Technology and Medicine in the AI and IoT Era

A lecture for medical students and early clinical trainees. Build a working mental model of what modern clinical AI—especially generative AI—can and cannot do, how AI is used today in radiology, pathology, and surgery, and how IoT and wearables are reshaping hospitals, home monitoring, and clinician workflow.

MEDICAL EDUCATION

AI & IOT

Lecture Goals & Learning Objectives

01

Generative AI Fundamentals

Explain what generative AI is, why it differs from traditional predictive models, and where it fits in clinical workflows.

02

AI in Clinical Specialties

Describe concrete examples of AI in radiology, pathology, and robotic surgery—including outputs, clinician roles, and failure modes.

03

CDSS & Governance

Define clinical decision support systems, recognize alert fatigue and automation bias, and outline governance basics.

04

IoT & Wearables

Describe how IoT and wearables change hospital operations, home monitoring, and decentralized clinical research.

05

AI Literacy

Articulate what AI literacy means for physicians and why early clinician involvement in AI development matters.

CHAPTER 1

The Digital Substrate of Modern Medicine

Modern medicine is increasingly **data-native**: diagnoses, images, labs, pathology slides, medication orders, and clinical notes are stored and exchanged digitally—creating the conditions for AI systems to plug into care.

Three Infrastructures of Digital Health

A useful framing: technology changes medicine through three core infrastructures that define where errors arise and how systems must be governed.

🖥️ Compute

Models that can transform and reason over data—the intelligence layer of clinical AI.

🌐 Connectivity

Networks and APIs that move data between systems and locations—the integration layer.

📡 Sensing

Devices that generate continuous or episodic measurements—the data acquisition layer.

ⓘ All three are emphasized in contemporary digital health guidance because they define where errors can arise and how systems must be governed.

Interoperability Standards: The "Plumbing" That Makes AI Work

FHIR

Fast Healthcare Interoperability Resources — the widely used standard for exchanging healthcare information electronically and enabling app-based access to EHR data.

DICOM

The international standard for medical images and related information, allowing imaging devices and systems to exchange clinically usable images.

OMOP CDM

Standardizes the structure of observational data to support reproducible analyses and multi-institution evidence generation.

Why "Plumbing" Matters

These standards often determine whether AI is a prototype or a durable clinical tool. Integration failures—missing context, wrong patient, wrong timestamp, wrong units—are **safety failures**, even if the model is accurate on a benchmark.

⚠ Interoperability is not a technical nicety. It is a patient safety requirement.

CHAPTER 2

AI & Generative AI Fundamentals for Clinicians

Generative AI systems—including large language models and multimodal foundation models—are trained to **generate** outputs (text, images, structured fields) rather than only classify or predict a predetermined label. This expands AI from "detect a pulmonary embolus: yes/no" into drafting reports, extracting structured fields, and summarizing longitudinal records.

Where Generative AI Fits in Clinical Work

Clinical Communication & Documentation

Drafting encounter notes, discharge instructions, and patient-friendly explanations (with clinician verification). Translating medical content into plain language—with caution, as subtle errors can change meaning.

Information Extraction & Structuring

Turning unstructured text into structured fields (problem lists, staging variables, quality measures)—powerful, but sensitive to hallucination and misinterpretation; requires evaluation and audit trails.

Clinical & Operational "Co-pilots"

Drafting ordersets, checklists, or referral letters; assisting with coding/billing; workflow routing. These often become **decision support** once they influence care pathways, making governance relevant.

Data Generation

Generating synthetic data to augment datasets or enable safer sharing—useful, but not automatically private. Re-identification and leakage risks must be tested rather than assumed away.

Why Generative AI Can Be Unsafe Without Guardrails

Key Risks

- **Hallucinations:** Confident but incorrect output
- **Bias propagation:** Reflecting dataset inequities
- **Over-trust:** Weakening critical thinking
- **Security:** Medical LLMs can be manipulated to inject incorrect biomedical "facts"

Generative AI is best treated as a **high-speed junior assistant**—excellent at drafting and organizing, unreliable as a sole source of truth, and always requiring accountable sign-off when used in patient care.

