DHCNI

Digital Health & Care Northern Ireland

HSC Goes Digital

DoH www.health-ni.gov.uk HSC Health and Social Care

HSC Artificial Intelligence Framework

January 2026

A shared reference to help us navigate the opportunities and challenges of Artificial Intelligence (AI) in Health and Social Care (HSC) in Northern Ireland.

Tom Simpson
Digital Health & Social Care NI

Document Version Control

Version	Date	Changes Made to previous version	Author
1.0	19/10/2025	1st draft presented to DHCNI Team	T Simpson
1.2	23/10/2025	Further iteration incorporating feedback from DHCNI Team	T Simpson
2.0	04/11/2025	Updated to incorporate further changes Presented to AI Steering Group, Enterprise Architecture Board, SET Digital Health & IG Sub-committee Group and Data Change Board	T Simpson
2.3	12/1/2026	Updated to include Table of Contents and version Control page	T Simpson
2.6	21/1/2026	Updated to incorporate further changes and suggestions from DHCNI Comms and include Glossary	T Simpson

Table of Contents

Introduction

Thanks for taking the time to read this.

This document is intended as a practical starting point - a shared reference to help us navigate the opportunities and challenges of Artificial Intelligence (AI) in Health and Social Care (HSC) in Northern Ireland. It's not a policy, it's not a strategy and it is most definitely not the final word. It's a framework to get us going - something we can build on together.

AI is already here. From tools that help clinicians interpret scans, to systems that support admin teams with correspondence or summarise meetings, we're seeing real-world use cases emerge across HSC. With that comes a responsibility to get it right - to make sure we're using these technologies safely, ethically and in ways that genuinely support our people and improve outcomes for patients and service users.

In principle, AI is just another digital tool - part of the broader transformation of how we deliver care and run our services, but it's also different. The speed of development, the complexity of the technology and the ethical and regulatory questions it raises mean we need to give it specific attention.

These technologies are here, and their use becoming increasingly more common. Colleagues across the HSC system have indicated the need for some guiding principles - a shared understanding of what "good" looks like, and some top cover (or perhaps more plainly, reassurance and alignment) to help us move forward consistently and with confidence.

That's what this framework is for. It draws on best practice from across the UK, Ireland and internationally - including NHS England, UK Government, the Information Commissioner's Office (ICO), Medicines and Healthcare products Regulatory Agency (MHRA), World Health Organisation (WHO) and others - but it's grounded in the realities of our system. It's designed to support innovation, not stifle it. To offer clarity, not complexity and to help us move forward with confidence, even when the path isn't always clear.

We've deliberately kept it high-level and non-prescriptive. Every organisation is at a different stage - and that's okay. What matters is that we're aligned on the principles, that we're learning from each other, and that we're building the capacity and culture to make the most of what AI can offer.

This is a living document. It will evolve as we learn more, as the technology develops and as we hear from you - the people doing the work. Please treat it as a resource, a conversation starter and a foundation for collaboration.

We don't need all the answers to get started - but we do need to move forward with purpose, curiosity and care. I'm looking forward to what we can achieve, together.

Tom Simpson

Director of DHCNI | Deputy Chief Digital Information Officer | HSC AI Steering Group Chair
October 2025

HSC AI Framework

0. Executive Summary

Artificial Intelligence (AI) presents transformative opportunities for Health and Social Care (HSC) in Northern Ireland - from improving clinical decision-making and operational efficiency to improving patient care and workforce support. However, realising these benefits safely and ethically requires a shared understanding of how AI should be approached across the system.

This AI Framework for HSC provides a high-level, non-prescriptive guide to support organisations in exploring, adopting, and scaling AI technologies. It is not a policy or strategy, but a practical reference document that offers guiding principles and ways of working to ensure AI is deployed responsibly and in alignment with public trust, professional standards and existing governance structures.

The framework draws on best practice from across the UK and internationally, including *NHS England's AI and Machine Learning guidance* and the *AI Playbook for the UK Government*, and is shaped by the needs and realities of HSC organisations. It is designed to:

- Support innovation while maintaining safety, privacy, transparency, and accountability.
- Enable local autonomy, allowing organisations to progress at their own pace.
- Provide clarity on governance, ethics and evaluation.
- Encourage collaboration and shared learning across the system.

