

Public Consultation Feedback Form

The Health Information and Quality Authority (HIQA) is an independent statutory body established to promote safety and quality in the provision of health and social care services. HIQA has a responsibility to develop standards, recommendations and guidance to support the Irish digital health and health information landscape to ensure safer, better care for people using health and social care services.

HIQA has been requested by the Department of Health to develop national guidance to promote and drive a responsible and safe approach to the use of Artificial Intelligence (AI) in the health and social care sector in Ireland. The national guidance aims to support services to promote and drive a responsible and safe approach to the use of AI. The main purpose of the guidance is to promote awareness and build good practice among services and staff about the responsible and safe use of AI in their services. The guidance will also be of use to people using services by educating and empowering them on what their expectations should be in respect of how AI can be used safely and responsibly while engaging with health and social care services.

The six-week public consultation gives people an opportunity to provide feedback on the draft national guidance and become involved in the development process by submitting their views.

Please note: the focus for this consultation is the content and structure of the draft National Guidance. The final design and layout of the guidance will be developed after the public consultation.

HIQA will carefully assess all feedback received and use it, where appropriate, along with other available evidence, to inform the final version of the National Guidance for the Responsible and Safe Use of AI in Health and Social Care. Before you complete this consultation feedback form, please read the instructions for submitting feedback on the following pages. The draft guidance has been informed by extensive stakeholder engagement. HIQA has also conducted an evidence review to inform the development of the draft national guidance and it is published on the HIQA website www.hiqa.ie.

The consultation closes at 5pm on 05 March 2026.

Data Protection and Freedom of Information (FOI)

This consultation is being conducted in accordance with data protection law, including the GDPR and Data Protection Act 2018.

HIQA will only collect and store personal information during this consultation for the purposes of verifying your feedback. For further information on how HIQA uses personal information, please see our Privacy Notice available [here](#). If you have any concerns regarding your personal information, please contact HIQA's Data Protection Officer on dpo@hiqa.ie.

Following the consultation, HIQA will publish a report summarising the responses received, which will include the names and types of organisations that submitted feedback. For that reason, it would be helpful if you could explain if you regard the information you have provided as being confidential or commercially sensitive.

Please note that HIQA is subject to the Freedom of Information (FOI) Act and the statutory Code of Practice for Public Bodies in relation to FOI. HIQA cannot give you an assurance that confidentiality can be maintained in all circumstances, due to the requirements of the FOI Act.

By submitting your feedback, you are agreeing to participate in this consultation.

Instructions for submitting feedback

- If you are commenting on behalf of a service or organisation, please combine all feedback from your organisation into one submission form and include the details of the service or organisation.
- Please do not paste other tables into the boxes already provided — type directly into the box.
- Hard copy: If you are handwriting responses, please feel free to use additional paper.
- Qualtrics: Please ensure that you click through all pages of the form, you will know you have reached the end of the questionnaire when you click the complete button.
- Please spell out any abbreviations that you use.
- When commenting on a specific section of the document, it would help if you can identify which part you are commenting on and the relevant page number.
- The questions are not intended in any way to limit your feedback, and other comments relating to the draft national guidance are welcome.

1. About you

1.1 Are you providing feedback as:

an individual

on behalf of an organisation (for verification purposes, please provide your name and your role in the organisation and your contact details).

Name of the organisation: **Medical Protection Society**

Your name: **Alice O Flynn**

Your role in the organisation: **Policy and Public Affairs Officer**

Your contact details: **alice.oflynn@medicalprotection.org**

1.2 Are you commenting?

In a professional capacity (*Please specify your role, discipline and the organisation you work for*)

Your role: **Policy and Public Affairs Officer**

Your discipline: **Policy and Public Affairs**

The organisation you work for: **Medical Protection Society**

As a member of the public/user of health and social care services

If you would like to provide any additional details, please share below.

2. General feedback on the draft National Guidance

2.1 Please provide any general feedback you have on the structure (layout, length, flow) of the draft National Guidance. Where applicable, please specify the section of the guidance document to which you are referring. Please note, the final guidance document will be professionally designed so the current format is only indicative.

We welcome HIQA's intention to provide national guidance on the responsible and safe use of AI, which is a positive step in clarifying expectations for services and staff.

However, the draft guidance is very long and often repetitive. Many concepts, particularly legislative references and descriptions of governance responsibilities, appear multiple times across sections, making it difficult to identify clear takeaway messages. To be more user-friendly, the document could be more succinct and to the point.

