

ARTIFICIAL INTELLIGENCE IN HEALTHCARE GUIDELINES

Version 2.0

Published March 2026



MINISTRY OF HEALTH
SINGAPORE



Contents

1. Foreword	4
2. Introduction	5
2.1. Objectives	5
2.2. AIHGle's Focus	6
3. Ethical Principles	8
4. Setting Clear Responsibilities between Developers, Deployers, and Users	11
5. Developers: Healthcare AI Developers	14
5.1. Compliance to Regulatory Requirements	14
5.2. Managing AI Solutions Through a Total Product Lifecycle (TPLC) Approach	16
5.3. Planning, Design, and Development	17
5.4. Model Evaluation	19
5.5. Post-Market Surveillance and Maintenance	19
5.6. End of Life	20
5.7. User Communications	20
6. Deployers: Healthcare Organisations	21
6.1. Organisational Governance and Oversight	21
6.2. Risk Assessment Framework	23
6.3. Pre-Deployment Testing and Validation	24
6.4. Deployment	25

6.5. Staff Training	26
6.6. Periodic Monitoring	26
6.7. Response to Adverse Events	27
6.8. End of Life	28
6.9. Patient Communications	28
7. Users: Healthcare Professionals	30
7.1. Deployment	30
7.2. Staff Training	31
7.3. Periodic Monitoring	32
7.4. Response to Adverse Events	32
7.5. Patient Communications	32
8. Emerging Developments	33
8.1. AI Solutions with Continuous Learning Capabilities	33
8.2. Generative AI	34
8.3. Direct-to-Consumer (DTC) AI Applications	36
9. Other Relevant Legislation and Guidelines	37
10. Acknowledgements	41

Foreword

Artificial Intelligence is rapidly transforming healthcare – from advanced diagnostic imaging to clinical voice-to-text solutions. When applied responsibly, AI enhances health systems, making them more resilient, efficient, and responsive to patient needs. Importantly, AI empowers our healthcare professionals and augments their clinical decision-making, in enhancing care delivery.

We are committed to upholding an agile governance framework that has patient safety at its core, whilst enabling innovation. The responsible development and deployment of AI in healthcare is a shared endeavour – one that calls for close collaboration among policymakers, healthcare professionals, technology developers, and patients.

That is why we updated the *Artificial Intelligence in Healthcare Guidelines (AIHGle 2.0)* in close partnership with representatives from across our healthcare family and government agencies, drawing on lessons from real-world deployment. We extend our gratitude to the Academy of Medicine Singapore, College of Family Physicians Singapore, Synapse, Infocomm Media Development Authority, the Personal Data Protection Commission, and various professional boards and associations, for their partnership in developing AIHGle 2.0.

AIHGle 2.0 addresses new challenges that have emerged as AI capabilities advance and become more deeply integrated into clinical practice. We will continue to refine these guidelines to keep pace with the evolving AI developments.

Together, we can shape a future where innovation and trust go hand in hand, and where AI truly serves our patients and strengthens our healthcare system and professionals.



Professor Kenneth Mak
Director-General of Health
Ministry of Health




Adjunct Professor (Dr) Raymond Chua
Chief Executive Officer
Health Sciences Authority

2 Introduction

2.1. Objectives

- 2.1.1. Artificial Intelligence¹ (“AI”) has the ability to enhance healthcare through improved efficiency, accessibility, quality, and affordability. To realise the potential of AI, it is important to manage risks in the development, deployment, and use stages, (**Figure 1**) – which may compromise patient care outcomes and undermine users’ confidence in the use of AI.
- 2.1.2. This Artificial Intelligence in Healthcare Guidelines (“**AIHGle**”) aims to **ensure patient safety and enhance trust in the use of AI in healthcare**, in line with Singapore’s national strategy to build a trusted and responsible AI ecosystem. AIHGle provides a consolidated set of recommendations and good practices for **healthcare AI developers**, **healthcare AI deployers**: healthcare organisations, and **healthcare AI users**: healthcare professionals. It complements prevailing legislation that governs these stakeholder groups.

Figure 1: Examples of Risks in the Development, Deployment, and Use Stages




¹ Refers to machine-based systems that are designed to function with varying levels of autonomy and adaptiveness after deployment and infer from the input received to generate output. OECD (2019) *Recommendation of the Council on Artificial Intelligence*.

2.2. AIHGle's Focus

- 2.2.1. AI is broad and spans from simple rule-based decision trees to more complex Machine Learning (ML) algorithms that enables computers to analyse data to discern patterns and predict future output, and Deep Learning (DL) algorithms such as Generative AI (GenAI) (**Figure 2**).
- 2.2.2. **While broadly applicable to all AI, AIHGle is targeted at the more complex subset of AI solutions that employ ML or DL algorithms**, which have amplified risks due to their complexity, opacity, and scalability.
- ML solutions include fixed algorithms² and continuous learning models. The latter update their algorithms based on new data encountered during deployment, potentially improving performance over time. However, it also presents added risks, like model drift where the model's performance degrades over time due to changes in data, as well as risks of personal data disclosure after input data is ingested by the model. This is also addressed in the developers' section, with guidance to comply with applicable personal data protection laws and to use Privacy Enhancing Technologies (PETs) (see **Chapter 5.3.3**).
 - DL has transformed healthcare analytics with artificial neural networks to analyse large-scale data and generate better predictions for screening or diagnostic. However, it also poses additional challenges given the sensitivity of health data and possible impact to patient care outcomes. Some concerns include issues with bias, hallucination, and data disclosure.

Figure 2: Relationship between AI, ML, and DL




- 2.2.3. In healthcare settings, we have identified three broad categories of AI use cases – Clinical, Clinical-Ops, and Ops. These are defined in **Table 1** and the examples serve to illustrate the three AI use cases. Please refer to **Chapter 5.1** for the definition and detailed explanation of Medical Devices (MD) and non-Medical Devices.

² Fixed algorithm ML solutions maintain a constant model after initial training and validation.

2.2.4. AIHGle is targeted at the first two categories of AI use cases, Clinical and Clinical-Ops, which are considered to have an impact on patient care outcomes, directly or indirectly. For Ops use cases, there are national, sector-agnostic guidelines (e.g. Infocomm Media Development Authority’s (IMDA) Model AI Governance Frameworks) for both traditional AI and GenAI which healthcare organisations can refer to instead. Researchers conducting AI-related research or using AI in human biomedical research should also refer to the applicable laws or regulations and relevant resources (e.g. Human Biomedical Research Act 2015 (HBRA), report on Ethical Use of Big Data and Artificial Intelligence In Biomedical Research published by the Bioethics Advisory Committee in 2025).

Table 1: Applications of AI in healthcare

		Definition	Examples
 Clinical	Healthcare AI used to support clinical work.	Where AI supports clinical decisions or judgment and has direct impact on patient care outcomes. Clinical decisions include diagnosis, monitoring, treatment or alleviation of a medical condition.	• AI software used to support doctors in identifying suspicious areas for cancer from chest X-ray images and making a diagnosis. • AI software used to support pharmacists in risk stratification of patients to determine appropriate levels of pharmaceutical care.
 Clinical-Ops		Where AI is part of a clinical workflow but does not directly impact clinical decision or judgment.	• AI software used to transcribe doctor’s consultation from speech-to-text and summarise into case notes for doctors’ review. • AI software used to conduct medication counselling and provide standardised education on medication use.
 Ops	Other types of AI used in healthcare settings by healthcare organisations.	Where AI is part of an operational process and is usually sector-agnostic.	AI software used to support scheduling or inventory management.



Ethical Principles

- 3.1. We must ensure that AI serves society's best interests without compromising human dignity or causing undue patient harm, and the physical healthcare professional-patient relationship remains valuable. Ethical principles are fundamental in this regard. These principles guide responsible development, deployment, and use of AI solutions.
- 3.2. AIHGle recognises and seeks to give effect to seven ethical principles – (i) safety, (ii) fairness, (iii) transparency, (iv) explainability, (v) robustness, (vi) security and data protection, as well as (vii) AI alignment to human values or goals. These principles, adapted from various authoritative sources on AI ethics and governance, fundamentally align with the medical ethics. They reflect the commitment to patient safety and high professional ethical standards that have long guided the healthcare profession. **Table 2** elaborates on these principles and translates them into the responsibilities of the key stakeholders.



Scan here for the Singapore Infocomm Media Development Authority's (IMDA) *Model AI Governance Framework*.



Scan here for the Association of South-East Asian Nations' (ASEAN) *Guide on AI Governance and Ethics*.



Scan here for the World Health Organisation's (WHO) *Ethics and Governance of AI for Health*.

