AI in UK General Practice

A Complete Framework

Patient Safety • Practice Governance • Regulatory Compliance

Combining harm reduction principles with UK regulatory requirements
for ChatGPT, Claude, Gemini, and AI-enabled clinical tools

Dr Krishnan Pasupathi

MBBS MBA MRCGP

NHS GP Partner | Founder, Aryash Health

January 2026 | Version 1.0

Contents

Executive Summary

Over **230 million people** now discuss their health with AI chatbots like ChatGPT every week. In UK general practice, approximately **28-30% of GPs** report already using AI tools in consultations or related tasks—often without formal governance frameworks in place.

This creates a dual challenge:

- **Patients** are using AI for health queries regardless of professional advice, creating risks from hallucinated medical information, delayed care-seeking, and medication errors
- **Clinicians** are experimenting with AI tools (sometimes quietly), creating risks around confidentiality breaches, data protection violations, and ungoverned clinical use

This framework addresses both challenges by combining:

Harm Reduction: Practical guidance to keep patients safer when they use AI for health queries, acknowledging they will do so regardless of advice to abstain

Practice Governance: UK-specific regulatory compliance covering GDPR, clinical safety, CQC expectations, and professional accountability

The document provides actionable tools including:

- Patient-facing materials for waiting rooms and websites
- Consultation approaches for discussing AI use
- A risk ladder for categorising AI use cases
- A draft AI Acceptable Use Policy for practices
- DPIA prompt questions
- Implementation checklist with regulatory citations

The goal is not to endorse AI for medical decision-making, but to acknowledge reality and work within it to protect both patients and practices.

PART ONE

Understanding the Landscape

Section 1: The Reality We Face

1.1 Three Different Conversations

When people discuss 'ChatGPT in healthcare', they're typically conflating three distinct scenarios with very different risk profiles:

Scenario	Description	Primary Risks
1. Patients at home	Using ChatGPT/Claude/Gemini for symptoms, explanations, lifestyle advice	Hallucinations, delayed care, medication errors
2. Clinicians experimenting	Using AI for admin, letters, coding, documentation (often informally)	Confidentiality breaches, ungoverned clinical use
3. Enterprise deployment	NHS-procured tools: AI scribes, workflow automation, knowledge search	Clinical safety, data governance, accountability

This framework addresses all three, but with different emphases: harm reduction for patients, governance for clinicians, and compliance frameworks for organisational deployment.

1.2 The Scale of Current Use

This is not a future problem—it is happening now:

- **230 million people** discuss health with ChatGPT weekly (OpenAI, January 2026)
- **28-30% of UK GPs** report already using AI tools in consultations or related tasks (Guardian/survey, December 2025)
- **91.8% of clinicians** surveyed globally have encountered AI-generated medical hallucinations (medRxiv, November 2025)
- **84.7% of clinicians** believe AI hallucinations are capable of causing patient harm

This 'shadow adoption' is why governance often lags behind reality. The challenge is not preventing AI use but making it safer.

1.3 What's Changed in the UK (January 2026)

OpenAI Healthcare Offerings

On 8 January 2026, OpenAI announced 'OpenAI for Healthcare' (enterprise) and 'ChatGPT Health' (consumer). Important: ChatGPT Health is not available in the UK at launch—UK patients continue using standard ChatGPT.

Anthropic Claude for Healthcare

On 11 January 2026, Anthropic announced 'Claude for Healthcare' with HIPAA-ready infrastructure. Currently US-focused with no UK-specific deployment. UK users continue using standard Claude.

NHS England Ambient Scribing Guidance

NHS England has published specific guidance for AI-enabled ambient scribing products, framing these as documentation/workflow support requiring clinical safety governance (DCB0129/DCB0160).

1.4 Why 'Just Don't Use It' Fails

Abstinence-only approaches have consistently failed in public health. The harm reduction approach recognises:

- People will continue using AI regardless of advice
- Judgmental responses drive behaviour underground
- Practical guidance on safer use is more protective than prohibition
- Maintaining engagement with healthcare services is paramount

The same applies to clinician use: if staff are using AI tools without governance, bringing this into the open is safer than pretending it isn't happening.