It is intended as a living document, to be iterated and refined through engagement with stakeholders, and informed by real-world experience, pilot projects and emerging evidence.

1. Purpose and Scope

This framework sets out a shared foundation for the safe, ethical and effective use of Artificial Intelligence (AI) across Health and Social Care (HSC) in Northern Ireland. It is designed to support organisations in exploring and adopting AI technologies in a way that aligns with public trust, professional standards and system-wide priorities - while respecting the autonomy of individual organisations.

Artificial Intelligence (AI) in this context refers to digital technologies that can perform tasks typically requiring human intelligence - such as analysing data, recognising patterns, generating content, or supporting decision-making - in ways that are increasingly embedded across clinical, operational, and administrative settings.

The framework is:

- Not a policy or strategy: It does not mandate specific actions or restrict innovation. Instead, it offers guidance and guardrails to help organisations make informed decisions.
- Flexible and enabling: It supports progress at different levels of maturity, allowing organisations to tailor their approach based on local context, capacity and ambition.
- System-aware: It recognises the complexity of HSC and the need for alignment with existing governance, regulatory and ethical structures.
- Outcome-focused: It aims to ensure that AI adoption contributes meaningfully to improved safety, quality, efficiency and experience across health and social care.

All HSC organisations, including Trusts, Arms-Length Bodies, and regional services, along with organisations working as part of the HSC system are encouraged to make use of this framework. It is intended to evolve over time, informed by experience, evidence and engagement with stakeholders.

2. Principles for Responsible AI in HSC

To ensure AI is adopted safely, ethically, and effectively across Health and Social Care, this framework sets out a series of guiding principles. These are intended to support decision-making, build public and professional trust, and promote consistency across organisations, without limiting innovation or progress.

Safety & Security

AI must enhance - not compromise - patient and service user safety. Systems should be rigorously tested, validated in real-world settings and monitored continuously for performance and unintended consequences.

Transparency & Clarity

AI systems should be understandable to clinicians, staff and service users. Decisions made or supported by AI should be clearly explained with suitable documentation of how outputs are generated and used.

Human Oversight

AI should support - not replace - professional judgement. Clear lines of accountability must be maintained and human oversight should be embedded in all AI-supported processes.

Fairness & Equity

AI should be designed and deployed in ways that avoid bias and promote equitable outcomes. This includes inclusive data practices, diverse stakeholder engagement and attention to underserved populations.

Accountability

Organisations should ensure clear ownership of AI systems, including responsibility for procurement, deployment, monitoring and outcomes. Governance structures should support ethical decision-making and risk management.

Privacy & Data Governance

AI should comply with data protection legislation and uphold the highest standards of confidentiality and security. Data uses in AI systems should be lawful, necessary, and proportionate.

Sustainability & Value

AI should deliver measurable value to patients, staff, and the system. Projects should be outcome-focused, scalable, and aligned with broader digital transformation goals.

User Readiness & Training

AI should be introduced in a way that builds staff confidence, capability and safe use. All users must have access to appropriate training, guidance and support to understand how AI systems work, how to use them responsibly, and how to identify issues or unintended consequences. Training should be proportionate, role-specific, and embedded into ongoing digital literacy and professional development. Clear support routes must be available so staff can raise concerns, report problems and access timely assistance.

3. Governance and Oversight

Effective governance is essential to ensure that AI is deployed safely, ethically, sustainably and in alignment with the values and responsibilities of Health and Social Care (HSC) in Northern Ireland. This section outlines key considerations and recommended practices for local governance, while allowing flexibility for organisations to tailor their approach.

Local Governance Structures

Organisations should establish or designate appropriate governance mechanisms for AI projects. This may include:

- AI Leads or Champions within digital, clinical, or operational teams.
- Multidisciplinary oversight groups involving clinical, technical, legal, and ethical expertise.
- Integration with existing governance bodies, such as digital boards, clinical safety groups, or ethics panels.

Local AI Policies and Organisational Autonomy

HSC organisations may choose to develop and maintain their own local AI policies, reflecting their specific operational needs, governance arrangements and maturity. This is encouraged. Local policies play an important role in supporting safe and context-appropriate adoption, and several organisations across HSC already do this effectively, including examples such as the BSO AI Policy and the Northern Trust AI Policy.