Table 2: Ethical principles for responsible healthcare AI and examples of how these may be operationalised by various stakeholders

Principles	Stakeholders		
	 Developers	 Deployers	 Users
	Healthcare AI Developers	Healthcare Organisations	Healthcare Professionals
<u>1. Safety</u> Patient safety must be prioritised and must adhere to the fundamental medical principle of non-maleficence.	Implement rigorous testing protocols, incorporate fail-safe features or construct value hierarchies in the design phase to prevent patient harm and medical errors.	Implement systems for robust monitoring, periodic evaluation of patient safety outcomes and avenues for reporting of adverse events by all stakeholders.	Understand AI's limitations and work with deployers to establish the appropriate contingency workflows to ensure uninterrupted and safe patient care if systems fail.
<u>2. Fairness</u> Healthcare AI should be made accessible and should not result in discriminatory or unjust clinical outcomes for patients across different demographic groups.	Use representative and accurate training datasets, validate AI solutions and ensure fair access to AI solutions to avoid discriminatory outcomes.	Validate AI solutions before deployment to their specific context and ensure equitable patient access and regular monitoring to address unintended biases.	Understand AI limitations and potential biases and apply clinical judgment to mitigate potential bias and prevent discriminatory clinical outcomes.
<u>3. Transparency</u> Every stakeholder should provide sufficient relevant information for each stakeholder in the AI lifecycle to support informed decision-making and users should be informed of interaction with AI, aligning with the medical principle of autonomy.	Be transparent about datasets, algorithms, evaluation process, limitations and potential risks. Wherever possible, provide clear front-end labels in user interfaces to highlight where recommendations, information, or content is generated by AI.	Promote transparency with policies for patient communications about AI, where appropriate.	Be proficient in the organisations' frameworks on patient communications about AI and know the information needed for patients to make informed decisions.

Principles	Stakeholders		
	 Developers	 Deployers	 Users
	Healthcare AI Developers	Healthcare Organisations	Healthcare Professionals
<u>4. Explainability</u> AI solutions should have appropriate levels of explainability tailored to user needs and associated risks.	Ensure AI output is interpretable and help users understand the AI decision-making process.	Ensure AI solutions meet users' expectations on explainability.	Understand the broad logic and limitations of AI solutions to interpret and validate output.
<u>5. Robustness</u> AI solutions should perform consistently under different deployment conditions.	Test rigorously for consistent output and model performance to ensure robust and reliable AI behaviour across all operational conditions.	Apply standardised evaluation and routinely assess accuracy and consistency of AI solutions to maintain system performance and reliability.	Check for consistency in output across time and different patient profiles and provide appropriate feedback to the deployer.
<u>6. Security & Data Protection</u> AI solutions should be secure-by-design, to maintain confidentiality, integrity and availability of the AI system itself and the organisation.	Ensure cyber and data security best practices across the lifecycle, including secure AI solutions from threats and enhance data protection.	Establish processes for data sharing and data protection in compliance with prevailing laws and data governance framework for data collected by AI solutions.	Adhere to prevailing best practices to protect patient data and confidentiality.
<u>7. AI Alignment to Human Values or Goals</u> Alignment to human values and goals to shape AI objectives and adjust AI functions.	Develop AI solutions with explicit consideration of human values, incorporating ethical frameworks and societal norms into the development process.	Establish clear guidelines on how AI solutions align with medical ethics, their organisation's mission, values, and patient care objectives. Assess the impact on clinical outcomes and patient welfare to guide deployment decisions, e.g. deployment of AI should not compromise clinical outcomes and patient welfare.	Maintain their role as primary decision-makers, using AI as an assistive tool that supports rather than replaces professional judgment.

Setting Clear Responsibilities between Developers, Deployers, and Users

4.1. The development and deployment of AI in healthcare is a collaborative effort that involves various stakeholders with different scope of roles and responsibilities depending on the extent of involvement at each stage. The details are in **Chapters 5, 6, and 7** respectively for each stakeholder. In short, the key responsibilities of each stakeholder are summarised below:



a. **Developers** who develop, integrate and/or maintain healthcare AI solutions.³ AI developers should familiarise themselves with the regulatory requirements⁴ and understand how these apply to the AI solutions they develop. Beyond the technical development and regulatory compliance, developers should support the users, by being transparent about the intended use and limitations of the AI as well as providing post-deployment user support.



b. **Deployers** who deploy AI solutions to augment healthcare service delivery⁵ and are licensed under the Healthcare Services Act 2020 (HCSA). They should maintain operational readiness, ensure healthcare professionals are adequately qualified and trained in appropriate AI use, implement effective governance frameworks focusing on risk management and quality assurance, as well as report any incidents that could impact care delivery. As they adopt AI solutions, they must also ensure that the standards of clinical outcomes and welfare of the patients and professionals under their purview are maintained or improved compared to if AI were not deployed.



c. **Users** who use AI solutions to support Clinical and Clinical-Ops work. They include clinicians, dentists, nurses, pharmacists, and allied health professionals, who may be registered with the relevant Professional Boards, e.g. under the Medical Registration Act 1997 (MRA), the Nurses and Midwives Act 1999 (NWA), the Allied Health Professions Act 2011 (AHPA), etc. All healthcare professionals, registered or not, are responsible for delivering safe and effective patient care and ensuring that they comply with the relevant professional obligations and align with professional standards.⁶ While AI solutions offer new capabilities for clinical work, they function much like other medical tools or devices (e.g. insulin monitoring devices or robotic surgical devices), as aids that enhance, not replace, professional judgment. In both Clinical and Clinical-Ops settings, AI should

³ This includes medical devices that are regulated under the Health Products Act 2007 (HPA), or non-medical devices used as part of clinical workflows.

⁴ Requirements include HSA's HPA for medical devices, PDPC's Personal Data Protection Act and Advisory Guidelines and IMDA's Model AI Governance Frameworks.

⁵ For AIHGLE, we will focus on healthcare services that are licensable under HCSA.

⁶ The professional obligations are set out by the respective councils and professionals boards, such as the Singapore Medical Council (SMC), the Singapore Dental Council (SDC), the Singapore Pharmacy Council (SPC), the Singapore Nursing Board (SNB), the Allied Health Professionals Council (AHPA), and the various Specialists Accreditation Boards (SAB).

augment and empower healthcare professionals, enabling them to deliver safer and more effective care. In this regard, human oversight is mandatory for all Clinical and Clinical-Ops use of AI. It also does not alter the responsibility of healthcare professionals to ensure appropriate and safe clinical care for patients.

- 4.2. The extent of responsibilities and how these should be carried out vary across the AI lifecycle. We illustrate some examples of the different responsibilities for each stakeholder below:



- a. **Plan and Design:** Developers should collaborate with deployers and/or users to define AI's intended use cases, validate the relevance and fairness of training datasets and set performance benchmarks. AI solutions should be fit-for-purpose and designed with consideration of deployers' requirements to integrate into existing technical infrastructure and clinical workflows. They should take appropriate steps to ensure compliance with applicable personal data protection laws and principles when training datasets comprise personal data, including whether the inclusion is reasonably appropriate and necessary for training purposes.

- b. **Develop:** Developers are responsible for the development of the AI solution. Developers should document the AI's development processes, standards compliance and technical specifications (e.g., relating to performance, security, personal data protection), and provide evidence demonstrating solution is safe for intended clinical use. Users are encouraged to guide the selection of clinically relevant, evidence-based, high-quality, current data, label data and assess data quality. If data collected by the deployers are used, deployers should ensure that any data ingested abides by the relevant data security and personal data protection requirements.



- c. **Evaluate:** Developers should validate the AI's performance according to the intended use and its functionalities. Developers should perform testing and validation, to ensure that the AI meets the required safety and security standards and specifications set out in the "Develop" phase. Deployers should conduct real-world validation to ensure the AI's performance within the deployed setting meets clinical requirements and ensures that patient outcomes meet or exceed existing care standards, compared to when AI was not deployed. Where necessary, the evaluation should involve deployers and users. Examples of additional potential risk mitigation measures include guardrails, which detect and filter inappropriate Large Language Model (LLM) input and output through rule-based constraints.

- d. **Deploy:** Deployers should establish robust governance to ensure safe and responsible deployment. This would include evaluating and approving deployment, adapting clinical workflows and protocols, as well as training staff before deployment of suitable AI solutions. Users should adhere to organisational protocols and professional standards when using AI solutions, and always maintain human judgment particularly in Clinical and Clinical-Ops settings.





- e. **Monitor:** Developers and deployers should monitor AI performance throughout deployment and report any adverse events to relevant regulatory authorities in accordance with prevailing regulations. Deployers should also establish internal incident management processes to encourage user feedback or reporting to organisations' governance committees and be ready to suspend AI deployment in the case of serious adverse events.

- f. **Review:** Developers and deployers should implement periodic performance assessments and safety reviews to optimise AI solutions as necessary. Reviews should also include feedback from users and clinically relevant, evidence-based, high-quality findings and recommendations from latest medical advancements, to improve workflows and clinical relevance and safety. As part of the review and maintenance process, AI developers should ensure that they properly document and take a snapshot of the AI model at each stage, before any changes are implemented, so that the changes can be reversed as necessary.



- g. **End of life:** Developers and deployers should establish an exit management plan to decommission AI solutions in accordance with data protection standards. Deployers and users would need to support such transitions to ensure care continuity.