Section 2: Understanding the Risks

2.1 The Core LLM Failure Mode

Large language models can generate text that sounds confident, fluent, and plausible but is factually wrong. NHS England's ambient scribing guidance explicitly flags hazards including:

- Missing critical information
- Incorrect information or context
- Delayed outputs
- Unintended function introduction

In GP terms, this manifests as:

- **Hallucinated facts:** Invented past medical history, medications, allergies, or guideline claims
- **Subtle distortion:** Misrepresentation when summarising long or complex notes
- **Bias/unequal performance:** Variable accuracy across patient groups due to training data
- **Automation bias:** 'It sounded right, so I signed it'—over-reliance on AI output
- **Confidentiality breaches:** The most common real-world risk in current UK use

2.2 Hallucination Statistics

Finding	Source
Hallucination rates in clinical AI: 8-20%	Multiple studies, 2024-2025
Hallucinations more likely to be 'major' errors (44%) vs omissions (16.7%)	npj Digital Medicine, May 2025
Most dangerous hallucinations occur in 'Plan' sections	npj Digital Medicine, May 2025
GPT-4 made unsupported clinical assertions in ~30% of cases	Stanford research, 2024
64-72% of residual hallucinations stem from reasoning failures, not knowledge gaps	medRxiv, November 2025

> **Critical point:** AI delivers incorrect information with the same confident, authoritative tone as correct information. Users cannot distinguish accuracy from fabrication by how the response sounds.

2.3 Categories of Harm

For Patients Using AI at Home

- **Direct harm:** Acting on fabricated treatment recommendations or drug information
- **Delayed care:** False reassurance leading to delayed presentation of serious conditions
- **Unnecessary anxiety:** AI suggesting serious diagnoses for benign symptoms
- **Medication errors:** Adjusting medications based on AI advice
- **Erosion of clinical relationship:** Acting on AI without informing clinicians

For Clinicians Using AI Informally

- **Confidentiality breaches:** Pasting patient-identifiable data into consumer tools
- **Ungoverned international transfer:** Data flowing to non-UK servers without appropriate safeguards
- **No audit trail:** Untraceable AI involvement in clinical decisions
- **Professional liability:** GMC accountability for AI-influenced decisions
- **CQC compliance gaps:** Using ungoverned tools that would not meet inspection standards

For Organisational Deployment

- **Clinical safety failures:** Systematic errors affecting multiple patients
- **Data protection breaches:** GDPR violations at scale
- **Regulatory non-compliance:** Medical device requirements, DSPT failures
- **Accountability gaps:** Unclear liability when things go wrong

PART TWO

Patient Safety (Harm Reduction)

Section 3: Patient Guidance

This section can be adapted for patient-facing materials, leaflets, websites, and waiting room displays.

3.1 The Absolute Red Lines

There are situations where AI should never be the first response:

> **Always Call 999 or Go to A&E For:**
>
> - Chest pain or tightness
> - Sudden weakness on one side of your body
> - Difficulty breathing
> - Severe bleeding that won't stop
> - Loss of consciousness or fitting
> - Signs of stroke (FAST: Face drooping, Arm weakness, Speech difficulty, Time to call 999)
> - Severe allergic reactions (swelling of face/throat, difficulty breathing)
> - Severe abdominal pain
> - A seriously unwell child
>
> **AI will not save your life in an emergency. The NHS will.**

Never Use AI To:

- Decide whether to stop or change prescription medications
- Diagnose symptoms in children under 5
- Assess worsening symptoms—if you're getting worse, see someone
- Replace follow-up appointments for serious conditions
- Interpret cancer screening results or serious test results

3.2 Safer Ways to Use AI for Health

If you choose to use AI for health queries:

Use It for Education, Not Diagnosis

SAFER: 'What does HbA1c measure?' or 'What are common causes of headaches?'
RISKIER: 'Do I have diabetes?' or 'What's causing my headaches?'

Cross-Reference Important Information

If AI tells you something significant, check it against NHS.uk or Patient.info before acting. AI should be one source, not the only source.

Bring It to Your GP

If you've looked something up, tell your doctor. Print it out or screenshot it. We'd rather discuss what AI said with you than have you act on it alone.

