30 January 2025

WMA Webinar #1 - Introduction to AI in Medicine - Key Points

Table of Contents

Introduction

- Overview of the role of AI in modern healthcare.
- Key concerns regarding AI adoption:
    - Appropriate use cases for AI in medicine.
    - Safety, ethical implications, and regulatory considerations.
    - The capabilities and limitations of generative AI and large language models.
    - The potential for AI to complement or replace physician roles.
- Introduction of the five-part AI in Medicine webinar series.

Foundations of AI in Healthcare

Artificial Intelligence (AI) is transforming healthcare by improving diagnostic accuracy, streamlining workflows, and enabling personalized treatment. As AI continues to integrate into healthcare systems, understanding its foundational principles is essential. This section explores the background, key terminology, historical advancements, and considerations related to safety and trust in AI-driven healthcare.

A. Background on Digital Health

- Adoption Trends:
    - 93% of physicians believe digital health tools are helpful and provide some advantage.
    - The average physician uses almost four digital tools, double the number from a few years ago.
    - The COVID-19 pandemic accelerated the adoption of digital health tools, particularly telemedicine.
- Impact on Healthcare Delivery:
    - AI tools are being used for remote patient monitoring and chronic disease management.
    - Personalized treatment plans are being enhanced through AI-driven data analysis.
    - AI is streamlining administrative tasks, such as prior authorization, scheduling, and supply chain management.
- Data-Driven Decision-Making:
    - AI is enabling predictive analytics to support early disease identification and risk assessment.
    - Real-time data access is improving clinical decision-making at the point of care.

B. Key Terminology & Definitions

- Digital Health: The use of technology to enhance healthcare services, efficiency, and patient outcomes.
- AI: The application of advanced algorithms and machine learning models to replicate human cognitive functions.
- Digital Health Components:
    - EHRs: Electronic health records that centralize and streamline patient information.
    - Telemedicine: Virtual delivery of medical care through video or audio consultations.
    - Telehealth: Broader digital health services, including remote monitoring and patient education.
    - mHealth: The integration of mobile devices and wearable technology into healthcare.
- AI-Specific Terms:
    - Augmented Intelligence: AI designed to support, rather than replace, human expertise in healthcare.
    - Machine Learning: AI systems that continuously learn from vast datasets to improve decision-making.
    - Deep Learning: Advanced neural networks capable of complex pattern recognition in medical imaging and diagnostics.
    - Natural Language Processing (NLP): AI-driven interpretation of human language in medical records and documentation.
    - Foundational Models: AI models trained on large datasets, adaptable for healthcare applications such as diagnostics and treatment planning.

C. Historical Evolution of AI in Healthcare

- Early Foundations (1950s-1960s): Introduction of AI concepts in computing and their initial application in medicine.
- Expert Systems (1970s-1980s): Development of rule-based AI models such as MYCIN for bacterial infection treatment.
- Machine Learning Advances (1990s-2000s): Early applications of data-driven AI in radiology, pathology, and predictive modeling.
- Deep Learning and Big Data (2010s-Present): Integration of large datasets, deep learning models, and real-time AI applications in clinical practice.

D. Key Points on Safety and Trust in Healthcare AI

- Physician Oversight: AI should assist, not replace, physicians, maintaining human expertise as a critical component of patient care.