- 4.3. As demonstrated above, stakeholders have different responsibilities and levels of involvement across the lifecycle. It is recommended to document and formalise their specific roles and responsibilities. For instance, developers and deployers may articulate these in Service Level Agreements (SLA) and deployers may document the responsibilities of users via Standard Operating Procedures (SOPs). These aim to ensure:

- a. **Clear accountability across the AI lifecycle:** Clear role definitions help identify the stakeholder accountable for each step of the process and ensure that they can fully understand their role(s) and specific responsibilities and effectively perform the same.



- b. **All responsibilities are appropriately tasked:** Responsibilities should be assigned to the stakeholder best equipped to handle them, and transparency in the role(s) and responsibilities will also help identify gaps and facilitate more informed development, deployment and use of AI.



- c. **Mitigation of risks and compliance with regulations:** A defined scope will help stakeholders identify potential risks and liability issues and ensure adequate safeguards and mitigation measures. This will also help in compliance with relevant regulations, especially where stakeholders are in different jurisdictions and different laws and regulations apply.



5 Developers: Healthcare AI Developers

This section will focus on the responsibilities of AI developers. Where relevant, it also references the Health Products Act 2007 (HPA) and its associated regulations and guidelines. As with other sections of AIHGle, the recommendations focus on AI solutions used in Clinical and Clinical-Ops settings.



5.1. Compliance to Regulatory Requirements

- 5.1.1. Developers are responsible for understanding the prevailing regulatory requirements and ensuring compliance for the AI solutions they develop. In general, devices that are intended for diagnosis, monitoring, treatment or alleviation of any medical condition are regulated as Medical Devices (MD) and are subject to regulatory controls for MDs under the HPA.⁷
- 5.1.2. To aid developers in understanding whether their AI solutions are subject to HPA regulations, we have provided some examples in **Table 3** below:



Scan here to see the definition of MD stipulated in the First Schedule of the HPA on HSA's website.

⁷ HSA published the Guidelines on Risk Classification of Standalone Medical Mobile Applications and Qualification of Clinical Decision Support Software (CDSS) for developers to determine the risk classification of Standalone Medical Mobile Applications that are Medical Devices (SaMD).

Table 3: Examples to show the distinction between Medical Devices and non-Medical Devices⁸

Medical Device (i.e. regulated by HPA)	Non-Medical Device (i.e. not regulated by HPA)
1. Software that uses real-time transcription of speech to text during medical consultation. The software provides medical diagnosis or treatment recommendation.	1. Software that uses for real-time transcription of speech to text during medical consultation and/or language translation. The software does not provide medical diagnosis or treatment recommendation.
2. Large Language Models (LLMs) that use patient data to provide clinical recommendations to clinicians. The software provides recommendations for medical diagnosis and evaluates patient suitability for treatment.	2. LLMs that explain side effects of medication prescribed to patients or take consent from patients. The software does not provide medical diagnosis or treatment recommendation.
3. Software that segments vessels and calculates the vessel area through intravascular imaging.	3. Conversational chatbots recommending general lifestyle recommendations (e.g. diet advice, exercise routines, sleeping tips).
4. Software that processes patient clinical parameters, medical history, and surgical variables to calculate individual risk scores and potential complications, enabling healthcare professionals to make informed decisions about surgical interventions and perioperative care management.	4. Software that analyses historical hospital admission patterns and operational data to forecast expected bed occupancy rates and staffing requirements.
5. Software that analyses patient-specific clinical data and laboratory results to provide recommendations for medication combinations and dosage adjustments, supporting clinical decision-making in patient treatment.	5. Software that analyses patient clinical parameters and medical history to calculate fall risk scores for individuals and alert healthcare professionals to individuals who may need additional supervision in the ward.
<u>Direct-to-consumer</u>	<u>Direct-to-consumer</u>
6. Wearable devices that monitor breathing patterns during sleep and use AI to detect signs of moderate to severe sleep apnoea and provide notification to user.	6. Wearable devices incorporating AI that track and analyse physical activity, sleep patterns, and general health metrics for general wellness and lifestyle management. ⁹

5.1.3. MD developers should familiarise themselves with HPA regulations and GL-04: Regulatory Guidelines for software medical devices – a lifecycle approach and may refer to HSA's website for details. Developers can seek regulatory advice during MD's development phase or seek feedback on completeness of the device dossier before product registration via the Health Sciences Authority (HSA) Premarket Consultation Scheme (PMC).¹⁰



Scan here to refer to HSA's website on digital health products and regulations.

⁸ This is a non-exhaustive list of examples. Developers may consult HSA on their proposed AI solution via an online feedback form or do a self-assessment on HSA's website.


⁹ Please refer to HSA Regulatory Guidelines on Telehealth Products. Link: www.hsa.gov.sg/docs/default-source/hprg-mdb/regulatory-guidelines-for-telehealth-products-rev-2-1.pdf

¹⁰ There are two types of pre-market consults: (a) the MD Development Consultation for all MDs, which provides regulatory advice in the development stages; and (b) the MD Pre-submission Consultation for MDs of higher risk class only, which provides feedback on the completeness of the device dossier prior to product registration. Please refer to HSA's website for more information and to schedule an appointment.

5.2. Managing AI Solutions Through a Total Product Lifecycle (TPLC) Approach

- 5.2.1. The Total Product Lifecycle (TPLC) approach¹¹ is a comprehensive and integrated strategy. This encompasses thorough risk assessment, rigorous software verification and validation, change management, and traceability to facilitate transparency and accountability throughout the solution's life cycle.
- 5.2.2. While multiple frameworks exist, they share fundamental principles focusing on development, post-market surveillance, maintenance, and decommissioning, with emphasis on risk management,¹² quality assurance, and regulatory compliance to meet the appropriate standards of safety, efficacy, and reliability. The following sections will elaborate on each aspect of TPLC illustrated in **Figure 3**.

Figure 3: TPLC



¹¹ HSA published the GL-04 Regulatory Guidelines for Software Medical Devices – A Life Cycle Approach to provide clarity on the regulatory requirements for software medical device (including MLMD) in its Total Product Life Cycle (TPLC) for developers. Developers may also refer to the ISO13485, which is the international standard for quality management systems in Medical Devices. This provides a structured framework for implementing TPLC principles by establishing requirements for regulatory compliance and quality assurance processes.

¹² Please refer to HSA Regulatory Guidelines for Software Medical Devices – A Life Cycle Approach for information on regulatory requirements for risk management and cybersecurity.

5.3. Planning, Design, and Development

5.3.1. Planning, Design, and Development require multidisciplinary expertise to provide context-specific insight and experience based on the intended use and indication for use. Developers should always ensure that ethical considerations are factored in from the start and that AI solutions developed are fit-for-purpose. This includes appropriate use of high-quality and representative datasets, proper design and choice of algorithms, applying secure-by-design principles, and clear documentation of the AI solution.



5.3.2. A multidisciplinary team can ensure that risk management comprehensively consider clinical, technical, and user perspectives. This is crucial because AI has added complexities related to algorithmic bias, model drift or model decay, interpretability of models, and the dynamic nature of data-driven decision-making. Developers should consider organisational requirements for integration of AI into existing technical infrastructure and clinical workflows, with due consideration of system interoperability, data flows, and impact on clinical processes. The iterative nature of AI development requires continuous collaboration between clinicians, data scientists, and regulatory teams to refine the real-world performance of the algorithms. Meaningful involvement of users throughout the development lifecycle, from initial planning, design, and development stages, ensures that risk management strategies are practical, user-centric, and aligned with actual clinical workflows. This participatory approach helps identify potential risks from the ground up and develops more effective mitigation strategies based on real-world usage patterns.

5.3.3. Data quality and management are fundamental in the development of robust and reliable AI solutions to ensure data integrity, representativeness, and appropriate use of data.

- a. Data obtained from a reliable, accurate, current, and clinically appropriate source used for model development and validation should be representative of diverse patient demographics and clinical conditions. Appropriate representation of the target patient population, medical condition, and other relevant variables help avoid bias and ensure model generalisability. Data processing should also be reviewed and documented to manage any unintended bias or data drift.
- b. The training and test datasets should be independent of each other, considering various sources of potential dependence such as patients, data acquisition methods, temporal dependencies, feature engineering, and site-specific factors. This prevents the model's performance from being artificially inflated due to data overlap. All potential sources of dependence should be addressed and managed.

- c. Where the training and test datasets contain or comprise of personal data, appropriate steps should be taken to ensure compliance with applicable personal data protection laws and principles. These include ensuring that any collection, use, or disclosure of personal data is only used for reasonably appropriate purposes, that data subjects' consent is obtained or there are valid grounds to proceed without consent; and that reasonable security arrangements are made to prevent unauthorised access and disclosure of personal data.
- d. Developers may consider implementing PETs, which include tools or techniques such as data obfuscation, encrypted data processing, and federated analytics. These allow for processing and analysing of insights from data while protecting sensitive information.
- e. When using synthetic data, developers should refer to PDPC's Proposed Guide on Synthetic Data Generation to understand the generation techniques, potential use cases, and good practices to guard against re-identification. While synthetic data can be useful for initial development, real-world data is essential for achieving model generalisation and avoiding performance degradation when deployed in real world setting. Hence, developers should prioritise training and evaluating the model with real-world data.