Ask AI About Its Uncertainty

Add to your prompts: 'What might you be wrong about?' or 'What else could this be?' This forces acknowledgment of limitations.

Don't Trust Confident Tone

AI delivers everything—correct and incorrect—with the same authoritative voice. How confident it sounds tells you nothing about accuracy.

3.3 What AI Cannot Do

- **It cannot examine you:** No pulse, blood pressure, palpation, or visual assessment
- **It cannot see trajectory:** Is this getting better or worse? AI sees only the snapshot you describe
- **It cannot know what you forgot:** The detail you dismissed may be critical
- **It cannot weigh your specific risks:** A clinician who knows you can integrate your personal history
- **It cannot take responsibility:** When things go wrong, there's no accountability
- **It cannot follow up:** It won't call tomorrow to check you're still alive

Section 4: Guidance for Clinicians

4.1 Normalising the Conversation

Patients are already using AI. Creating a safe space for disclosure is more protective than judgment.

Replace Judgment with Curiosity

Instead of: 'You shouldn't trust Dr Google'
Try: 'Have you looked anything up about this? What did you find?'

This opens dialogue, allows correction of dangerous misunderstandings, maintains trust, and keeps you informed about patient thinking.

Create a 'Bring Your AI' Culture

Explicitly invite patients to share AI research:

> *'If you've asked ChatGPT or Claude about this, bring me what it said. I'd rather we discuss it together than have you wondering if it was right.'*

4.2 Using AI Outputs as Teaching Moments

When a patient brings AI-generated content:

1. **Take it seriously:** Read it with them, don't dismiss it
2. **Acknowledge what's correct:** 'This bit is actually pretty reasonable'
3. **Explain what's wrong and why:** 'This part is wrong because...'
4. **Explain what's missing:** 'What this doesn't account for is...'
5. **Build critical thinking:** 'Next time, here's what to watch out for...'

4.3 Red Flags Suggesting Harmful AI Use

Be alert to patterns indicating AI use may be causing harm:

- Patients self-adjusting medications based on AI advice
- Delayed presentation because AI provided false reassurance
- Excessive health anxiety driven by AI-suggested diagnoses
- Requests for investigations based solely on AI recommendations
- Distrust of clinical advice because it contradicts AI

4.4 Documentation Considerations

When patients disclose AI use or bring AI-generated content:

- Document that AI was consulted (briefly)
- Note any incorrect information that was corrected
- Document your clinical assessment and how it differed
- Record safety netting advice given

PART THREE

Practice Governance

Section 5: UK Regulatory Framework

5.1 Professional Accountability (GMC)

The GMC position is clear: doctors remain responsible for decisions they take when using AI tools, and professional standards still apply.

Key points from GMC guidance:

- Professional standards apply regardless of whether AI is involved
- Doctors must be able to justify clinical decisions made with AI assistance
- Good Medical Practice 2024 treats software/digital tools (including AI) as examples of medical devices for adverse incident reporting
- Responsibility cannot be delegated to an AI system

5.2 CQC Expectations

The CQC GP Mythbuster 109 on AI in GP services sets clear expectations:

Human Oversight

- AI must be a support tool—not a replacement for human oversight
- Evidence of monitoring and evaluation required
- Learning from errors expected, including reporting to developers and MHRA Yellow Card where relevant

Data Protection

- UK GDPR compliance including DPIAs and Records of Processing Activities (ROPA)
- Cyber security arrangements
- Data Security and Protection Toolkit (DSPT) completion

Transparency

- Patients should be told AI is being used (especially for scribes)
- Focus on transparency and ability to object
- Explicit consent may not be required for direct care, but transparency is

Staff Training

- Staff must be trained and competent to use AI tools
- Training should cover failure modes, not just 'how to prompt'

5.3 MHRA and Medical Devices

Whether an AI tool is regulated as a medical device depends on its intended purpose and how it's used.