This stance is consistent with WHO guidance emphasizing human oversight and risk management for high-impact AI.

CHAPTER 3

AI in Radiology, Pathology & Robotic Surgery

AI in Radiology: A Natural Home

Radiology is a natural home for AI because it is already digital, standardized (DICOM), and workflow-driven. The RSNA has emphasized safe and effective implementation of AI tools, including guidance on deployment and oversight.

Triage & Worklist Prioritization

Flagging suspected critical findings to ensure urgent cases are read first.

Detection & Segmentation

Assisting with lesion, hemorrhage, and nodule identification.

Quality Control

Identifying motion artifacts and protocol mismatches before interpretation.

Measurement Automation

Tumor burden quantification and volumetric measurements.

Generative AI in Radiology: New Layers

Report Drafting & Impression Generation

LLM/VLM approaches to generating or improving radiology report components—with persistent concerns about fidelity and clinical reliability.

Report Simplification for Patients

A systematic review/meta-analysis evaluated LLMs for simplifying radiology reports—generative AI at the "communication boundary" between clinician and patient.

Image Reconstruction & Modality Translation

Generative models used for denoising, reconstruction, synthesis, and cross-modality translation—improving efficiency but demanding strong validation because image artifacts can directly change interpretation.

⚠ **Key teaching point:** In imaging, "generative" can mean generating clinically plausible images. Plausibility is *not* the same as truth.

AI in Pathology: From Glass Slides to Computational Analysis

Pathology is undergoing a foundational shift: **glass slides → whole slide imaging (WSI) → computational pathology.** The FDA authorized the first digital pathology WSI system for primary diagnosis in surgical pathology in 2017—a regulatory milestone. WSI systems continue to mature, including new clearances such as the **Philips IntelliSite Pathology Solution 5.1**.

Where AI Shows Up in Pathology Today

Cancer Detection

Assistive "second read" on biopsies to flag suspicious foci.

Metastasis Detection

AI can reduce reading time and improve detection performance in certain screening tasks.

Quantification & Grading

Tumor burden proxies and grading support tasks.

QC & Triage

Workflow quality control and slide-level triage to prioritize cases.

Regulated Pathology AI: Examples Students Should Know

Paige Prostate (Paige.AI)

Received **De Novo FDA authorization** as an assistive tool to help pathologists detect foci suspicious for cancer on prostate biopsy WSIs. Provides coordinates highlighting the most suspicious location. Explicitly framed as **adjunctive—not a primary diagnosis**.

Galen™ Second Read™ (Ibex Medical Analytics)

FDA-cleared digital pathology algorithm device (510(k) K241232)—illustrating the growing category of regulated "assistive" computational pathology tools.

Emerging Generative AI Uses

- Extracting structured variables from pathology reports for registries and research
- Creation of synthetic pathology images (still high-risk and evidence-sensitive)

Robotic Surgery & Perioperative AI

Robotic Assistance

Teleoperation, enhanced dexterity and vision—the current mainstream in many specialties.

AI-Enabled Autonomy

Systems that interpret context and act with conditional or higher autonomy—still limited. A systematic review found most FDA-cleared surgical robots remain at basic assistance levels, with few reaching conditional autonomy.

ⓘ A concrete regulatory example: Intuitive Surgical's **da Vinci 5** platform received FDA 510(k) clearance (K232610).

How AI Is Used Around the OR

Preoperative Planning & Risk Prediction

Models built on EHR + imaging + labs stratify risk—but must be monitored because drift and practice changes can degrade performance.

Intraoperative Scene Understanding

Computer vision models interpret laparoscopic/robotic video (steps, instruments, anatomy) to support workflow analytics, coaching, and safety checks.

Skill Assessment & Training

AI analyzes video/kinematics to support objective feedback and scalable surgical assessment.

Toward Partial Autonomy

Research describes AI-supported automation of sub-tasks (e.g., suturing) and foundation models for surgical video—promising, but still early and evidence-intensive.