- Patient Safety as a Fundamental Priority:
    - AI tools must be rigorously validated to ensure they do not harm patients.
    - Physicians must maintain oversight of AI tools to ensure they are used safely and effectively.
    - The concept of "human in the loop" is critical, meaning that AI should assist, not replace, physician judgment.
- Transparency and Explainability:
    - AI systems must be transparent, meaning that physicians and patients should understand that a decision is impacted by AI.
    - The output of AI tools should be explainable, allowing clinicians to trace back the reasoning behind recommendations.
    - The "black box" problem (where AI decisions are not understandable) must be avoided to maintain trust in AI systems.
- Regulatory and Ethical Considerations:
    - Regulatory frameworks for AI in healthcare are still evolving, and there is a need for global standards to ensure safety and efficacy.
    - Ethical concerns, such as bias in AI algorithms and the potential for misuse of patient data, must be addressed.
    - Physicians should be involved in the design, development, and validation of AI tools to ensure they align with clinical needs and ethical standards.
- Data Privacy and Security:
    - Patient data used to train AI models must be protected, and patients should have the opportunity to opt out of data sharing.
    - Compliance with data protection regulations is essential to maintain patient trust.
    - There is a need for clear consent processes to ensure patients understand how their data is being used.
- Physician Oversight and Accountability:
    - Physicians remain ultimately responsible for patient care, even when using AI tools.
    - AI should augment, not replace, physician judgment, especially in complex or uncertain cases.
    - Liability frameworks must be clarified to determine responsibility when AI tools are involved in decision-making.
- Global Collaboration and Standards:
    - There is a need for international collaboration to develop consistent standards for AI in healthcare.
    - The World Medical Association (WMA) and other organizations play a key role in advocating for ethical and safe AI practices globally.

Generative AI and the Future of Medical Intelligence

Generative AI distinguishes itself from traditional AI applications by its ability to create new data rather than merely analyzing existing datasets. Unlike predictive analytics and standard machine learning models that enhance decision-making, generative AI produces clinical reports, synthesizes patient histories, and even formulates research hypotheses, making it highly valuable in medical documentation, research synthesis, and patient communication.

A. Overview of Generative AI in Healthcare

- AI models that generate novel medical content, including diagnostic reports, treatment recommendations, and research insights.
- Enhances medical research by synthesizing vast amounts of data into actionable insights.
- Supports AI-driven clinical decision-making and automates documentation workflows.

B. Large Language Models (LLMs) in Healthcare

- AI systems trained on extensive medical literature to provide contextually accurate recommendations.
- Automates administrative documentation, improving physician efficiency.
- Challenges include ensuring model accuracy, addressing biases, and maintaining patient privacy.

Will AI Replace Doctors? The Evolving Role of Physicians

Despite concerns that AI will replace physicians, its primary function is to augment medical practice by increasing efficiency and improving decision-making. While AI excels in automation and data analysis, essential aspects of healthcare—such as patient interaction, ethical decision-making, and complex case management—remain irreplaceable.

A. WMA's Statement on Augmented Intelligence

- AI is designed to enhance medical expertise, not replace human judgment.
- AI systems should be developed in alignment with ethical, regulatory, and clinical best practices.
- Continuous professional education is essential for integrating AI into medical practice effectively.

B. AI Use Cases

- Administrative automation: AI assists with scheduling, billing, and documentation.
- Medical imaging diagnostics: AI enhances the accuracy of radiology, pathology, and dermatology assessments.
- Predictive analytics: AI improves early disease detection and risk stratification for preventive care.

C. The Future of Doctor-Patient Relationships

- AI enables enhanced personalized treatment, optimizing medical decisions based on patient data.
- Physicians may increasingly act as AI interpreters, guiding patient care using AI-driven insights.
- Ethical considerations must ensure AI integration does not compromise patient trust and transparency.

Questions and Answers

1. Who is responsible for AI-driven medical decisions?

- This is still an evolving area with multiple factors influencing accountability.
  Countries are implementing processes to establish clear and consistent guidelines.
  Efforts are underway to define legal responsibility among physicians, AI developers, and healthcare institutions.
  Regulatory frameworks are evolving to address accountability concerns.

2. What ethical considerations are involved in AI use?

- AI must be developed and used transparently to maintain patient trust.
- Ethical concerns include bias in training data, patient privacy, and accountability.
- Informed consent is crucial when AI influences diagnosis or treatment decisions.

3. How can AI improve patient care without replacing doctors?

- AI can analyze vast amounts of data to support evidence-based decision-making.
- It enhances workflow efficiency by automating administrative and diagnostic tasks.
- The human element—empathy, ethical reasoning, and nuanced decision-making—remains irreplaceable.

4. What are the biggest challenges in implementing AI in healthcare?

- Ensuring AI reliability and avoiding biased or incorrect recommendations.
- Integrating AI into existing clinical workflows without disrupting care.
- Compliance with regulatory standards such as GDPR and HIPAA to protect patient data.