Local or shared policies should complement - not duplicate or conflict with - system-wide principles. The HSC AI Framework is designed to support and enable these local arrangements by providing a common reference point, ensuring broad alignment while allowing flexibility for organisations to tailor governance, decision-making, and documentation to their own context.

Organisations are encouraged to share local policies, templates and learning through the HSC AI Steering Group to support consistency, reduce duplication and promote collective improvement across the system.

Decision-Making and Accountability

AI-related decisions should be made transparently, with clear documentation of:

- Purpose and intended outcomes of AI use.
- Roles and responsibilities across teams.
- Risk assessments, including Data Protection Impact Assessments (DPIAs), clinical safety cases, and ethical reviews.

Organisations should ensure that human accountability is maintained for all decisions supported by AI.

Alignment with System Governance

AI governance should align with broader HSC governance frameworks, including:

- Information governance and data protection standards (e.g. ICO guidance).

- Clinical safety and regulatory compliance, including MHRA expectations for Software and AI as a Medical Device.
- Ethical oversight, potentially involving regional bodies such as the Ethics Advisory Group.

AI systems classified as medical devices may be subject to EU regulations under the current Northern Ireland Protocol arrangements. Devices must typically carry CE marking, and UKCA-marked devices may require additional approvals. The MHRA remains the competent authority. Organisations should seek advice on regulatory status early in the development process.

Continuous Oversight and Learning

Governance should not be a one-off exercise. Organisations should:

- Monitor AI systems in real-world use
- Capture feedback from staff and service users.
- Share lessons learned across the system to support collective improvement, demonstrating effectiveness through appropriate evaluation.

4. Guiding Parameters for Safe Use

To support safe and responsible AI adoption across HSC, this framework outlines a set of guiding parameters - practical boundaries and expectations that help organisations manage risk, uphold standards, and maintain public trust. These are not rigid rules, but guiding practices that should be adapted to local context, risk profile and maturity.

Ethical Decision-Making

Organisations are encouraged to adopt structured approaches to ethical decision-making. The INFORMED framework (Medical Protection Society) offers a useful model, prompting consideration of:

- Intent: What is the purpose of the AI system?
- Need: Is AI the right tool for this problem?
- Fairness: Who benefits, and who might be disadvantaged?
- Oversight: How will decisions be governed?
- Risk: What are the potential harms?
- Monitoring: How will performance be tracked?
- Evaluation: What does success look like?
- Documentation: Is the process transparent and auditable?

Safety and Risk Management

AI systems should be subject to appropriate safety assessments, proportionate to their intended use and potential risks. This may include:

- Clinical Safety Cases: While DCB0129/0160 standards are not formally mandated in Northern Ireland, they offer a recognised and structured approach to clinical safety that organisations may choose to adopt, particularly where AI systems interface with clinical decision-making or patient care.
- Hazard Logs and Mitigation Plans: Organisations should identify potential risks and document how these will be managed throughout the lifecycle of the AI system.
- Post-Market Surveillance: Especially relevant for AI systems classified as medical devices, organisations should monitor performance in real-world use and report any safety concerns in line with MHRA guidance.

Where AI systems are considered medical devices, organisations should be aware of the regulatory landscape. CE marking remains the primary route for device approval in Northern Ireland and access to UKCA-marked devices may require exceptional approval from the MHRA. This is an evolving area, and early engagement with regulatory advice is recommended.

Data Protection and Information Governance

AI projects must comply with data protection legislation and HSC information governance standards. This includes (but is not limited to):

- Data Protection Impact Assessments (DPIAs).
- Lawful basis for data use, especially where personal or sensitive data is involved.
- Minimisation - anonymisation of personal data, or pseudonymisation where appropriate.

- Management of information across the life cycle of the project, including application of retention and disposal controls.
- Appropriate monitoring and review.

Documentation and Auditability

Organisations should maintain clear and accessible records to support transparency, accountability, and continuous improvement. This may include:

- AI system design and intended use - including scope, functionality, and expected outcomes.
- Decision-making processes and oversight arrangements - who was involved, what was considered and how risks were managed.
- Performance metrics and evaluation outcomes - including baseline comparisons, user feedback, and any unintended consequences.
- Data inventory of data used across AI projects - providing a structured catalogue or registry of all data assets that are used, generated or processed by AI systems.

Where appropriate, organisations may choose to use structured tools to support documentation - such as internal templates, risk registers or impact assessments. The goal is not to create unnecessary bureaucracy, but to ensure that AI use is auditable, explainable and aligned with organisational governance.