The "AI" in surgery often appears **before and after** the incision—in planning, safety monitoring, documentation, and outcome prediction—at least as much as inside the robot.

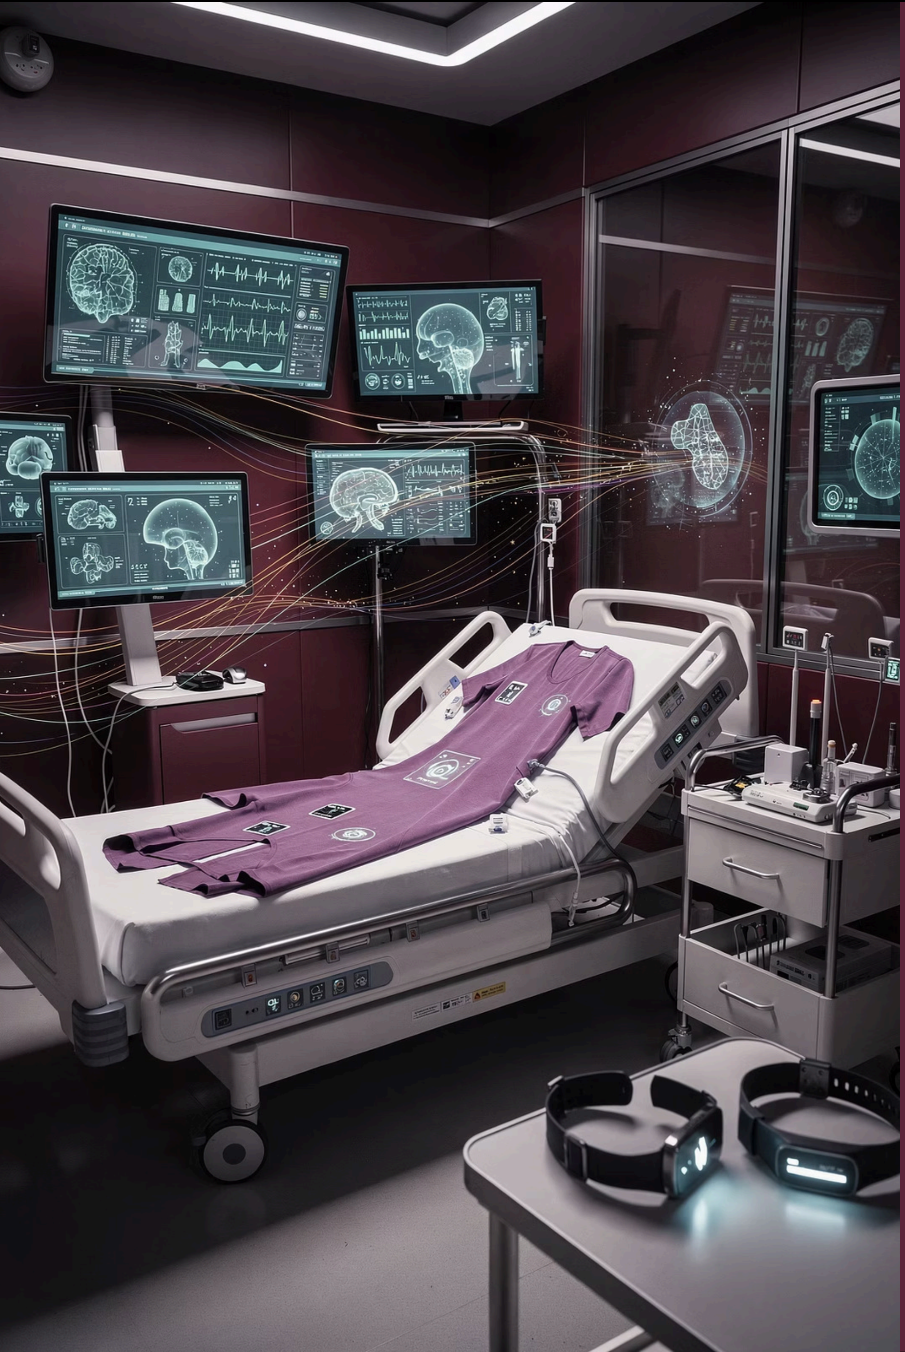

CHAPTER 4

IoT and Wearables: Smart Hospitals, Home Monitoring & Clinician Support

What IoT (and IoMT) Means in a Hospital Context

The **Internet of Medical Things (IoMT)** refers to networks of connected sensors and devices—patient-worn, room-based, and equipment-based—that capture data and can trigger actions or decision support. This is central to "smart hospital" visions because operational and clinical data become continuous and machine-readable.

Real-Time Location Systems (RTLS)

Tracking patients, staff, and equipment (BLE/RFID) to reduce delays, improve throughput, and enhance safety visibility.

Asset Tracking & Inventory

Locating infusion pumps, portable ultrasound, ventilators—reducing "hunt time" and enabling preventive maintenance scheduling.

Environmental Sensing

Room occupancy, airflow proxies, hand-hygiene workflow augmentation (evidence quality is mixed and context-dependent).

⚠ IoT is not "just convenience." Because it touches devices affecting patient care, **cybersecurity becomes clinical safety**. The FDA emphasizes cybersecurity across the full medical device lifecycle.

Wearables for Patients: Care Beyond the Clinic

Two Major Shifts

Clinical care shift: Continuous or frequent measurement outside the hospital.

Research shift: Remote endpoint capture in trials and observational studies.

FDA guidance explicitly addresses using digital health technologies to acquire data remotely from clinical investigation participants—reflecting how mainstream remote measurement has become.

Clinical Examples

- Heart rate/rhythm monitoring, activity/sleep proxies, fall detection
- Continuous glucose monitoring and connected medication adherence devices
- Home blood pressure monitoring integrated into chronic disease management

Sensor-based digital biomarkers are moving toward regulator-accepted digital endpoints in both the U.S. and Europe.

Wearables for Physicians: Decision Support at the Point of Action

Hands-Free Information Access

Smart glasses for telemedicine support, hands-free viewing, and situational awareness in time-critical settings.

Procedure Guidance & Training

Wearables used in health professional education and skills assessment—promising, but with privacy and workflow challenges.

Contextual Decision Support

Devices surfacing checklists, dosing guidance, or risk summaries in a workflow-aware way—useful, but susceptible to automation bias if poorly designed.

Clinician wearables can become "CDSS delivery channels"—their safety is inseparable from CDSS design principles: relevance, timing, interpretability, and minimizing cognitive load.