5. How can physicians stay informed about AI advancements?

- Engaging in continuous education programs on AI in medicine.
- Collaborating with AI developers from the very beginning to ensure practical, ethical applications.

Closing Remarks & Next Webinar Announcement

- Key Takeaways:
    - AI's transformative role in healthcare, including its limitations and ethical considerations.
    - The importance of physician oversight in AI-driven decision-making.
    - Generative AI's role in improving efficiency, research, and personalized medicine.
- Next Webinar:
    - Topic: Ethics, Legal, and Regulatory Aspects of AI in Healthcare
    - Date: February 27th
    - Why Attend? Learn about the evolving regulatory landscape and ethical frameworks guiding AI adoption in medicine.

We invite all participants to continue the discussion in the upcoming session and play an active role in shaping the future of AI-driven healthcare.

27 February 2025

Summary Document:

Ethics, Legal, and Regulatory Aspects of AI in Healthcare

WMA Educational Webinar - Medical Technologies Workgroup – 27 February 2025

Introduction

This document provides a summary of key points discussed in the second installment of the WMA Educational Webinars on Artificial Intelligence in Medicine. The session focused on ethical, legal, and regulatory aspects of AI in healthcare, exploring patient rights, accountability, and risk mitigation strategies.

Key Topics Discussed

1. Ethical Considerations in AI-Powered Healthcare

- Autonomy & the Doctor-Patient Relationship
    - AI can either **empower** patients by enhancing autonomy or **reduce** it if decisions are overly dependent on AI-generated recommendations.
    - Different global perspectives influence ethical views (e.g., Western consumerism vs. paternalistic models in China).
- Informed Consent Challenges
    - Traditional informed consent involves doctor-patient discussions.
    - AI's "black-box" nature makes it difficult for physicians to explain how AI arrived at a decision, raising concerns about whether true informed consent is possible.
    - Some legal frameworks (e.g., EU regulations) mandate disclosure when AI is used in patient care, while others (e.g., US law) do not explicitly require it.
    - The level of AI intervention significantly impacts informed consent.
    - It's essential to consider whether patients should be informed about AI use in their care, balancing transparency with the risk of overwhelming them with information.

    - The "explainability" of AI is crucial when AI significantly influences healthcare decisions.
- Loss of Human Touch in Medicine
    - AI-driven healthcare may reduce direct physician-patient interaction, leading to concerns about dehumanization.
    - There is potential for over-reliance on AI, which can introduce confirmation bias.

2. Legal and Regulatory Challenges for AI in Healthcare

- Liability & Accountability
    - Three primary perspectives on liability:
        - Physician Liability – Doctors are traditionally held accountable if AI-driven decisions harm patients.
        - Institutional Liability – Hospitals may be responsible if they implement AI tools in ways that contribute to harm.
        - Developer Liability – AI manufacturers and software developers could be held accountable, though legal precedents are still evolving.
    - In the future, AI could potentially influence the standard of care, shifting liability from doctors to institutions or developers.
- Regulating AI Across Its Lifecycle
    - Three key stages where regulation is needed:
        - Research & Development: Ensuring ethical AI design, data bias mitigation, and proper validation.
        - Market Approval: Regulatory approvals vary across countries (e.g., FDA in the US, MDR in the EU, SFDA in Saudi Arabia).
        - Post-Market Oversight: Mechanisms to hold AI accountable after deployment (e.g., WHO global governance frameworks).
- Public Sector vs. Private Sector Regulation
    - Different countries adopt different approaches:
        - Free-market-driven approach with minimal AI regulation to encourage innovation.
        - Strong regulatory frameworks prioritizing patient protection.
        - AI governance is largely centralized, with government oversight playing a key role.
- Documentation: Hospitals may need subcommittees to set internal documentation standards for AI use.
- Best practices: Until clear guidelines emerge, physicians must exercise clinical judgment and understand the standard of care, as relying solely on AI may not be a legal defense.