Public and Professional Engagement

Guiderails should be informed by appropriate engagement with the public and across the HSC system. This includes, but again is not limited to:

- Staff, patients and service users, to understand concerns and expectations.
- Professional bodies, to ensure alignment with clinical and operational standards.
- Ethics advisory groups, to support reflection and challenge.

5. Use Cases and Deployment Guidance

AI is already being trialled and deployed across Health and Social Care in Northern Ireland, with promising results in both clinical and administrative domains. This section provides guidance on identifying and deploying AI use cases, with a focus on real-world relevance, measurable value and safe experimentation.

Identifying Suitable Use Cases

Organisations should prioritise use cases that:

- Address clear pain points or service challenges.
- Are data-rich, repeatable, and measurable.
- Can be piloted safely and scaled if successful.

Deployment Principles

To ensure safe and effective deployment:

- Start small: Pilot projects should be modest, with clear objectives and evaluation criteria, but with clarity on potential approach to scaling at the outset.
- Co-design with users: Engage frontline staff, clinicians, and operational teams early.
- Integrate with workflows: AI should enhance, not disrupt, existing processes.
- Monitor and iterate: Use real-world data to validate performance and refine the approach.

Evaluation and Learning

Each pilot should include:

- Baseline and outcome measures (e.g. time saved, accuracy improved, user satisfaction).
- Feedback loops to capture experience and performance.
- Documentation to support transparency and future scaling.

Sharing findings across the system will support collective learning and avoid duplication.

Support for Pilots

The AI Steering Group is keen to support a small number of pilot projects in 2026/27, with guidance on:

- Selection criteria and submission process.
- Evaluation frameworks.
- Access to shared resources and expertise.

Further details will be developed through iterating the HSC AI Steering Group.

6. Data, Privacy and Security

AI systems rely on data to function effectively. In the HSC, data is often personal and sensitive, (health-related). Ensuring that this data is handled lawfully, securely and transparently is essential to maintaining public trust and meeting regulatory obligations. This section outlines key principles and expectations for data governance in AI projects across HSC.

Legal and Regulatory Compliance

All AI deployments must comply with the appropriate legal and regulatory requirements. This includes, but is not limited to:

- UK GDPR and Data Protection Act 2018
- ICO guidance on AI and data protection
- BSO and HSC information governance policies
- Clinical safety standards, where applicable.

Organisations must complete Data Protection Impact Assessments (DPIAs) for any AI system that processes personal data, and ensure appropriate legal bases for data use, in consultation with their Data Protection Officers (DPOs).

Data Minimisation and Protection

AI systems should use the minimum necessary data to achieve their intended purpose. Where possible:

- Use pseudonymised or anonymised data.
- Avoid sharing identifiable data with external vendors unless explicitly approved.
- Ensure data sharing agreements and contractual safeguards are in place.

Local considerations should be made as to use of data specific to data ownership and control. As an example, the BSO AI Policy explicitly prohibits the use of sensitive HSC data with public-facing large language models (LLMs), due to risks of data leakage and model training.

Secure Deployment

AI systems should be deployed within managed HSCNI platforms, including secure cloud environments approved by BSO and DHCNI. External vendors must meet minimum standards, such as:

- Cyber Essentials Plus or ISO 27001 accreditation.
- Recent penetration testing.
- Confidentiality agreements and SIRO approval for data use.

All outputs from AI systems must be subject to human review, and any externally-facing content generated by AI should be clearly attributed.

Transparency and Opt-Out

People have the right to challenge and request human intervention in relation to automated decision-making under UK Data Protection legislation. Any AI system that influences decisions about individuals must:

- Provide clear information about how decisions are made.

- Offer a route for human review or override.
- Ensure transparency in data use and model behaviour.

Monitoring and Audit

Organisations must maintain records of AI use, including:

- Purpose and scope of each system.
- Data sources and processing methods.
- Performance metrics and evaluation outcomes.

Local organisations are responsible for monitoring AI use within their service areas. Local organisations can choose to reserve the right to audit AI activity on organisational devices and networks.

HSC organisations are also expected to contribute to the central AI Register, ensuring timely coordination and updates in partnership with the HSC AI Steering Group, which maintains overall ownership and oversight.