There is also a heavy reliance on lengthy text; visuals or summary formats would make the guidance easier to interpret. A table outlining the different regulations, who they apply to, oversight bodies and jurisdiction, with links rather than references, would significantly improve clarity and ease of use.

The flow of each principle could be strengthened by reordering the principles e.g. by bringing the "what this means for people using services" section to the top. This would allow each section to begin with the aim or goal, followed by how to achieve it, and then what it looks like in practice. The current structure, where the purpose appears only at the end, feels counterintuitive.

The guidance also repeatedly cites the same regulations (e.g., IVR, MVR, EU AI Act). It would be simpler to present these once at the beginning as a consolidated list and then signpost only the relevant ones at the start of each section rather than re-explaining them throughout.

2.2 Is the content provided clear and understandable?

Yes No

If no please provide detail on how the clarity of the content can be improved

We do not believe the content is fully clear and understandable in its current form. The volume of background information presented before the core guidance makes it difficult for readers to identify the most important points. A separate executive-summary-style document highlighting the actionable elements would significantly improve clarity.

The explanations within each principle are lengthy, and by the time readers reach the examples, it can be hard to recall the specific point the example is intended to illustrate. Case studies are very wordy; presenting them in bullet-point form aligned clearly with the stated aims of each principle would improve comprehension.

The content would also benefit from additional visual aids, especially given the complexity of the topic and the need for quick understanding in situations where AI-related concerns may affect patient safety. Although the jigsaw visual is helpful, more illustrations would support users with different learning styles and enhance accessibility.

2.3 Do you think the language used in the draft national guidance is clear and easy to read and follow?

Yes **No**

If no please provide examples of where you think we can improve the language

While the language itself is not particularly complex, the overall volume of text makes the guidance feel overwhelming. The high density of information, combined with some repetition, creates a sense of information overload, which makes it harder to grasp the key messages.

It is important to consider the full range of people who will need to use and interpret this guidance. Although the current level of detail may suit those familiar with policy documents, many staff in health and social care settings will require much clearer, more straightforward language to understand what is expected of them in practice. Without simplifying the presentation, there is a risk that users may lose sight of the core requirements.

A more concise approach would help ensure that the guidance is workable, practical, and accessible for all intended users.

2.4 Do you think the introduction (section 1) to the document will help staff understand the wider context regarding the responsible and safe use of AI in health and social care services in Ireland?

Yes No

If no please specify how the guidance can be improved

The introduction contains useful material, particularly the overview of the current regulatory landscape and the explanations of activity at European level. These elements provide important context and help readers understand *why* national guidance is needed and *who* the key actors are in the governance of AI.

However, the introduction would benefit from being more balanced and more concise. It presents the benefits of AI upfront but does not highlight the risks until much later in the document. Including a brief summary of the key risks in the introductory section would ensure a more even and accurate framing of the topic.

The introduction is also particularly long, at 15 pages, which may be challenging for many users to navigate. Given the wide range of staff who will rely on this document, it is important that the introduction is pitched at an appropriate level and is accessible to those with varying levels of AI knowledge.

The section describing different sub-types of AI is detailed, but it is unclear how this technical information supports the subsequent guidance. Many users will not need to understand concepts such as “multi-layered artificial neural networks,” and such descriptions may be too technical for a general health and social care audience. This content may be more effective if simplified into broader categories (e.g., diagnostic, predictive, administrative systems) or moved to an appendix for readers who wish to explore the subject in more depth.

Overall, while the introduction contains relevant contextual information, it could be significantly strengthened by being shorter, more balanced, and more focused, ensuring that staff can quickly understand the key contextual points without feeling overwhelmed.

3. Feedback on the principles underpinning the draft National Guidance

In this section we want to find out what you think of the guidance under each principle. The questions in this section are not intended to limit your feedback and other comments relating to the guidance is welcome.

The draft National Guidance is underpinned by four principles:

- Accountability
- A human rights-based approach
- Safety and wellbeing
- Responsiveness

Under each principle there is:

- An explanation of what the principle means for the responsible and safe use of AI in health and social care.
- Examples of how a service (including people who are responsible and accountable for managing the service) can uphold the principle in practice.
- Examples of how staff can uphold the principle in their day-to-day work.
- Examples of what it means for people who use health and social care services when the principle is upheld.
- A case study to provide context by showing what the responsible and safe use of AI looks like in a health and social care setting.