Case Study 1 : Synapxe

Synapxe's Assisted Chronic disease Explanation using AI (ACE-AI) risk prediction model was trained on NEHR data of about 3.6 million derived patient records (2015-2021), covering diverse ethnicities (e.g., Chinese, Indian, Malay, Caucasian) and health conditions (e.g., diabetes, hypertension) to reduce unintended biases.

- 5.3.4. Developers should also apply secure-by-design principles when developing the AI solution. Developers should refer to CSA's Guidelines and Companion Guide on Securing AI Systems to help system owners secure AI throughout its lifecycle. These guidelines will help to protect AI systems against classical cybersecurity risks and novel risks such as Adversarial Machine Learning.
- 5.3.5. Developers should prepare comprehensive Instruction for Use (IFU)¹³ tailored to the intended user, e.g. healthcare professionals and/or patients. The IFU should encompass, but not be limited to, the AI solution's intended use, indications for use, intended patient population, user and use environment, input and corresponding output of the model, limitations, user interface interpretation, clinical workflow integration, and performance. The IFU must also include prominent warnings about the limitations of the device, potential risks, and situations where the solution should not be used.



¹³ Please refer to GN-23-R2 Guidance on Labelling for Medical Devices (September 2022).

- 5.3.6. Developers should maintain thorough documentation of the AI solution, including data characteristics, sources, version history, change control procedures, and the rationale behind any modifications. The documentation ensures transparency and traceability through software versioning, facilitates troubleshooting and future updates, promotes continuity in the development process. To achieve effective traceability, it is crucial to document and link key artifacts, including requirements, design specifications, test cases, and validation results, throughout the development process.

5.4. Model Evaluation

- 5.4.1. Model evaluation involves assessing the performance and generalisability of the model.¹⁴ Developers should test the AI solution for accuracy and reliability to ensure safety.
- 5.4.2. Developers should evaluate how the model performs by comparing its output against industry benchmarks or reference standards to measure performance metrics. These metrics should be compared to those obtained during pre-clinical validation to assess how well the model generalises to new, real-world data.
- 5.4.3. AI applications (e.g. DL and ML-based systems to predict pre-defined labels) should be tested with measures of performance such as sensitivity and area under the ROC curve (AUC), and their performance should be compared with the standard of care.

Case Study 2: Synapxe

Synapxe conducts retrospective validation using unseen hold-out test dataset and clinical validations to assess real-world performance, gather user feedback, and refine workflows to minimise patient safety risks. Outputs are compared with benchmarks like clinical risk calculators e.g. Framingham to assess the onset risk of specific diseases.

5.5. Post-Market Surveillance and Maintenance

- 5.5.1. Monitoring and evaluation of the AI solution's performance after deployment should be conducted periodically to detect any anomalies or unexpected behaviour. This includes establishing robust data collection (including complaints and feedback) and monitoring processes through a feedback loop for user to report any issues or concerns. This feedback loop can be a built-in function within the software or incorporated as part of the clinical workflow recommended by the developer.
- 5.5.2. For all MDs, post-market reporting of Adverse Event or Field Safety Corrective Action is a legislative requirement under the HPA.
- 5.5.3. For any changes to registered MDs (such as change to the IFU, change in manufacturing sites, etc.), the manufacturer should refer to HSA's guidance^(15 and 16) and determine if the change requires prior approval from HSA before implementation. Developers should note that certain changes made to an MD may require a new pre-market application. For example, a change to the intended use will require a new product registration or could alter an AI solution's classification from a non-MD to an MD. Such transitions can occur when additional functionalities that impact patient care outcomes or clinical decision-making are introduced.

¹⁴ External validation is part of the Clinical Evaluation assessment. For more information on Clinical Evaluation, please refer to the GL-04 Regulatory Guidelines for Software Medical Devices – A Life Cycle.

¹⁵ Regulatory Guidelines for Software Medical Devices – A Life Cycle Approach.

¹⁶ GN-21 Guidance on Change Notification For Registered Medical Device.

5.6. End of Life

- 5.6.1. End of life refers to the decommissioning of the AI solution. Developers should work with users to establish an exit management plan before decommissioning an AI solution. The plan should include knowledge and data transfer, archival and disposal, asset recovery, and alternative clinical processes during and after decommissioning. There should be secure disposal or destruction of all data and models in accordance with data protection standards and/or relevant rules and regulations that apply to the organisation. Examples of disposal methods include data sanitisation, software de-activation, physical destruction.

5.7. User Communications

- 5.7.1. Developers should disclose accurate information on the AI solutions to users in healthcare settings, i.e. healthcare organisations and professionals, based on the ethical principle of transparency. Information that may be disclosed includes datasets, algorithms, evaluation processes, limitations, potential risks and optimal operating contexts (e.g. intended use, users, and patient population). This helps organisations make informed buying decisions. For MDs regulated under the HPA, this refers to the IFU referenced in **Chapter 5.3.5**, which are submitted to HSA for product registration.
- 5.7.2. Wherever possible, developers should integrate instructional elements directly into the software interface to provide clear instructions on setup, use, and maintenance as well as limitations, escalation protocols (e.g. instances when users should seek additional or professional advice), and situations where the AI solution should not be used. Developers should provide information to support user training, catering to varying levels of technological proficiency and health literacy. Training should enable users to understand and integrate the AI output with overall health management. This supports responsible AI use in healthcare settings and is especially important for direct-to-consumer applications often deployed outside of usual care settings, where patients or citizens engage with these applications independently. Such applications offer unprecedented, direct access to health recommendations, but it also removes critical safety layers that exist in healthcare settings. Hence, educating users on responsible use is paramount.



Deployers: Healthcare Organisations

This section will focus on the responsibilities of deployers – i.e. healthcare organisations – in ensuring safe and responsible AI deployment. They are aligned to the responsibilities articulated in HCSA and its associated regulations and guidelines and aim to demonstrate how these responsibilities translate to AI deployment in healthcare.



6.1. Organisational Governance and Oversight

6.1.1. Organisations should set up internal governance platforms to maintain oversight of all healthcare AI solutions deployed in their organisation throughout the AI solutions' lifecycle. Such platforms should establish processes to:



a. Assess and approve¹⁷ AI solutions deployed in the organisation. This includes assessing that the AI is fit-for-purpose, performs at least as well as current standard processes,¹⁸ ensuring proper documentation of the decision to deploy the AI solutions as well as adequate deliberation on risk assessment and potential ethical implications arising from AI use.¹⁹



b. Maintain a registry of all AI solutions deployed in the organisation and monitor these. The registry should include the AI solution's risk classification. This will enable the organisation to calibrate and appropriately deploy its monitoring and risk mitigation measures, until the end-of-life.



c. Provide guidance on the responsible and safe deployment of AI solutions, including integration with existing technical infrastructure and systems as well as clinical workflows, adverse event and incident management procedures, and approach to patient communications.

¹⁷ Depending on the intended use and associated risks, this may need to be approved by Organisational Leadership. This refers to those who are responsible for the overall leadership and governance of the healthcare service and varies based on the organisational size and structure. For large healthcare organisations (i.e. hospitals, nursing homes, laboratories) this could be the Board of Directors, Clinical Director, and Chairman, Medical Board, or equivalent. For solo practitioner clinics or organisations, this could be the business owner and/or the clinical lead.

¹⁸ AI solutions deployed must at least perform as well as the current standard processes. Any performance difference should be within clinically acceptable limits.

¹⁹ GSM Association's (GSMA) AI Ethics Playbook and ASEAN Guide on AI governance and ethics.



d. Ensure that AI deployment meets prevailing cybersecurity and data protection requirements. Organisations should specify the permissible data types that may be input for each AI solution. This should take into account the data security classification framework and sensitivity framework.

- 6.1.2. Governance platforms should have the relevant clinical, operational, technical and legal knowledge to make an informed decision over the adoption and deployment of AI solution in the organisation. Where feasible, it is recommended that these platforms contain multi-disciplinary representatives.
- 6.1.3. In addition to setting up a platform to set governance policies for AI solutions deployed, organisations should strive to establish clear separation of responsibilities in AI governance. While existing units may take on these roles, three distinct functions should be maintained: (a) a governance platform; (b) a unit to handle the day-to-day implementation, monitoring, enforcement of compliance, and operational aspects of AI governance; and (c) a unit to check and audit for compliance.
- 6.1.4. Organisations may decide to procure an existing healthcare AI solution (“buy”), develop it (“build”) or co-develop it with an external developer (“co-develop”). Regardless of approach, organisations should ensure that the AI solution meets the prevailing regulatory requirements, e.g. AI-MDs should be registered with HSA under the HPA. Organisations should also ensure appropriate use of the AI solutions in accordance with the intended use defined by the developers in the IFU. The “build” or “co-develop” approach will mean that organisations take on all or some of the responsibilities as a developer. Organisations should refer to **Chapter 5** to familiarise themselves with best practices for developers.

Case Study 3: National University Health System (NUHS)

NUHS’ multi-disciplinary AI Governance Committee (AIGC) has members with technical, clinical and legal/regulatory expertise and from across different specialities.