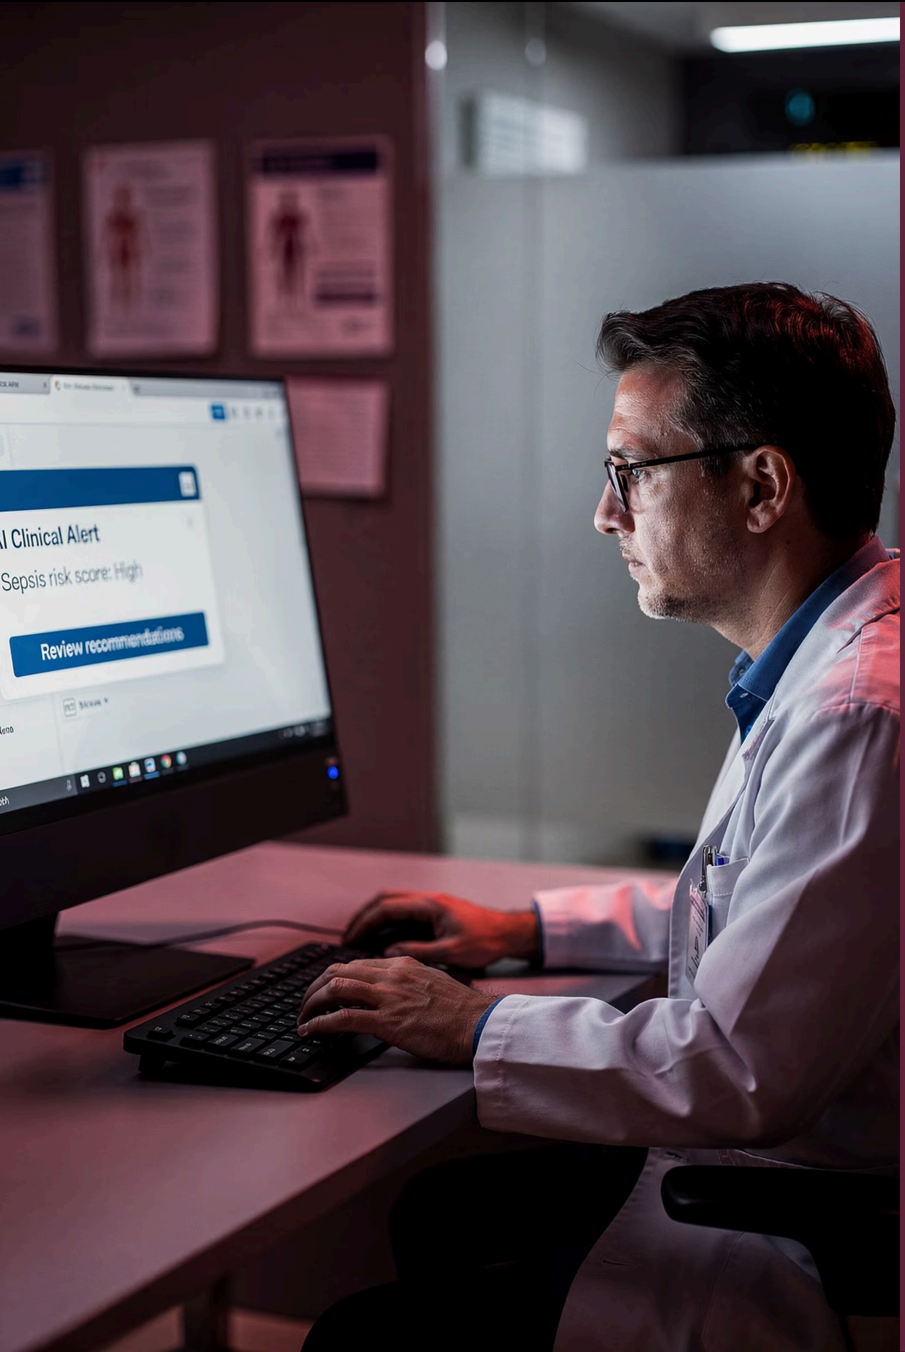

CHAPTER 5

Clinical Decision Support Systems & Human–AI Teamwork

What Is a CDSS?

A **clinical decision support system (CDSS)** is a computerized tool—often embedded in or integrated with the EHR—that provides information, recommendations, alerts, or predictions to support clinical decision-making.

Knowledge-Based CDS

Medication interaction alerts, duplicate therapy flags, guideline reminders, order sets.

Data-Driven / AI-Based CDS

Risk prediction (deterioration, sepsis, readmission), imaging triage outputs, pathology second-read highlights, OR complication prediction.

Generative AI-Enabled CDS

Summarizing longitudinal charts, generating differential diagnosis drafts, drafting consult recommendations—powerful but especially exposed to hallucination risks; requires verification and accountability.

The FDA's guidance clarifies which CDS software functions are regulated as medical devices—an important boundary when CDS meaningfully drives diagnosis or treatment decisions.

Core Failure Modes: Alert Fatigue & Automation Bias

⚠️ Alert Fatigue

When clinicians are inundated with warnings, they become desensitized and override or ignore many alerts—including potentially important ones. This is a known risk in CDS-heavy environments and a major patient safety concern.

🤖 Automation Bias

Clinicians may over-rely on automated outputs, especially when a system is highly (but imperfectly) reliable—creating new error modes such as failing to notice contradictory evidence. Systematic reviews describe this as a persistent risk for decision support tools.