Please consider the following questions as part of your review of the guidance and case studies under each principle:

1. Have all important areas relating to this principle been addressed? Are there any other areas that should be included?
2. Do you think the case study helps understand the principle in practice? Is it reflective of what would happen in practice?

3.1 Please provide your feedback on the principle of accountability below:

When commenting on a specific aspect of the principle it would help if you can identify which part you are commenting on and the relevant page number.

Overall, we support the emphasis on accountability and the central role of human oversight in the safe use of AI. However, several areas within this section would benefit from greater clarity and consistency.

First, there is inconsistent wording regarding human oversight. On page 20 the guidance states that there “*should*” be human oversight, while page 21 states there “*must*” be. This inconsistency appears again later in the document. We recommend choosing a one word, preferably “*must*” where the expectation is mandatory and applying it consistently throughout the guidance.

In the section that defines roles and responsibilities, the guidance states that staff hold responsibility for interpreting AI outputs and ensuring their appropriate use. We believe this is only reasonable if staff have been adequately trained and educated in how an individual AI tool works including its limitations.

Currently, page 23 states that training “*should be made available*”; however, we do not believe this goes far enough. Training should be mandatory, and the responsibility for delivering it should rest with the healthcare organisation, not with individual clinicians, and the content of the training itself rests with the AI provider. Otherwise, clinicians may unfairly bear responsibility for decisions made without sufficient support, a challenge which is already experienced in Ireland where new technologies and services may be introduced without robust governance in place.

A further gap is the absence of any reference to liability. AI providers sometimes include contractual clauses limiting or excluding their responsibility for the performance of their systems. This is a significant issue in practice. The guidance should either address the issue directly or advise clinicians and services to avoid entering into contracts that absolve AI providers of liability and accountability.

Finally, the examples given under this principle are exclusively hospital-based. Given that the guidance applies to a wide range of providers, including smaller and community-based services such as dental practices, the examples should be broadened to reflect the diversity of settings covered by the scope of the guidance.

3.2 Please provide your feedback on the principle of a human rights-based approach below:

When commenting on a specific aspect of the principle it would help if you can identify which part you are commenting on and the relevant page number.

We welcome the commitment to transparency and informed decision-making reflected in this principle. However, several important elements require further clarification to ensure that rights are meaningfully upheld in practice.

The guidance does not address what should happen if a patient chooses to decline the use of AI in their care. While it states that people should be made aware when AI is used, it does not clarify whether opting out is possible, nor whether doing so would affect the standard of care offered. In reality, many patients may not wish to engage with AI-enabled tools, and the guidance should explicitly confirm that declining AI must not negatively impact their access to safe, timely, and equitable care.

Similarly, the guidance does not consider the scenario in which a clinician does not wish to use AI that has been introduced by their organisation. Clarity is needed on how such situations should be managed, and what responsibilities fall to the service provider versus the individual clinician.

We also note that the example of informed consent provided under this principle is extremely detailed. In practice, clinicians already face significant time pressures, and the level of information suggested may be unrealistic without substantial organisational support. It would be helpful for the guidance to outline a more practical and proportionate consent process, including the role of the service provider in supplying standardised information materials (e.g., patient leaflets) so that the burden does not fall entirely on individual clinicians. Additionally, it is important to note that the detail and depth of the consent process may vary according to the complexity and use of the tool.

3.3 Please provide your feedback on the principle of safety and wellbeing below:

When commenting on a specific aspect of the principle it would help if you can identify which part you are commenting on and the relevant page number.

We support the strong emphasis on using AI tools safely and in line with their intended purpose. However, this section would benefit from acknowledging that AI technologies may evolve rapidly, and that safe use requires more than one-off training or a one-time understanding of an AI system's instructions for use.

AI systems are likely to undergo updates, refinements, or expanded capabilities over time. The guidance should therefore highlight the need for ongoing monitoring, updated training, and refreshed guidance so that staff remain aware of changes to functionality, limitations, or risk profiles.

Linking the principle explicitly to the need for continuous learning and updated safety protocols would strengthen its applicability in real-world settings, where AI systems do not remain static.

3.4 Please provide your feedback on the principle of responsiveness below:

When commenting on a specific aspect of the principle it would help if you can identify which part you are commenting on and the relevant page number.