Key considerations:

- Software intended for diagnosis, prevention, monitoring, or treatment may be a medical device
- Decision support tools may or may not qualify depending on design
- MHRA Yellow Card reporting applies for safety issues with AI medical devices
- UK government has announced a National Commission to rewrite AI regulatory framework

5.4 NHS Clinical Safety Standards

For digital tools affecting care, NHS clinical safety standards (DCB0129/DCB0160) may apply:

DCB0129 (Manufacturers)

Requirements for manufacturers of health IT systems including hazard identification, risk management, and safety case development.

DCB0160 (Deployment)

Requirements for organisations deploying health IT systems including:

- Nomination of a Clinical Safety Officer (CSO) with appropriate training
- Hazard logs and safety case documentation
- Integration with incident reporting systems

Section 6: Information Governance & GDPR

6.1 The Core 'GDPR Trap' for Clinicians

> If a clinician pastes patient-identifiable data (or easily re-identifiable detail) into a consumer chatbot on a personal account, you may have:
>
> - An unauthorised disclosure (confidentiality breach)
> - Uncontrolled international transfer risk
> - No Data Processing Agreement in place
> - No DPIA, no audit trail
> - Unclear retention and potential training use of data
>
> **This is where practices get burned.** The most common real-world risk in current UK use is confidentiality breaches through informal AI tool use.

6.2 What 'Good' Looks Like (IG Requirements)

NHS Transformation Directorate guidance establishes the shape of lawful AI deployment:

Data Protection Impact Assessment (DPIA)

A DPIA must legally be completed prior to implementing AI-based technologies. This is not optional—it's a legal requirement for processing that is likely to result in high risk to individuals.

Legal Basis for Processing

For direct care, UK GDPR special category condition 9(2)(h) will generally apply. Consent may be implied under the common law duty of confidentiality for direct care purposes.

Re-identification Risk

Be careful about 'anonymised' data. Combinations like local area + rare disease + young age can re-identify individuals. NHS guidance explicitly flags this risk.

6.3 DSPT Requirements

If your practice handles NHS patient data, the Data Security and Protection Toolkit is part of the expected assurance landscape. CQC explicitly points to DSPT completion as what they'll look for regarding data protection and third-party vendor assurance.

Section 7: Clinical Safety Requirements

7.1 When Clinical Safety Standards Apply

NHS clinical safety standards (DCB0129/DCB0160) apply when deploying digital tools that could affect patient care. This includes:

- AI scribes and ambient documentation tools
- Clinical decision support systems
- Workflow tools that affect clinical pathways
- Any AI system whose output enters the patient record

7.2 Clinical Safety Officer Role

Organisations deploying relevant tools should nominate a Clinical Safety Officer (CSO):

- Must be a registered healthcare professional
- Should have appropriate training (e.g., NHS Digital Clinical Safety training)
- Responsible for overseeing clinical safety case development
- Signs off on clinical safety documentation

7.3 Hazard Identification

Key hazards to consider for AI tools in general practice:

Hazard	Potential Harm	Mitigation
Hallucinated information	Incorrect clinical decisions	Mandatory clinician review
Missing critical information	Incomplete records, missed diagnoses	Comparison with source, spot audits
Incorrect attribution	Wrong patient, wrong history	Identity verification checks
Delayed output	Care delays	Fallback processes defined
Bias/unequal performance	Inequitable care quality	Monitoring across patient groups
Automation bias	Reduced clinical vigilance	Training, audit culture

PART FOUR

Implementation

Section 8: The Risk Ladder

Use this framework to categorise AI use cases and apply appropriate governance:

8.1 Green Zone: Generally OK with Normal Review

Use Case	Notes
Drafting generic patient information (no personal data)	Review before use; ensure accuracy
Admin templates, policies, meeting summaries	No patient-identifiable information
Training materials and educational content	Verify accuracy of clinical content
Personal reflection/learning (your own, non-identifiable)	No patient details

8.2 Amber Zone: Governance Required

Use Case	Requirements
Summarising long letters with identifiers removed	Must remove ALL identifiable info first; approved tool preferred
Creating consultation note drafts	Line-by-line clinician review mandatory; approved enterprise tool only
Translating patient-facing text	Review for nuance; consider professional translation for complex/critical content
Drafting referral letters	Full review; use approved tool with DPA if patient data involved
Differential diagnosis prompts	Thinking support only; never sole basis for decisions