> **CDSS shifts error patterns.** It can reduce some errors (missed interactions) while introducing others (over-trust, missed contradictions, degraded situational awareness).

Keeping AI and CDS Safe: Lifecycle Governance

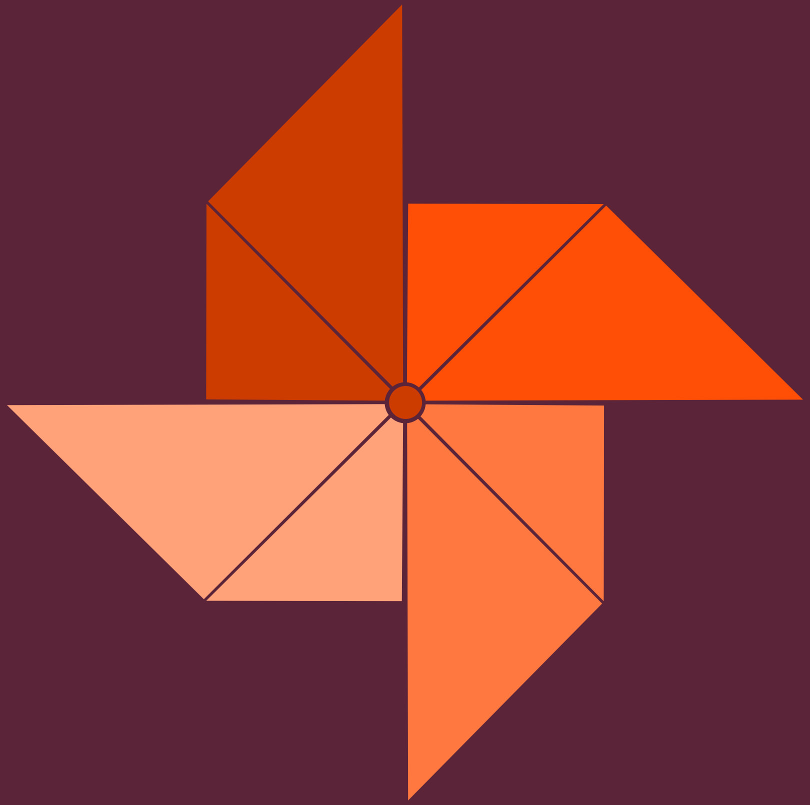

Many clinical AI tools degrade over time due to **model drift**—changes in data distributions, clinical practice, or patient populations. The FDA has issued guidance on **Predetermined Change Control Plans (PCCPs)** for AI-enabled devices, supporting iterative improvement while maintaining safety. Good Machine Learning Practice principles emphasize total product lifecycle thinking and human–AI team performance.

Hospital-Level AI Safety: Key Takeaways

1

Demand Clarity

Require clear documentation of intended use and known failure modes before deploying any AI tool.

2

Implement Monitoring

Track both performance and drift continuously post-deployment—not just at launch.

3

Treat Cybersecurity as Patient Safety

Connected medical devices and AI systems are attack surfaces; security failures are clinical failures.

4

Design for Cognition

Reduce cognitive load and alert overload through thoughtful interface design and alert governance.

CHAPTER 6

AI Literacy & the Future of Research and Clinical Trials

What "AI Literacy" Means for Physicians

AI literacy is **not** "learning to code first." It is the ability to understand what an AI tool is intended to do, evaluate evidence quality and generalizability, recognize bias, safety, and privacy risks, and communicate clearly with patients and teams about AI-assisted care.

AAMC

Developing national AI competencies across the medical education continuum to standardize foundational skills and ethical expectations.

AMA

Adopted policy advancing AI literacy in medical education and expanding AI-focused CME resources for practicing physicians.

World Medical Association

Explicitly stated that physicians must maintain appropriate AI literacy and that it should be systematically integrated into undergraduate medical curricula.

Why Early Clinician Involvement Matters

AI systems reflect the assumptions of their developers and the datasets they are built on. Early clinician involvement improves:

Outcome Definition

"What matters clinically?" must be answered by clinicians, not engineers alone.