The AIGC:

- a. Formulates and implements policies on the development and deployment of AI for clinical and non-clinical use in NUHS;
- b. Developed a registry and monitoring framework to ensure all clinical AI systems meet quality and safety standards;
- c. Monitors clinical outcomes, including complications or adverse events/issues, and have incident monitoring and management processes to ensure safety and protect patients’ best interests; and
- d. Conducts post-implementation audit of AI systems.

6.2. Risk Assessment Framework

6.2.1. In deploying AI, organisations must ensure that healthcare professionals remain in charge of patient care outcomes. Organisations should adopt a risk-based approach to determine the appropriate deployment model and governance measures proportionate to the assessed risk, taking into account:

- Potential harm to the patient's health, safety or wellbeing. The assessment should consider the level of human oversight²⁰ and the potential harm that could arise from using the AI solution, including adverse events, diagnostic errors, and compromised patient care outcomes.
- Probability of harm occurring to the patient. The assessment should consider the use case (i.e. whether deployed for Clinical or Clinical-Ops use), the AI model's maturity, the quality of training data, the validation of real-world performance, the intended use, and workflow integration.

6.2.2. Based on these considerations, an illustrative example is provided in **Table 4** below:

Table 4: Risk Assessment Framework

Broad Categories on Human Oversight	Clinical	Clinical-Ops
<u>Human-in-the-loop</u> Suggests that human oversight is active and involved, with the human retaining full control and the AI only providing recommendations or input. Decisions cannot be exercised without affirmative actions by the human, such as a human command to proceed with a given decision.	<u>Minor to moderate risk</u> ²¹ E.g. AI software in imaging systems that highlights potential abnormalities in X-ray images to support doctors in making diagnoses. Doctors retain full control and make final decision on diagnoses.	<u>Negligible to minor risk</u> ²¹ E.g. AI software that uses real-time transcription of speech to text during medical consultation and summarise into case notes for doctors' review. Doctors retain full control and make final decision on the case notes saved.
<u>Human-over-the-loop</u> Suggests that human oversight is involved to the extent that the human is in a monitoring or supervisory role, with the ability to take over control when the AI encounters unexpected or undesirable events (such as model failure).	<u>Moderate to severe risk</u> Deployers should not deploy AI solutions that make clinical decisions independently without human making the final decision. E.g. AI software in imaging systems that highlights potential abnormalities in X-ray images and make diagnoses with doctors monitoring/supervising the diagnosis recommendations. Doctors take over control and make final decision on diagnoses, only if and when AI flags unexpected events.	<u>Moderate risk</u> ²¹ E.g. AI software that educates on medication use and does not provide diagnosis or treatment recommendations. Pharmacists would monitor patient interactions with AI software and intervene where necessary. Pharmacists take over full control and make final decision on the replies, only if and when the AI flags unexpected events.

²⁰ This references the broad approaches to classify the levels of human oversight in the IMDA's Model AI Governance Framework Second Edition (2020) and OECD Framework for the Classification of AI Systems (2022).

²¹ Deployers should assess the risks, train users, and implement adequate governance measures for each AI solution to be deployed.

Broad Categories on Human Oversight	Clinical	Clinical-Ops
<u>Human-out-of-the-loop</u> Suggests that there is no human oversight over the execution of decisions. The AI evaluates input and acts upon its recommendations or output without human involvement.	<u>Severe risk</u> Deployers should not deploy AI solutions that make clinical decisions independently without human oversight/intervention. E.g. AI-driven robotic system performs simple surgical procedures without real-time human oversight/intervention.	<u>Moderate to severe risk</u> Deployers should not deploy AI solutions that make decisions independently without human oversight/intervention. E.g. AI-driven robotic system reads prescriptions, packs medication and dispenses them to patients directly without human oversight/intervention.

6.2.3. Based on the risk assessment, organisations should consider these approaches to deploying AI solutions:



- a. Deploy AI solutions, starting with less critical health situations or contexts. This reduces the risk of harm. In Clinical or Clinical-Ops settings, autonomous AI should not be deployed without human oversight.



- b. Adopt robust, high-performance, and explainable AI models. This lowers the risk of harm and also provides greater explainability and transparency to users. For black-box models, organisations should work with developers to ensure explainability and implement additional measures to detect bias or anomalies.

Case Study 4: MOH Office for Healthcare Transformation (MOHT)

To ensure safety, MOHT deploys AI models that may only escalate cases to clinicians and can never lower the severity of existing alerts. This sets a robust minimal standard of care and AI can raise additional cases to the clinician review on further action required.

6.3. Pre-Deployment Testing and Validation

- 6.3.1. Thorough testing and validation of AI solutions ensure that they are appropriate for the intended use as specified in the IFU for MDs, and any potential risks identified are mitigated before deployment in clinical settings. Unlike performance validation performed by developers as part of the development process, the implementing organisation usually performs this testing and validation. This should consider the specific context and clinical workflows in which the AI solution is incorporated in. For AI solutions that have higher risk, organisations may consider a pilot phase to assess the AI's real-world performance before mainstreaming its deployment to clinical workflows.
- 6.3.2. A pilot phase is typically done in a controlled environment, with a smaller patient population. Organisations should pay special attention to unusual scenarios or clinic-

specific contextual cues that might challenge the AI solution to reveal limitations or potential failures not immediately apparent (e.g. certain demographic groups). These insights enable organisations to design AI-augmented clinical workflows with appropriate mitigations to enhance overall safety and minimise patient risk.

- 6.3.3. There are various toolkits for testing and validation of AI solutions, such as AI Verify testing framework and software toolkit which evaluates AI solutions against internationally recognised AI governance principles. These assess the technical measures of the algorithm's performance, like resilience and fairness of the AI solution. In addition to measuring the algorithm's performance, organisations should measure clinical outcomes²² to ascertain the effectiveness of the AI solution (i.e. impact on patients when using the AI solution in their care). Organisations should minimally ensure that healthcare professionals and patients benefit from or maintain current care standards with the deployment of the AI solutions.
- 6.3.4. Organisations may also wish to consider subjecting their testing and validation results and the associated testing methodology for peer-review.²³ This provides valuable third-party validation and contributes to the broader knowledge base in healthcare AI, benefitting other organisations and improving overall standards in the field. Organisations may also refer to **Chapter 5.4** on testing requirements for developers.

Case Study 5: Health Sciences Authority (HSA)


To support the responsible deployment of AI solutions, MOH and HSA develop regulatory sandboxes as a tool to facilitate evaluation in real-world healthcare settings.

One such example is the AI-MD Exemption Sandbox, that was launched in February 2026 to allow selected public healthcare entities to deploy low-to-moderately low-risk AI-MDs across multiple institutions without full licensing and registration requirements.

This flexibility helps accelerate the deployment of beneficial AI solutions while maintaining rigorous safeguards.

6.4. Deployment

- 6.4.1. Successful AI deployment requires comprehensive change management. This encompasses technological adoption, workflow transformation and organisational culture shifts. Organisations should clearly articulate how AI will integrate into existing technical infrastructure and clinical workflows and consider appointing AI champions to facilitate AI adoption through communication and development of assessment and surveillance tools.



²² Clinical association of the outcomes of an AI solution to its intended use can be established through existing evidence (e.g. literature, original clinical research, guidelines), or generating new evidence (e.g. data analysis, clinical trials) – SaMD: Clinical Evaluation (IMDRF, 2017).

²³ Peer-reviews of AI-MDs (especially those trained on demographic data comprising e.g. racial, gender, etc.) can mitigate against compounding bias. Peer-reviewers may include internal and external care providers, researchers, educators, and diverse groups of data scientists other than AI algorithm developers.

Case Study 6: SingHealth

To support adoption of NoteBuddy across SingHealth institutions, user champions, who are from the data analytics team and proficient users, are appointed to assist new users during onboarding, provide support for basic queries and troubleshooting, and act as a liaison between implementation team and clinical users.

SingHealth monitors implementation progress through:

- a. Utilisation and adoption metrics on active users, frequency of logins and token consumption; and
- b. User feedback from satisfaction surveys, onboarding, and training.

6.4.2. While AI solutions offer benefits for both patients and staff, organisations should ensure that healthcare professionals maintain professional competencies and continue to upskill, particularly during contingencies (see **Chapter 7.2.4**). To measure deployment progress and success, organisations may consider monitoring: (i) AI adoption rate; (ii) AI utilisation rate (e.g. session duration, frequency, features, etc); and (iii) user satisfaction.

6.5. Staff Training

- 6.5.1. Organisations should develop comprehensive training programmes to ensure all relevant staff understand how to use AI solutions responsibly and effectively. They may wish to partner with developers to curate the training, and developers can present on the AI solutions' performance specifications and limitations as well as the intended use for the AI solutions.
- 6.5.2. The staff training should comprehensively cover roles and responsibilities, including but not limited to the relevant professional standards and regulatory requirements, workflow integration, contingency plans, specific use cases, AI solution limitations, output interpretation, patient communication protocols, data input restrictions (including data security / classification), abnormality detection and incident reporting. Organisations should also re-train staff when there are significant updates to the AI solutions or clinical workflows. To train staff in AI solutions, organisations can provide training through multiple channels, including:
 - a. Self-paced online courses and learning materials that staff can access anytime;
 - b. Hands-on demonstrations and knowledge sharing sessions led by experienced colleagues; and
 - c. Structured seminars and comprehensive training programmes.