Privacy and Rights in AI-Assisted Healthcare

- Data privacy laws: Legal frameworks, such as GDPR, impact AI development and deployment, but they also have gaps, such as exceptions for research and public benefit.
- Data security: Blockchain technology may offer a path to securing health data, with systems allowing patients to approve data access requests.
- Patient data ownership: Patients should maintain ownership over their data, raising questions about data removal from AI systems and models. Balancing the desire to create comprehensive AI systems with the need for patient autonomy and control is a real tension.
- Intellectual property: AI developers may protect their algorithms as trade secrets, complicating transparency and regulation.
- Global variability: There is tremendous variation in what developers do to ensure data is unbiased and patients have recourse if harm is caused.

3. Risk Mitigation Strategies for AI in Clinical Practice

- Best Practices for Physicians Using AI Tools
    - Physicians should document AI-assisted decision-making in patient records.
    - Human-in-the-loop (HITL) models should ensure physician oversight in AI-generated recommendations.
    - AI tools should augment, not replace clinical judgment.
- Patient Data Protection & Security
    - AI relies heavily on patient data, raising privacy concerns.
    - Key data risks:
        - Bias in Training Data – AI can perpetuate biases if trained on unrepresentative datasets.
        - Data Privacy – Regulations in different regions attempt to protect patient rights, but gaps remain.
        - Security Threats – Blockchain technology is being explored for securing patient data.
    - Should patients have the right to remove their data from AI models?
        - While AI requires large datasets, patients should maintain ownership and be informed about how their data is used.
- Human Oversight & AI Transparency
    - AI should not function as an independent actor in patient care. Instead:
        - Explainability: AI systems must provide reasoning for decisions where possible.
        - Regulatory Alignment: Healthcare institutions should align AI usage with existing medical guidelines.

    - **Ethical Oversight:** AI should be reviewed by hospital ethics committees before clinical deployment.

Closing Thoughts & Next Steps

- AI in healthcare presents **unprecedented opportunities and challenges**.
- Legal and ethical frameworks are still evolving, with **no single global standard yet**.
- Clinicians must remain engaged in discussions about AI governance to ensure **safe, ethical, and effective AI integration**.

Next Webinar in the Series

📅 **March 27, 2025** – *"Current and Future Applications of AI in Medicine"*

Speaker Introduction: TBD

Topics will include:

- Breakthrough AI applications in diagnostics and treatment
- AI's role in precision medicine and drug discovery
- Challenges in AI adoption across healthcare settings

Feedback & Further Reading

📩 **Feedback Form:** Please share your thoughts via the link.

📖 **Recommended Reading:**

- *Research Handbook on Health, AI, and the Law* – Available online (Open Access).
- *Informed Consent & AI in Healthcare* – Article by Prof. Glenn Cohen.

For further inquiries, contact **WMA at wma@wma.net**.

FAQs:

Ethical, Legal, and Regulatory Aspects of AI in Healthcare

1. What are the primary ethical considerations when using AI in medicine?

Core ethical principles like autonomy, beneficence, non-maleficence, and justice are crucial. AI's impact on the doctor-patient relationship is a key concern. Does AI increase patient autonomy by giving them more control over their data and care, or does it lead to a more paternalistic dynamic where doctors overly rely on AI-driven insights? There is also a risk of ceding too much authority to AI developers who set the standards for AI device usage if regulations are inadequate. The potential loss of the human touch is also a major consideration and must be addressed to maintain patients' trust and the therapeutic alliance with their doctors.

2. How does AI impact informed consent in healthcare?

The traditional informed consent process involves a doctor explaining treatment options and the reasoning behind their recommendations to a patient. When AI is involved, especially in decision-making, the explainability of the AI's reasoning becomes crucial. If the AI's decision-making process is too complex to understand (even for its developers), how can a doctor adequately explain the basis of the recommendation to a patient, fulfilling the requirements for informed consent? The level of AI involvement in the medical decision greatly impacts informed consent. Transparency about the AI's role and limitations is essential.

3. Should patients be informed about the use of AI in their clinical care, and to what extent?

There are varying opinions. Some argue that informing patients about every instance of AI use (e.g., in radiology for dose reduction) would be overwhelming and ineffective. Others believe transparency is necessary, especially when AI significantly influences treatment recommendations.

4. