, but to do so ethically you should communicate explicitly and prominently that the report was fully AI-generated and requires validation.

In summary, you should be transparent about what part AI played in any output you produce so that users can interpret it appropriately.

3. Look for the signals that AI might miss

AI tends to privilege what is most clearly articulated or easiest to parse in data. That means well documented perspectives, populations, markets and pathways can receive disproportionate attention, while less visible groups or problems can be neglected, and cultural differences may be missed. This issue is not unique to AI – it exists in research more broadly – but AI can intensify it because it can convincingly amplify what is most visible in the available data.

In contrast, being able to identify a rare but strategically important signal is a valuable skill for an insights professional.


Conducting ethical market research means reporting the diverse range of respondent perspectives, whether AI is used for analysis or not. Consider noting unusual or minority viewpoints throughout research, then check that they have been carried through in any AI-supported analysis.

4. You are accountable for the decision, even if AI recommended it

In future, AI may become better than humans at reasoning, prediction and many forms of analytical judgement. That does not remove human responsibility. Deciding what matters, which trade offs are acceptable, and who bears the consequences remains with you and your organisation, as does responsibility for the research you produce, and the actions taken from it. That means ensuring the work meets the ethical standards the BHBIA and your organisation expect, and where an ethical judgement is made, making it explicit and owning it – even when AI has supported the decision.

5. Be wary of AI analysis that ignores the need for ethical judgement

AI models will try hard to answer your prompt, but if you don't include any ethical considerations in your prompts, AI will likely not use that lens in the response it provides. This means that a decision that has an ethical component (e.g. whether



to focus on the private insurance market vs. aim for wider access to a medicine) may be presented as a mathematical optimisation problem. This is not an issue in itself – it could be a valuable analysis and both options may be defensible – but we need to ensure that we continue to highlight the ethical considerations in our business intelligence work, rather than avoiding those discussions.

6. Account for downstream influence of decisions

Even when used upstream in business insights and analytics, AI can shape downstream decisions that affect patients, healthcare professionals, access and resource allocation. Ethical scrutiny should reflect downstream consequences – this means that we should document when (and how) AI has been used to inform ethically sensitive downstream decisions. Poorly grounded AI-driven work comes with the risk of taking decisions that can damage trust in healthcare research and the wider industry, which is exactly what the BHBIA code asks us to protect.

Consider limiting use of AI at certain points in the decision making process to ensure, where needed, human judgement remains at the forefront.

7. Analytical rigour still needs expert stewardship

AI can produce apparently sophisticated analysis very quickly, including for people who are not trained analysts. That creates a real risk of poor framing, weak assumptions, bad data fit, or superficial interpretation being turned into polished outputs that look robust. Even if AI becomes very strong analytically, organisations still need highly trained analysts and researchers – or strong embedded guardrails designed by them – to ensure that outputs are methodologically sound and interpreted appropriately. The ethical issue is ensuring analysis is being produced under conditions that make it accurate and reliable.

Note that for many of these challenges, AI can be part of the solution – AI tools could be used to assess business intelligence work for each of these ethical considerations and provide guidance on how to action them. But ultimately, the ethical accountability ends with us.


Why Are These AI Ethical Principles Needed?

Business intelligence professionals should embrace new methods and technology, but to use these responsibly we need to understand how they are affecting our work and our thinking.

For instance, consider an insights project flow. There are many benefits in using AI to increase efficiency and quality of results, but there are associated challenges:

Stage 1 – Scoping the project

AI challenge: AI will scope a problem based on what it can find. It can be great at writing a brief – but if it is also doing the background research, AI may over rely on easily accessible, mainstream sources and under weight the less documented patient groups and pathways. The question will be framed around what data the tool can easily access and could miss crucial perspectives.