8.3 Red Zone: Do Not Do in Consumer Tools

Use Case	Why Not
Pasting identifiable patient history asking for diagnosis	Confidentiality breach; uncontrolled data transfer; no audit trail
Asking AI to propose prescribing for 'this patient'	Confidentiality + clinical risk; no professional accountability
Letting AI output enter record without verification	Hallucination risk; CQC non-compliance; professional liability
Using AI for safeguarding concerns	High-risk decisions require human judgment only
Mental health crisis assessment via AI	Safety-critical; requires human relationship and judgment

The red zone is where confidentiality breaches and clinical risk cluster. This is not about AI being bad—it's about consumer tools not having appropriate safeguards for clinical use.

Section 9: Implementation Checklist

A prioritised checklist for GP practices implementing AI governance:

9.1 Immediate Actions (Do This Week)

1. Stop Shadow AI

Write a one-page practice AI acceptable use statement covering:

- No patient-identifiable data in consumer tools (personal accounts)
- Approved tools only (once assessed)
- What counts as identifiable (rare disease + location + age, etc.)

This alone prevents most real-world confidentiality mishaps.

2. Communicate the Policy

- Brief all clinical staff
- Include in induction for new staff
- Display in staff areas

9.2 Short-Term Actions (This Month)

3. Start with Low-Risk Uses

Trial AI for:

- Drafting non-clinical policies/SOPs
- Patient leaflets that you fully review
- Summarising your own non-identifiable reflections

4. Patient Communication

- Add guidance to practice website about safer AI use
- Display waiting room information (see Appendix A)
- Consider including in new patient packs

9.3 Medium-Term Actions (This Quarter)

5. If Considering AI Scribe or Documentation Tool

Treat it like a clinical system:

- Assign a Clinical Safety Officer
- Maintain hazard log / safety case mindset
- Define review processes and failure modes
- Use NHS England ambient scribing guidance as your framework

6. Information Governance Package

Minimum requirements:

- Complete DPIA (legally required—see Appendix C for prompt questions)
- Define controller/processor roles
- Establish Data Processing Agreement with vendor
- Update privacy notices and transparency materials
- Check vendor posture against DSPT expectations

7. Human Oversight Mechanisms

Build in:

- Mandatory clinician sign-off for anything entering the record
- Spot audits (e.g., 10 notes/week) looking for omissions/distortions
- Incident reporting pathway including MHRA Yellow Card where relevant

9.4 Ongoing Requirements

8. Staff Training

Training should include:

- Hallucinations and AI overconfidence
- Data minimisation and re-identification risk
- When to stop using the tool and escalate
- How to report concerns and incidents

9. Patient Transparency

Especially for scribes:

- Tell patients you're using AI (transparency, not theatre)
- Offer ability to object
- Document patient informed/objection status

10. Continuous Monitoring

- Regular review of AI tool performance
- Track incidents and near-misses
- Monitor regulatory updates (ICO, MHRA, NHS England)
- Annual DPIA review

Appendix A: Patient-Facing Materials

Ready to adapt for waiting room posters, website content, or patient leaflets.

Using AI for Health Questions?
Here's How to Stay Safe

We know many people ask ChatGPT, Claude, or Google about symptoms before (or instead of) seeing a doctor. That's understandable—it's quick, it's private, it's available at 3am when you're worried.

But AI has real limitations that could put you at risk.

AI cannot examine you. It can't feel your pulse, check your blood pressure, or look at a rash properly.

AI doesn't know your full history. Even if you tell it things, it doesn't know what you've forgotten.

AI can be confidently wrong. It sounds authoritative whether it's right or wrong. You can't tell the difference.

If You're Going to Use AI:

- ✓ Use it to understand conditions, not to diagnose yourself
- ✓ Check anything important against NHS.uk
- ✓ Bring what you found to your GP—we'd rather discuss it with you
- ✓ Never delay seeking care because AI said it's 'probably nothing'

Always Call 999 or Go to A&E For:
Chest pain • Sudden weakness on one side • Difficulty breathing • Severe bleeding • Loss of consciousness • Signs of stroke • Severe allergic reactions • Severely unwell child

AI won't save your life in an emergency. The NHS will.