6.6. Periodic Monitoring

- 6.6.1. Organisations should establish a process to monitor the AI solution's performance, both on an ad-hoc and regular basis, to ensure the continued safety, efficacy, and robustness of the model during deployment.

6.6.2. The process for periodic evaluation and validation of AI solution's performance should ensure it minimally meets the clinical practice baseline (i.e. AI solution's performance is equivalent to or better than current practice for patient safety or welfare), and verify the accuracy, and reproducibility of the AI solution's algorithmic decisions. Additionally, user feedback channels should be established and communicated. The feedback should be conveyed in a timely manner to the (i) organisational leadership, (ii) relevant authorities, and (iii) developers or technical teams. Some features of the feedback loop should include (i) roles & responsibilities, (ii) reporting timelines, (iii) performance measurements or indicators such as model drift, (iv) response matrix, etc. Feedback on user satisfaction should also form part of the holistic feedback loop (see **Chapter 7.3.1**) to optimise AI deployment.



6.7. Response to Adverse Events

- 6.7.1. Organisations should establish a process to detect, respond, investigate, and report adverse events or other device issues resulting from the use of the AI solution.
- a. Detect and respond: Organisations should train staff to detect adverse events and respond in a timely and appropriate manner. Contingency plans should include shutting down the AI solution and switching to contingency protocols (e.g. where the AI solution is not involved) to ensure the impact to the patient is minimised, while working with developers to implement both interim and permanent measures.
 - b. Report adverse event: Organisations should establish internal adverse event and incident reporting mechanisms, and abide by existing legislative frameworks like:
 - i. Data-related incidents by Personal Data Protection Commission (PDPC);²⁴
 - ii. Adverse events from the use of medical devices by HSA;²⁵ and
 - iii. Critical cybersecurity incidents by Cyber Security Agency of Singapore (CSA).²⁶

²⁴ Organisations are legally required to report data breach incidents that are likely to cause significant harm to the affected individuals or affect a significant scale of individuals. Refer to PDPC's guidance on reporting a data breach.

²⁵ Refer to HSA's guidance on Adverse Event reporting.

²⁶ Prescribed cybersecurity incidents which affect Critical Information Infrastructure (CII), or inter-connected systems that are connected to or communicate with CII, should be reported in accordance with the Cybersecurity Act 2018 and Cybersecurity (Critical Information Infrastructure) Regulations 2018. Other AI-related cybersecurity incidents may be reported to the Singapore Cyber Emergency Response Team (SingCERT).

- d. Investigate and understand: As the root cause of an adverse event may not be immediately linked to the AI solution, organisations should properly document and investigate adverse events. Where appropriate, organisations may also involve the AI developer in investigating the issue.

6.8. End of Life

- 6.8.1. Organisations should work with developers to establish an exit management plan before decommissioning an AI solution or changing the provider before integrating the AI solution into a workflow.

6.9. Patient Communications

- 6.9.1. Patient autonomy is a fundamental principle in medical ethics and must be respected. Patients should be provided accurate and sufficient information for them to be able to make informed decisions about their medical management.



- a. An important part of patient autonomy involves ensuring that patients give their valid consent (if they are able to do so) to any treatment or procedure prior to their undergoing such treatment or procedure.²⁷



- b. This involves patients understanding: (i) the purpose of tests, treatments or procedures to be performed on them; (ii) potential implications; as well as (iii) alternatives available to them, before making voluntary decisions on their medical care.

- 6.9.2. Not all instances of AI use in healthcare settings are relevant to a patient's medical management. For instance, the use of AI in automating billing or scheduling healthcare professionals' time-off are not relevant.

- 6.9.3. However, there are instances whereby AI use is relevant to a patient's medical management because AI use has significant impact and require informed decision-making (see **Chapter 6.9.1(b)**). Clear and open communication also helps patients understand and feel comfortable with technological changes in their healthcare services.

- 6.9.4. In such instances, organisations are encouraged to establish a process



²⁷ Section C5 and C6 of the Singapore Medical Council's *Ethical Code and Ethical Guidelines* for details.

for patient communications to support patient's decision-making about medical management. The information provided should focus on the benefits (particularly when compared to non-deployment of AI such as missed early detection of warning signs), significant limitations, material risks (particularly those relevant to patients' individual circumstances), and possible complications as well as alternatives available to them of the tests, treatments or procedures, rather than the technical details of the AI solution.

- 6.9.5. Organisations should establish a streamlined process to facilitate effective patient communications. This would include training healthcare professionals to communicate AI use to patients and, where appropriate, making resources about AI use available to patients.
 - a. Organisations may wish to calibrate the extent of information provided based on the benefits to be accrued and/or risks posed by AI use in specific instances. This should be deliberated and decided by the organisation's internal governance platform (see **Chapter 6.1**).
 - b. Multiple AI solutions may be used in a single workflow. There is no expectation for the use of each individual solution to be communicated to patients as this would not be meaningful. Instead, organisations should review the clinical workflow holistically and work with healthcare professionals to decide on how to operationalise this based on the specific context.
 - c. There are various modalities to patient communications – e.g. delivered directly via written communications or information pamphlets, publication on websites. Organisations should work with healthcare professionals to decide on the most appropriate modality.
- 6.9.6. Organisations that deploy direct-to-consumer AI applications for patients' independent use outside of usual care settings must ensure that patients are clearly informed that they are interacting with AI applications. Some key information for patients includes clear instructions on responsible use, limitations, and escalation protocols (details in **Chapter 5.7.2**).
- 6.9.7. For the avoidance of doubt, organisations and healthcare professionals must continue to comply with all applicable laws as well as ethical and professional codes of conduct.

7 Users: Healthcare Professionals

This section provides guidance on the responsibilities of users of AI – i.e. healthcare professionals. This includes the required competencies in AI use and managing adverse events as well as recommendations for patient communications.



7.1. Deployment

7.1.1. Healthcare professionals remain responsible for delivering quality and safe patient care, regardless of whether AI solutions are used. When using AI solutions, healthcare professionals must use them appropriately and exercise professional judgment in (i) verifying that the data entered into an AI solution is accurate and appropriate; and (ii) carefully evaluating AI output data before incorporating them into their workflows. Examples of healthcare professionals' responsibilities regarding input and output data are in **Table 5** below:

Table 5: Examples of Responsible Use regarding Input and Output Data

	 Input data Information and material provided/uploaded by users, e.g.: prompts, code, files, images, audio, videos, etc.	 Output data Information and material generated by applications in response to input, e.g: text, code, files, images, audio, videos, etc.
Users are accountable for:	1. Accurate data classification of all input data. 2. Compliance with existing data privacy, security standards, and handling requirements in accordance with the classification of such data. 3. Appropriate use of input data, in compliance with legal requirements and obligations.	1. Accuracy and quality of all work produced with the assistance of AI. 2. Ensuring the responsible and appropriate use of AI-generated output, accounting for the specific context of use.

Users must:

- a. Review all input data (including non-text input, e.g. files and images) to ensure it is within the organisation's permitted security and sensitivity classification for the AI solution in question. Anonymise personal data and other sensitive information whenever necessary or operationally possible.
 - b. Include instructions in prompts for AI-generated output to include source citations or references, if possible, to facilitate fact- and source-checking.
 - c. Avoid using prompts which carry a high risk of infringing intellectual property (IP) rights (e.g. avoid prompts such as "produce an image similar to [X's] works", or "create an image incorporating [X's] logo").
- a. Assess the suitability of AI-generated output (e.g. AI-generated text, images, audio, videos, etc) for the use case in question, considering factors such as whether:
 - i. The output may be inappropriate for the task or context at hand
 - ii. The output may contain falsehoods, inaccuracies, distortions, or reflect biased or discriminatory views
 - iii. The output may contain confidential or sensitive information
 - iv. The output may infringe IP rights, or give rise to liability or reputational harm to the organisation
 - b. Fact-check, source-check, proof-read, verify, and adapt all output appropriately given the context and purpose for which the output is intended to be used.
 - c. Ensure that the use of such output complies with relevant legislation and policies (e.g. Personal Data Protection Act 2012 (PDPA)).

7.2. Staff Training


- 7.2.1. It is crucial that healthcare professionals maintain adequate understanding of the AI solution's intended use, limitations, clinical workflow integration (e.g. whether personal data may be input into the AI solution, how the output may be factored into clinical decisions), including but not limited to relevant professional standards and regulatory requirements, ethical codes and ethical guidelines, patient communication protocols, abnormality detection, and incident reporting before deployment. Similar to getting trained on tools like robotic arms to perform surgeries, healthcare professionals should continually update their understanding of the AI solutions through self-study or formal trainings provided by their employer and/or AI developer. This ensures that they can confidently integrate AI into their practice and maintain high standards of care.
- 7.2.2. Healthcare professionals should minimally be able to recognise limitations such as when the AI solution is not performing as expected. This is no different from the expectations on clinicians when using any other tool or device. For AI, there are various abnormalities that healthcare professionals ought to recognise, including hallucinations.
- 7.2.3. When abnormalities are detected, healthcare professionals must adapt and respond to these issues. They remain accountable for patient care and patients' outcomes while AI is the supportive tool to enable care delivery.
- 7.2.4. Healthcare professionals must maintain proficiency in contingency and containment measures for AI solutions, especially where AI is relied on as a standard practice. This includes understanding fail-safe mechanisms and maintaining skills in non-AI-augmented workflows to ensure continuity of care should any abnormalities be detected in the AI solution.