Labeling Quality

Accurate ground truth requires clinical expertise at the annotation stage.

Workflow Fit

Tools designed without clinician input often fail at the point of care.

Monitoring Plans

Clinicians help define what "degraded performance" looks like in practice.

Decentralized Clinical Trials & Remote Endpoints

FDA Guidance on Decentralized Trials

Outlines how trial activities can occur remotely—via telehealth, in-home visits, and local providers—reflecting a durable shift in trial operations.

Digital Health Technologies for Remote Data

FDA guidance supports using connected devices and software to collect study data outside the clinic, potentially improving efficiency and participation.

European Guidance

EU recommendations formalize decentralized elements while emphasizing patient rights/safety, data reliability, and cross-border considerations. Early involvement of participants and investigators is explicitly highlighted as key to trust, recruitment, adherence, and scientific value.

Synthetic Data, Federated Learning & Privacy

Synthetic Data

Increasingly explored to bridge data access barriers for AI development and research. However, privacy is **not guaranteed**—peer-reviewed work describes re-identification risks and attack methods. Synthetic data must be evaluated with explicit privacy testing and governance.

Federated Learning

Enables model training across institutions without centralizing raw patient data. Reviews highlight privacy-preserving frameworks and the continuing need for robust security and validation.

Standards-Driven Research Infrastructure

Networks such as OHDSI and OMOP CDM enable distributed analytics and reproducible observational research at scale—an increasingly important complement to trials.

Governance & Regulation Students Should Know

FDA AI-Enabled Device List

Identifies AI-enabled devices authorized for marketing in the U.S., supporting transparency for providers and patients.

EU AI Act (2024)

AI-based medical software classified as "high-risk" with requirements for risk management, high-quality data, user information, and human oversight.

IMDRF SaMD Framework

Defines Software as a Medical Device and provides approaches to SaMD clinical evaluation, supporting regulatory convergence across jurisdictions.

Practical AI Literacy Behaviors for Medical Students

01

Ask What the Tool Is For

Intended use, population, setting—and what it is *not* for.

02

Treat Outputs as Hypotheses

Not conclusions, unless validated for that purpose with appropriate oversight.

03

Request Evidence Quality

Trial reporting standards like CONSORT-AI/SPIRIT-AI exist because low-quality evaluation is common.

04

Expect Monitoring

Drift happens; post-deployment surveillance is part of patient safety.

05

Understand Human Factors

Alert fatigue and automation bias are predictable and preventable design failures.

The Clinician as Competent Evaluator & Steward

> In the AI era, "being a good doctor" increasingly includes being a **competent evaluator and steward of clinical technology**—not only a user.

Major medical organizations (AAMC, AMA, WMA) are formalizing this expectation. AI literacy is now a professional competency, not an optional technical skill.

The Core Message

Not AI replacing doctors.

The next era of medicine is **doctors practicing in AI-saturated systems**—where clinical excellence includes knowing how to evaluate, supervise, and safely integrate new tools.

Evaluate

Assess evidence quality, intended use, and limitations of every AI tool.

Supervise

Maintain accountable human oversight—especially for high-impact decisions.

Integrate Safely

Design workflows that reduce cognitive load, prevent bias, and support monitoring.

Key Takeaways

Interoperability is a safety issue

FHIR, DICOM, and OMOP CDM determine whether AI tools are prototypes or durable clinical instruments.

Generative AI requires guardrails

Hallucinations, bias, and over-trust are predictable risks—treat AI as a high-speed junior assistant, not a sole source of truth.

Regulated examples exist now

Paige Prostate, Galen Second Read, da Vinci 5—real tools with real FDA clearances and defined intended uses.

IoT cybersecurity = patient safety

Connected devices and AI systems are attack surfaces; governance must span the full product lifecycle.

AI literacy is a professional obligation

AAMC, AMA, and WMA are formalizing AI competencies. The competent physician of tomorrow evaluates and stewards AI—not just uses it.