Appendix B: AI Acceptable Use Policy Template

Adapt this template for your practice. This is a starting point, not legal advice.

[PRACTICE NAME]
Artificial Intelligence Acceptable Use Policy

Version 1.0 | [Date] | Review Date: [Date + 12 months]

1. Purpose

This policy establishes rules for the safe use of artificial intelligence tools by staff at [Practice Name], protecting patient confidentiality and ensuring compliance with regulatory requirements.

2. Scope

This policy applies to all staff using AI tools including but not limited to ChatGPT, Claude, Gemini, Copilot, and any AI-enabled features within clinical or administrative software.

3. Approved vs Consumer Tools

Approved Tools: [List any approved enterprise tools here]

Consumer Tools: Personal accounts on ChatGPT, Claude, Gemini, etc. These may only be used in accordance with Section 4.

4. Absolute Prohibition: No Patient-Identifiable Data

Staff must NEVER enter patient-identifiable information into consumer AI tools. This includes:

- Names, NHS numbers, dates of birth, addresses
- Combinations that could identify someone (rare disease + age + location)
- Any information from the clinical record
- Photos, scans, or documents containing patient information

Breach of this rule may constitute a data protection incident and will be treated as a serious matter under practice disciplinary procedures.

5. Permitted Uses (Consumer Tools)

Consumer AI tools may be used for:

- Drafting generic templates, policies, or educational materials
- Research and learning (no patient-specific queries)
- Administrative tasks with no patient-identifiable content

All outputs must be reviewed for accuracy before use.

6. Clinical Decision-Making

AI tools must not be the sole basis for clinical decisions. Clinical judgment remains with the registered professional. GMC accountability applies to all decisions regardless of AI involvement.

7. Incident Reporting

Report immediately if:

- Patient-identifiable data has been entered into a consumer tool
- AI output containing errors has entered the clinical record
- AI tool behaviour causes concern about patient safety

Report to: [Practice Manager / Caldicott Guardian / named person]

8. Review

This policy will be reviewed annually or when significant changes occur in AI tools, regulatory guidance, or incident patterns.

Approved by: _____ Date: _____

I have read and understood this policy:

Staff member: _____ Date: _____

Appendix C: DPIA Prompt Questions

Use these questions when completing a Data Protection Impact Assessment for AI tools. This is not a complete DPIA template—consult ICO guidance and your DPO.

Nature of Processing

1. What personal data will the AI tool process?
2. What special category data (health data) is involved?
3. How will data flow to/from the AI system?
4. Where is data processed and stored geographically?
5. Is data used for AI model training? (Check vendor terms)
6. What is the data retention period?

Necessity and Proportionality

7. What is the lawful basis for processing? (For direct care, likely GDPR 6(1)(e) and 9(2)(h))
8. Is AI processing necessary for the stated purpose?
9. Could the purpose be achieved with less data or less intrusive means?
10. How does this processing relate to the common law duty of confidentiality?

Risks to Individuals

11. What are the risks if data is breached or misused?
12. What are the risks from inaccurate AI outputs?
13. Are there risks of discrimination or bias?
14. What happens if the AI system fails or is unavailable?
15. Could individuals be re-identified from supposedly anonymised data?

Mitigation Measures

16. What security measures does the vendor provide?
17. What contractual protections are in place (DPA)?
18. How will human oversight be maintained?
19. What audit and monitoring will occur?
20. How will incidents be detected and reported?
21. What staff training will be provided?

Transparency

22. How will patients be informed about AI use?
23. How can patients exercise their rights (access, objection)?
24. Is the privacy notice updated to reflect AI processing?

Vendor Assessment

25. Does the vendor have appropriate certifications (ISO 27001, SOC 2)?
26. Does the vendor support UK GDPR compliance?
27. What is the vendor's incident response process?
28. Does the vendor meet DSPT-relevant requirements?

Appendix D: Key UK Regulatory Sources

Essential references for AI governance in UK general practice.