7.3. Periodic Monitoring

- 7.3.1. Healthcare professionals should periodically provide feedback to developers and organisations on AI system performance and clinical outcomes, risks and vulnerabilities relating to security and data protection. This proactive approach facilitates early warning and provides developers and organisations with valuable insights into potential issues or areas for improvement.

7.4. Response to Adverse Events

- 7.4.1. Healthcare professionals should also abide by their organisations' internal reporting processes for adverse events and report them timely. It is recommended that any abnormalities or unexpected behaviours observed in AI systems are recorded, even if these are minor and do not cause adverse outcomes. The root cause of the adverse event may not be immediately apparent, and further investigation may be needed, depending on the incident.




7.5. Patient Communications

- 7.5.1. Healthcare professionals have a duty to provide accurate and adequate information to their patients to allow them to make informed choices about their medical management and give their valid consent (if they are able to do so) to any treatment or procedure prior to their undergoing such treatment or procedure.²⁸ This is aligned with prevailing practice stipulated in clinical practice guidelines.



- 7.5.2. When using AI solutions to support care delivery, healthcare professionals should take reference from the organisational processes (see **Chapter 6.9.5**) but also exercise discretion to share more information, if appropriate. The extent of information a patient requires to make an informed choice on their medical management would differ, depending on their priorities, preferences, expectations, and concerns.



- 7.5.3. If healthcare professionals are involved in deploying direct-to-consumer AI applications for patients' independent use, they must also ensure that patients are clearly informed that they are interacting with AI applications and take reference from the organisational processes in **Chapter 6.9.6**.



²⁸ This is aligned with the Singapore Medical Council Ethical Code and Ethical Guidelines on clinicians' ethical duty to respect patient autonomy. Other healthcare professionals are also guided on respect for patient autonomy in their professional codes of conduct (refer to Chapter 9).



Emerging Developments

8.1. AI Solutions with Continuous Learning Capabilities

8.1.1. AI solutions with continuous learning²⁹ capabilities can adapt, and update based on new data ingested without reprogramming. This means they can change their behaviour or performance post-deployment which can be beneficial in maintaining relevance to changing patterns. However, these models pose specific risks, such as model drift where performance degrades over time as the underlying data distribution changes and unintended data disclosure.

8.1.2. This requires stakeholders to adopt appropriate risk mitigation strategies, such as:

- a. Developers should implement robust monitoring systems that track model performance. This is to ensure that the AI solution's performance is still within the acceptable limit and does not deteriorate over time.
- b. Deployers should design workflows and implement periodic assessment to support identification of any substantial changes to the intended capabilities and function. They should ensure healthcare professionals are trained to understand additional risks of AI solutions with continuous learning capabilities – e.g. to be able to identify signs of model drift and to be able to respond appropriately.
- c. Users should be vigilant of any substantial changes and report or escalate issues as stipulated in the organisational workflows.




²⁹ IMDRF/AIMD WG/N67 (Edition 1): 2022 *Machine Learning-enabled Medical Devices: Key Terms and Definitions*.

8.2. Generative AI

8.2.1. GenAI has the capability to create various forms of content (e.g. text, images, videos) from user prompts, offering significant potential in healthcare to synthesise new output, support data-driven diagnoses and transform patient care with virtual health assistants. Yet it also amplifies existing risks, like:



- a. Hallucination: GenAI models may generate and present incomplete or fictitious content as facts and users may not notice such inaccuracies.



- b. Undesirable content: GenAI models may produce harmful content that can potentially compromise clinical outcomes or patient safety (e.g. mental health chatbot providing dismissive responses to user distress).



- c. Data disclosure: GenAI models could inadvertently leak sensitive data when users submit sensitive data during use or when models are prompted in specific ways.



- d. Vulnerability to adversarial prompts: GenAI models are vulnerable to adversarial attacks, which are deliberate attempts to manipulate the model's output or behaviour, leading to misinformation and security breaches.

8.2.2. This requires stakeholders' careful management of the amplified risks, and compliance with prevailing guidelines and policies, where applicable.³⁰ Some examples include:

- a. Developers should validate GenAI models through methods that assess the quality and diversity of the generated content, such as human evaluation studies or specific metrics. Automated metrics, such as Recall-Oriented Understudy for Gisting Evaluation (ROUGE), Bilingual Evaluation Understudy (BLEU), and Bidirectional Encoder Representations from Transformers (BERT)-score can be

³⁰ MOH Holdings entities should refer to relevant guidelines applicable to the Public Healthcare Sector.

used but they exhibit significant limitations when applied to the evaluation of healthcare-related content. They should use qualitative and quantitative evaluation methodologies specific to the use case of the GenAI³¹ to assess the model's performance, taking into account both objective metrics and subjective human judgment. Perhaps the most promising recent technique is to compute "certainty" of the outputs



of LLMs. This works by examining the falloff or probabilities in the "logits" in the output layer of the LLM. In medical applications, uncertainty or errors are often categorised as either aleatoric or epistemic. Aleatoric uncertainty in a medical model follows from inherent randomness in a process, or errors in collecting information, whereas epistemic errors are shortcomings of the model, often due to improper or incomplete summarisation of the knowledge which it is trying to represent. While noisy data is present in healthcare, LLM's hallucinations consist of information that is "made up", is more a form of epistemic uncertainty, and is inherent to LLMs and hard to cure. Therefore, developers should subject GenAI models to red teaming to "break" the model and thereby induce safety, security, and other violations for testing purposes.

- b. Where possible, developers should design GenAI applications with the function to support fact-checking by citing sources, identifying unknown entities, or prompting users to review the output against sources. Developers should reduce hallucination using appropriate techniques such as Retrieval-Augmented-Generation (RAG).
- c. Deployers should employ relevant testing strategies to improve LLM safety and performance such as baseline and specific tests using public benchmarking and red teaming; and component testing to identify failure points for further mitigation.
- d. Deployers should design workflows to support fact-checking of outputs and ensure comprehensive training of healthcare professionals on the responsible use of GenAI, including redlines. They should also support with patient communication to address concerns on the use of GenAI. Wherever possible, they should provide clear and front-end labels in user interfaces to highlight where recommendations, information or content is generated by AI.
- e. Users should be well-trained on how to critically evaluate GenAI outputs. They should always maintain clinical judgment as the primary decision-making tool and use GenAI solutions as a supportive resource. Where possible, healthcare professionals should also include instructions for the model to cite sources for its responses to facilitate fact- and source-checking, and review output data for any sensitive or identifiable information and scrub such data where necessary.

³¹ Please contact HSA on the regulatory requirements for medical device incorporating GenAI. Developers may arrange a Pre-market Consultation for regulatory inputs (see Chapter 5.9).

Case Study 7: MOH Office for Healthcare Transformation (MOHT)

MOHT deploys GenAI to support mental health helplines operators to listen empathetically, assess risk, and provide timely support and de-escalation. It provides evidence-informed guidance for operators to decide whether to adopt AI suggestions, enhancing human judgment while safeguarding empathy, safety, and trust.

8.3. Direct-to-Consumer (DTC) AI Applications

- 8.3.1. DTC or patient-facing applications are increasingly deployed outside traditional care settings, democratising access to healthcare services through patients' personal devices. These applications enable patients to access health information and self-help tools independently, without direct healthcare professional supervision. All DTC applications that fall under the definition of an AI-MD or SaMD must be registered with HSA. This regulatory oversight remains fundamental regardless of the consumer-facing nature of the application.
- 8.3.2. DTC applications have additional safety considerations as they operate outside conventional healthcare environments, with critical safety layers (e.g. medical device oversight, healthcare settings, professional guidance) reduced or removed. Further, individuals may have limited medical knowledge, and additional safeguards may be essential to ensure safe and appropriate use.
- 8.3.3. Stakeholders must recognise how DTC applications will be used and implement appropriate guardrails, including comprehensive consumer education. Best practices include:
 - a. Developers must comply with all applicable regulations and guidelines, whether developing DTC MD or non-MD. For non-MD DTC applications, developers should carefully consider the complete patient journey: how individuals will interact with the application, interpret outputs and act upon the information provided to them. Application interfaces must be designed with simplicity and clarity, presenting information in language easily understood by laypersons. Critically, DTC applications should not generate outputs requiring clinical expertise to interpret safely, even when accompanied by disclaimers about not providing medical advice.
 - b. Deployers that deploy DTC applications for independent consumer use outside of healthcare settings should provide clear communication about intended use, limitations, and escalation protocols.
 - c. Healthcare professionals that recommend the use of a DTC application to a patient, must educate the patient on proper use and advise them to seek professional advice when unclear, in accordance with organisational processes. This educational role is crucial for ensuring safe and effective engagement with the technology.