The emphasis on recognising concerns—such as bias or system flaws, is welcome, but the guidance would benefit from greater clarity on how staff should respond to concerns, and within what time frame.

Current risk-management pathways within the HSE can be slow, with some risk committees meeting infrequently. As a result, concerns about AI that pose an urgent risk may not be escalated promptly if they enter routine reporting cycles. The guidance should therefore highlight the need for timely escalation processes aligned with the level of risk, including clear routes for urgent action where required.

While the document rightly notes the need for robust risk management and risk registers, it is important to recognise that developing comprehensive risk registers can take significant time. Clear, practical steps for interim escalation, especially in high-risk scenarios, would be valuable.

The guidance mentions at an earlier point that AI-related device issues should be reported to the HPRA. We believe this should be reiterated here, as many smaller service providers (including dental and community services) may be unfamiliar with where or how to report concerns. Providing a simple flow chart that outlines what types of concerns should be reported, to whom, and in what circumstances would greatly improve clarity. Even if placed in an appendix, such a tool would support staff and organisations with varying levels of AI literacy.

Finally, it may help to distinguish more clearly between what staff are responsible for and what organisations are responsible for, ensuring that the guidance remains accessible and practicable across both groups.

3.5 Do you think the guidance will be implementable in practice? What will support services and staff to implement the guidance?

When commenting on a specific aspect of the guidance it would help if you can identify which section you are commenting on and the relevant page number.

While the aims of the guidance are clear, its implementation in practice would be significantly strengthened through the inclusion of practical tools, such as flow charts, checklists, and embedded links that enable users to quickly access relevant information.

The examples and case studies provided are lengthy and focused primarily on hospital settings. Shorter, more varied examples, tailored to different environments such as dental practices, community services, and primary care, would make the guidance more inclusive and easier for a wider range of providers to apply.

Splitting case studies into sections describing what the organisation must do and what staff must do would also improve clarity.

Given the length of the document, users may struggle to identify the parts most relevant to their role. An interactive digital version of the guidance, allowing users to select their role or service type and be directed to tailored content, would greatly support practical implementation. Alternatively, a short set of initial screening questions could guide users to the sections most relevant to them, reducing the need to read the entire document in detail.

3.6 Any final thoughts or feedback to add:

We welcome the development of this national guidance and appreciate the opportunity to provide input.

We would also like to highlight MPS's wider programme of work on AI, which we believe is directly relevant to the development of effective national guidance.

Our [AI White Paper](#) was developed by the MPS Foundation in partnership with the University of York's Centre for Assuring Autonomy and the Improvement Academy at the Bradford Institute for Health Research. This paper examines how AI decision-support tools affect clinicians and identifies barriers to safe adoption, including concerns that clinicians could become "liability sinks" under current legal frameworks. A liability sink in AI refers to a situation in which legal responsibility for outcomes or harms caused by artificial intelligence systems is shifted onto parties with the least control over the design, deployment, or operation of these systems.

Complementing this research, our [AI Safer Practice Framework](#), developed by Dr Raj Rattan, offers a practical and memorable structure to support the safe integration of AI in clinical settings. This framework is made up of two components, INFORMED and RECORDS. INFORMED guides clinical decision-making using AI, while RECORDS documents AI-assisted decisions for accountability and clinical rationale. The framework has been structured around these acronyms to ensure it is practical and memorable.

This work demonstrates our expertise in the area and our ongoing commitment to supporting responsible AI adoption in healthcare. We would be happy to discuss this work further and to contribute to the development of future iterations of HIQA's guidance.

Thank you for taking the time to give us your views on the draft National Guidance for the responsible and safe use of AI in health and social care services.

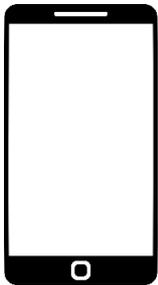


You can **email** the completed form to hist@higa.ie

OR

Print the consultation feedback form and **post** the completed form to:

National Guidance for the Responsible use of
AI in Health and Social Care Services



[If you have any questions on this document, you can contact the HIQA Health Information Standards Team either by:](#)

[Phoning: \(01\) 814 7400](#)

[Or](#)

[Emailing: hist@higa.ie](mailto:hist@higa.ie)

Please ensure that you submit your form online or return it to us either by email or post by 05 March 2026