7. Workforce and Culture

AI adoption will reshape how people work across Health and Social Care. To ensure this transformation is positive, inclusive and sustainable, organisations must actively support staff through the change - building confidence, capability, and clarity around the role of AI, as well as ensuring that staff are provided with appropriate training.

Augmentation, Not Replacement

AI should be seen as a tool to support and enhance human work – not to replace it. This principle should be clearly communicated to staff to avoid fear, resistance, or misunderstanding.

Building Capability

Organisations should invest in:

- Training and awareness for staff at all levels.
- Co-design approaches, involving users in shaping AI tools and workflows.
- Communities of practice to share learning and build confidence.

This includes both technical skills (e.g. prompting, validation) and broader digital literacy.

Change Management

AI adoption is a change process, not just a technology rollout. Organisations should:

- Use structured change management approaches.
- Engage staff early and often.
- Be transparent about what's changing and why.
- Support teams to adapt roles, workflows, and responsibilities.

Workforce planning should consider how AI affects job design, workload, and career pathways.

Organisational Readiness

Before deploying AI, organisations should assess:

- Staff readiness and engagement.
- Process maturity - AI should not be used to patch broken workflows.
- Leadership support - senior leaders must champion responsible adoption.

AI should be introduced in ways that empower staff, reduce friction, and improve experience.

Role of the HSC AI Steering Group

The HSC AI Steering Group will play a key role in supporting workforce readiness at a regional level, including:

- Developing and sharing guidance on safe, ethical AI use.
- Supporting organisations to manage change and build confidence.
- Facilitating system-wide collaboration and shared learning.

8. Evaluation and Monitoring

To ensure AI delivers real value and remains safe, organisations must build in evaluation and monitoring from the outset, enabling learning, improvement and maintaining trust.

Build Evaluation into Pilots, Trials and Projects

Every initiative should include:

- Baseline measures - what does the current process look like?
- Outcome measures - what impact is expected (e.g. time saved, accuracy, experience)?
- Success criteria - what does "good" look like?

Evaluation should be proportionate to the scale and risk of the project.

Monitor Performance Over Time

AI systems should be monitored in real-world use, including:

- Accuracy and reliability of outputs
- Safety and unintended consequences
- User experience and feedback
- Equity and fairness in outcomes

This supports continuous improvement and early identification of issues.

Feedback Loops

Organisations should create simple mechanisms for:

- Staff and service users to share feedback
- Technical teams to respond and iterate
- Governance groups to review and advise

Feedback should be captured, acted on, and used to inform future deployments.

Share Learning Across the System

Evaluation isn't just local - it's a system asset. Organisations are encouraged to:

- Share findings with the HSC AI Steering Group
- Contribute to communities of practice
- Help build a shared evidence base for what works

This supports smarter investment, faster scaling, and better outcomes across HSC.

9. Ethics and Public Trust

AI adoption in health and social care must be grounded in ethical principles and public trust. These are not optional extras - they are essential to safe, sustainable, and socially acceptable use of AI.

People at the Centre

AI should serve people - not replace or disempower them. This includes:

- Respecting dignity, agency, and professional judgement
- Ensuring AI supports better care, not just efficiency
- Designing systems that reflect real-world needs and values.

Fairness, Transparency, Accountability

Organisations should ensure that AI systems:

- Are free from bias and promote equitable outcomes
- Are transparent in how decisions are made and data is used
- Include clear lines of human accountability

Ethical oversight should be built into governance structures, not bolted on.

Public and Professional Engagement

Trust is built through dialogue. Organisations should:

- Engage staff, patient sand service users early in AI projects
- Be open about what AI is being used for and why
- Create space for challenge, feedback, and reflection

This includes working with ethics advisory groups and professional bodies.

Alignment with Standards and Governance

AI adoption should align with:

- National and international ethical frameworks (e.g. WHO, Ada Lovelace Institute, National Institute of Standards & Technology (NIST), ISO/IEC 42001)
- HSC governance and accountability structures
- Legal and regulatory requirements

Ethical use of AI is not just about avoiding harm - it’s about actively promoting fairness, trust, and human-centred design.