NHS England

- **Guidance on AI-enabled ambient scribing:** england.nhs.uk/long-read/guidance-on-the-use-of-ai-enabled-ambient-scribing-products-in-health-and-care-settings/
- **Digital clinical safety assurance (DCB0129/DCB0160):** england.nhs.uk/long-read/digital-clinical-safety-assurance/

NHS Transformation Directorate

- **AI and Information Governance:** transform.england.nhs.uk/information-governance/guidance/artificial-intelligence/

Care Quality Commission

- **GP Mythbuster 109 - AI in GP Services:** cqc.org.uk/guidance-providers/gps/gp-mythbusters/gp-mythbuster-109-artificial-intelligence-gp-services

General Medical Council

- **AI and innovative technologies:** gmc-uk.org/professional-standards/learning-materials/artificial-intelligence-and-innovative-technologies

MHRA

- **Software and AI as a medical device:** gov.uk/government/publications/software-and-artificial-intelligence-ai-as-a-medical-device
- **Yellow Card reporting:** yellowcard.mhra.gov.uk

Information Commissioner's Office

- **AI and data protection guidance:** ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/
- **AI risk toolkit:** ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/guidance-on-ai-and-data-protection/ai-and-data-protection-risk-toolkit/

NHS Digital / NHS England

- **Data Security and Protection Toolkit:** dsptoolkit.nhs.uk

Appendix E: Glossary

AI Hallucination: When an AI system generates confident, authoritative-sounding information that is factually incorrect or fabricated.

Ambient Scribing: AI-powered documentation tools that listen to consultations and generate clinical notes.

Automation Bias: The tendency for humans to over-rely on automated systems and accept their outputs without sufficient scrutiny.

Caldicott Guardian: Senior person in NHS organisations responsible for protecting patient information confidentiality.

Clinical Safety Officer (CSO): Registered healthcare professional responsible for ensuring clinical safety of health IT systems.

Consumer AI Tools: General-purpose AI chatbots (ChatGPT, Claude, Gemini) accessed via personal accounts without enterprise agreements.

DCB0129: NHS clinical safety standard for manufacturers of health IT systems.

DCB0160: NHS clinical safety standard for organisations deploying health IT systems.

DPA (Data Processing Agreement): Contract between data controller and processor establishing GDPR-compliant processing terms.

DPIA (Data Protection Impact Assessment): Assessment required before processing likely to result in high risk to individuals.

DSPT (Data Security and Protection Toolkit): NHS self-assessment tool for data security and information governance.

Enterprise AI Tools: AI products with business agreements, DPAs, and governance features suitable for organisational use.

GDPR Article 9(2)(h): Legal basis for processing special category health data for medical diagnosis and healthcare provision.

Harm Reduction: Public health approach accepting that risky behaviours will continue while focusing on minimising associated harms.

Human-in-the-Loop: Design principle requiring human oversight and approval for AI-generated decisions.

LLM (Large Language Model): The underlying AI technology behind ChatGPT, Claude, Gemini and similar chatbots.

MHRA: Medicines and Healthcare products Regulatory Agency—UK regulator for medical devices including some AI software.

Re-identification: The ability to identify individuals from supposedly anonymised data through combination of attributes.

Safety Netting: Clinical practice of providing patients with guidance on warning signs and when to return.

Shadow AI/IT: Use of technology tools by staff without formal organisational approval or governance.

Yellow Card: MHRA reporting system for adverse incidents with medicines and medical devices.

About This Document

This framework combines harm reduction principles for patient safety with UK regulatory requirements for practice governance. It draws on:

- Published guidance from NHS England, CQC, GMC, MHRA, and ICO
- Research on AI hallucinations in healthcare settings
- Clinical experience in UK general practice
- Public health harm reduction methodology

This document is intended as practical guidance and a starting point for local adaptation. It does not constitute legal advice and should be reviewed against current regulatory requirements, which continue to evolve.

AI Acknowledgment: This framework was developed collaboratively using multiple AI tools. Claude (Anthropic) drafted the initial content, ChatGPT (OpenAI) provided critical analysis and refinements, and the whole was reviewed by Dr Pasupathi to ensure clinical accuracy and practical applicability for UK general practice.

Contact

Dr Krishnan Pasupathi
feedback@aryashhealth.com
aryash.health