Other Relevant Legislation and Guidelines

- 9.1. In addition to complying with HSA's HPA regulations, organisations or individuals involved in the development or deployment of AI-MD should also consider the requirements set out in other related legislation and guidelines covering the provision of healthcare services, professional responsibilities, product safety, cybersecurity, and data protection, amongst others. Examples of the relevant legislation and guidelines (as at the date of publication of the Guidelines) are indicated in **Table 6**.

Table 6: Examples of Key Legislation and Guidelines³²

Legislation or Guidelines	Regulatory Context
Legislation	
1. Personal Data Protection Act 2012 (PDPA)	Obligations in the collection, use, disclosure, protection, and access of personal data used in AI solutions.
2. Healthcare Services Act 2020 (HCSA) and Guidelines	Healthcare Institutions licensed under the HCSA, including those which use AI solutions, must comply with its requirements (e.g. controls on the maintenance and security of medical records) and all licensing terms and conditions.
3. Legislation governing professions (e.g. Medical Registration Act 1997)	All registered healthcare professionals, including those who use AI solutions to deliver healthcare services (e.g. clinical decision support tools to read Computer Tomography (CT) scans), must comply with the requirements under their respective professional Acts.
4. Civil Law (Amendment) Act 2020 (see section 37)	Sets out the statutory test for the standard of care for providing medical advice.
5. Health Products Act 2007 (HPA)	The HPA is Singapore's primary legislation governing the manufacture, import, supply, presentation, and advertisement of health products including medical devices, to ensure their quality, safety, and efficacy.

³² References to an Act include subsidiary legislation made under the authority of that Act.

Legislation or Guidelines	Regulatory Context
6. Health Products (Medical Devices) Regulation 2010	Regulatory controls over medical devices (e.g. dealer's licensing, product registration, change management for registered medical devices, notification of medical devices used in clinical trials), and post-market surveillance.
7. Human Biomedical Research Act 2015 (HBRA)	Researchers conducting healthcare research involving AI, which fall under the definition of human biomedical research, must comply with the requirements under the HBRA and/or refer to the ethical guidance document for researchers and those involved in human biomedical research in Singapore.
Guidelines	
8. GL-04: Regulatory Guidelines for software medical devices – a life cycle approach	Provides clarity on the regulatory requirements for software medical devices in its entire life cycle including machine learning enabled medical device.
9. Guidelines on Risk Classification of Standalone Medical Mobile Applications and Qualification of Clinical Decision Support Software (CDSS)	This guideline takes reference from the IMDRF's Framework for Software as a Medical Device (SaMD) to determine the risk classification of Standalone Medical Mobile Applications that are Medical Devices (commonly referred as SaMD). This guideline also provides clarity on the qualification of Clinical Decision Support Software (CDSS) as regulated medical devices or otherwise, as well as the current regulatory approach and requirements for such software that are regulated by HSA.
10. GN-13: Guidance on the Risk Classification of General Medical Devices	Provides information on classification of general medical devices using risk based classification rules.
11. GN-14: Guidance on the Risk Classification of In vitro Diagnostic Medical Devices	Provides information on classification of in vitro diagnostic (IVD) medical devices using risk based classification rules.
12. GN-15: Guidance on Medical Device Product Registration	Provides information on general requirements for product registration for medical devices.

Legislation or Guidelines	Regulatory Context
13. GN-17: Guidance on Preparation of a Product Registration Submission for General Medical Devices using the ASEAN Common Submission Dossier Template (CSDT), OR GN-18: Guidance on Preparation of a Product Registration Submission for In Vitro Diagnostic (IVD) Medical Devices using the ASEAN CSDT	Provides guidance on the preparation of a product registration submission using the ASEAN CSDT.
14. Revised Guidelines for the Retention Periods of Medical Records (2022)	Guidelines on the retention periods of medical records apply to all licensees under HCSA.
15. National Telemedicine Guidelines (2015)	Specific guidelines for the provision of telemedicine services.
16. Cybersecurity and Data Security Essentials (2026) (by MOH, in consultation with CSA, IMDA, and PDPC)	The guidelines provide guidance to all healthcare providers on security measures for the proper storage, access, use, and sharing of health information.
17. Professional Ethical Codes and Guidelines, e.g. • Singapore Medical Council Ethical Code and Ethical Guidelines (2016) • Singapore Dental Council Ethical Code and Ethical Guidelines (2018) • Singapore Pharmacy Council Code of Ethics (2015) • Singapore Nursing Board Code for Nurses and Midwives (2018) • Allied Health Professionals Council Code of Professional Conduct (2013) • Singapore Association of Social Workers Code of Professional Ethics (2021)	The respective professional ethical guidelines to enable healthcare professionals to navigate the ethical issues encountered in practice.
18. Bioethics Advisory Committee (BAC) Report on Ethical Use of Big Data and Artificial Intelligence in Biomedical Research (2025)	The report covers ethical issues arising from the use of big data and AI in human biomedical research, such as responsible data usage, data ownership, custodianship and stewardship, data privacy, accessibility and security, data anonymisation, and other ethical considerations and issues specific to AI.

Legislation or Guidelines	Regulatory Context
19. BAC Ethics Guidelines for Human Biomedical Research (2021 rev. ed.)	The guidelines serve as a one-stop resource for researchers and members of ethics committees, or any interested individual seeking guidance on best practices for the ethical conduct of human biomedical research in Singapore.
20. IMDA Model AI Governance Framework (2 nd Edition)	The Model AI Governance provides a set of best practices that organisations can readily adopt to develop and deploy traditional AI solutions responsibly.
21. PDPC Implementation and Self-Assessment Guide for Organisations (ISAGO)	A companion to complement the voluntary Model AI Governance Framework and aims to help organisations assess the alignment of their AI governance processes with the Model Framework, identify potential gaps in their existing processes and address them accordingly.
22. PDPC Advisory Guidelines for the Healthcare Sector (2023)	The guidelines clarify how the Data Protection Provisions in the PDPA apply to healthcare institutions' collection, use, and disclosure of personal data, as well as suggest good data protection practices in certain scenarios.
23. PDPC Advisory Guidelines on use of Personal Data in AI Recommendation and Decision Systems (2024)	The guidelines clarify on use of personal data to develop AI, information to provide to consumers when seeking consent and best practices to support compliance with the PDPA.
24. PDPC Proposed Guide on Synthetic Data Generation (2024)	The proposed guide assist organisations to understand synthetic data (SD) generation techniques and potential use cases, particularly for AI.

10 Acknowledgements

The Ministry of Health would like to express our appreciation to the following partners who have co-developed AIHGle 2.0 with us:



We would also like to express our gratitude to the following organisations and individuals that have contributed to AIHGle 2.0.

Organisations

Alliance of Patients' Organisations Singapore	NHG Health	Singapore Pharmacy Council
Amazon Web Services	Pharmaceutical Society of Singapore	Singapore Physiotherapy Association
Asia Pacific Medical Technology Association	Raffles Hospital	Singapore Society of Radiographers
Crawford Hospital	Singapore Dental Association	SingHealth
Google	Singapore Medical Association	The Farrer Park Company
Mount Alvernia Hospital	Singapore Medical Council	Thomson Medical
National University Health System	Singapore Nurses Association	Traditional Chinese Medicine Practitioners Board

Individuals

Alastair Denniston

Professor of Regulatory
Science and Innovation,
University of Birmingham

Harvey Castro

Chief Medical AI Officer,
Helpp.ai

M Kamala Devi

Associate Professor,
Programme Director Nursing,
University of Glasgow

Andy Greenfield

Honorary Research Fellow,
Nuffield Department of
Women's & Reproductive
Health, Institute of
Reproductive Sciences,
University of Oxford

James Kingsland OBE

Chairperson,
Digital Clinical Excellence
Forum (DiCE)

Rifat Atun

Professor Global Health Systems,
Harvard T.H. Chan School of
Public Health

Charlene Li

Founder and Chief
Executive Officer,
Quantum Networks Group

Julian Savulescu

Chen Su Lan Centennial
Professor in Medical Ethics,
National University Singapore

Simon Chesterman

Senior Director of AI Governance,
AI Singapore

Christopher Hodges OBE

Emeritus Professor of
Justice Systems,
Centre for Socio-Legal Studies,
University of Oxford

Keith McNeil

Commissioner,
Commission on Excellence
and Innovation in Health,
South Australia

Tan Boon Gin

Chief Executive Officer,
Singapore Exchange Regulation

Emmanuel Eckard

Senior Expert on AI,
Institute for Regulatory
Innovation, Delivery and
Effectiveness,
European Public Law Organisation

Kuah Boon Theng SC

Managing Director,
Legal Clinic LLC

Yeong Zee Kin

Chief Executive,
Singapore Academy of Law

Florentin Blanc

Director,
Institute for Regulatory
Innovation, Delivery and
Effectiveness,
European Public Law Organisation

Leung Pak-yin

Honorary Clinical Professor,
School of Public Health of the
University of Hong Kong

Government Agencies

Ministry of Digital Development
and Information

Cyber Security Agency of
Singapore