Glossary

Term	Definition
AiMD	Active Implantable Medical Device
BSO	Business Service Organisation
CE Marking	Conformité Européenne (European Conformity) - products to show they meet the European Union's (EU) health, safety, and environmental protection requirements
DHCNI	Digital Health & Care Northern Ireland
DPIAs	Data Protection Impact Assessments
DPO	Data Protection Officer
EPMO	Enterprise Project Management Office
ICO	Information Commissioner's Office
LLM	Large Language Model
MHRA	Medicines and Healthcare products Regulatory Agency
NIST	National Institute of Standards and Technology
SaMD	Software as a Medical Device
SIRO	Senior Information Risk Owner
ToR	Terms of Reference
UK GDPR	United Kingdom General Data Protection Regulation
UKCA Marked	UK Conformity Assessed - shows that the product meets UK-specific safety, health, and environmental protection requirements.
WHO	World Health Organisation

Annex 1.

Next Steps and Implementation Support

This framework is a starting point - not a finished product. The next phase will focus on embedding the principles, supporting practical adoption and enabling system-wide learning, with collaborative, system leadership being directed through the HSC AI Steering Group. Some practical next steps to be undertaken through the steering group are listed below.

Finalise Governance Structures

- Confirm Terms of Reference (ToR) for the HSC AI Steering Group
- Establish Core Leadership Group, Wider Reference Network and supporting sub-groups
- Define roles, responsibilities and meeting cadence

Publish and Share the Framework

- This framework is to be shared across HSC organisations
- Feedback and suggestions for iteration will be collated through the HSC AI Steering Group
- It will remain as a living document to support ongoing development

Launch the AI Register

- Develop and publish a live register of AI initiatives across HSC
- Include pilots, active deployments, and planned projects
- Support visibility, coordination, and shared learning

Support Practical Application

- Identify and support 2-3 priority pilots in 2026/27
- Provide guidance on evaluation, governance, and scaling
- Reference existing EPMO processes and explore external funding routes

Enable Collaboration and Capacity Building

- Use the Steering Group to convene stakeholders, share insights and build capability
- Promote communities of practice and cross-sector engagement
- Align with national and international developments

Annex 2.

References and Further Reading

This section provides a curated list of key frameworks, guidance documents and resources that have informed the development of this framework. It is intended to support further exploration, implementation, and continuous learning across the HSC system.

Note: The below list is not exhaustive

UK and Northern Ireland Guidance

- *NHS England Artificial Intelligence (AI) and machine learning* - Practical guidance for the safe and effective development, deployment and use of AI in health and care.
  https://www.england.nhs.uk/long-read/artificial-intelligence-ai-and-machine-learning/

- *NHS England AI in practice* – Knowledge repository covering best practice in AI adoption across the health & care system.
  AI in practice - NHS England Digital

- Artificial Intelligence Playbook for the UK Government - a practical guide for the safe and effective development, deployment and use of AI across government.
  Artificial Intelligence Playbook for the UK Government (HTML) - GOV.UK

- *Information Commissioner's Office (ICO) Guidance on AI and Data Protection* - Key principles and practical steps for ensuring data protection compliance in AI systems.
  https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/guidance-on-ai-and-data-protection/

- *MHRA Guidance on Software and AI as a Medical Device (SaMD/AiMD)* - Regulatory requirements for AI technologies used in clinical settings.
  https://www.gov.uk/government/publications/software-and-artificial-intelligence-ai-as-a-medical-device/software-and-artificial-intelligence-ai-as-a-medical-device

- *BSO AI Policy (v2.1)* - Internal policy guidance on the use of AI within the Business Services Organisation and wider HSC context.

International Frameworks and Standards

- *NIST AI Risk Management Framework (AI RMF)* – A voluntary framework to improve the ability to incorporate trustworthiness into AI systems.
  https://www.nist.gov/itl/ai-risk-management-framework

- *ISO/IEC 42001:2023 - AI Management System Standard* - International standard for managing risks and responsibilities in AI systems.
  https://www.iso.org/standard/81228.html

- *WHO Guidance on Ethics and Governance of AI for Health* - Global principles and recommendations for ethical AI in health.
  https://www.who.int/publications/i/item/9789240029200

Ethics and Public Engagement

- *Ada Lovelace Institute Reports* - Research and recommendations on data, AI, and society, including public attitudes and ethical considerations.
  https://www.adalovelaceinstitute.org

- *Medical Protection Society - INFORMED Framework* - A structured approach to ethical decision-making in AI.
  https://www.medicalprotection.org