Stage 2 – Fieldwork & data collection

AI challenge: AI can (potentially) make a great interview moderator – it can mimic human speech and appearance, it can probe, build rapport, and can feel incredibly empathetic and reactive. But it is a simulation of a moderator – it doesn't truly respect a respondent's time, nuance, or emotional boundaries, or read the room the way an experienced moderator does. This becomes an ethical issue when researching sensitive topics, such as end of life care.

AI can still be used in research with vulnerable respondents, but it should be used carefully, with risks assessed and with the human researcher holding the duty of care to the human respondent. The safety and respect of respondents must underpin any interaction – be it an AI interaction, or human.

Stage 3 – Analysis & interpretation

AI challenge: With AI, rigour may be skipped. AI can produce sophisticated looking analysis quickly, including by people who aren't trained analysts or who are under time pressure. Weak assumptions, incomplete data or shallow interpretation can be presented as a robust output, and could damage the reputation of the individual or company who uses it. By being transparent in how AI is used, and training researchers well, we can avoid these pitfalls.


The same AI tool can give very different results depending on the skill of the user. This includes how well a prompt is written, the level of context given, and the right choice of data and sources. The better we train business intelligence professionals to use these tools – and to know when to question what comes back – the higher quality our work will be, and the more likely it is we meet our high ethical standards.

Stage 4 – Building a report and narrative

AI challenge: In certain circumstances, AI can make a thin finding sound well evidenced and comprehensive, or even fabricate a response. A polished narrative travels well, and the underlying uncertainty may be forgotten along the way. AI derived findings should be validated before being included in a report (or explicitly labelled as AI generated and not yet validated).

Stage 5 – Recommendation & decision making

AI challenge: Moral and value choices could be presented as technical outputs by AI. For example, if you ask an AI to generate an optimal HCP targeting list, it might mathematically optimise for the highest immediate revenue – quietly filtering out rural doctors or harder to reach patient populations. Deciding whether to trade equitable access for commercial efficiency should be a human, strategic choice. Using AI risks that these trade offs are hidden inside a statistical black box, presenting a moral choice as an objective "best practice" output. By documenting how AI is used when making these decisions, it can help us identify when to limit use or bring the decision explicitly back to human judgement.



AI Ethics and Compliance – Complementary Concepts

This guidance has focused on using AI ethically. The AI compliance and governance layer is also essential but has a different purpose. The compliance layer asks whether AI use is permitted, secure, documented and controlled. Ethics deals with what happens next – how AI is changing the way we generate business intelligence and make decisions. By considering both together, we ensure that the work we do is of high quality and of high ethical standards.

The BHBIA maintains a dedicated page for [AI related guidance](#) and regulation on the BHBIA website if you would like more guidance on the topic.

Acronyms and definitions

AI – Artificial Intelligence. Note that artificial intelligence covers a wider range of technologies – in this article we focus on the outputs of tools powered by large language models (LLMs), but the guidance is intended to be relevant to most AI technologies.

BHBIA – British Healthcare Business Intelligence Association.

Fluency bias – The tendency to over trust AI outputs because they sound clear, confident and well articulated, regardless of the underlying evidence.

HCP – Healthcare professional.

Synthetic respondents – Simulated research participants built on an existing AI model, usually a large language model, and grounded in source material such as interview transcripts, published sources and web content. Used in place of, or alongside, real respondents.

The BHBIA's Ethics & Compliance Committee is providing this guidance as general information for its members. It is not legal advice and should not be relied upon as such. Specific legal advice should be taken in relation to any specific legal problems or matters. Whilst every reasonable effort is made to make sure the information is accurate, no responsibility for its accuracy or for any consequences of relying on it is assumed by the BHBIA.

Produced by the BHBIA's Ethics & Compliance Committee September 2026

British Healthcare Business Intelligence Association
Fountain Precinct, 4th Floor Orchard Lane Wing, Balm Green, Sheffield, S1 2JA
t: 01727 896085 • admin@bhbia.org.uk • www.bhbia.org.uk
A Private Limited Company Registered in England and Wales No: 9244455

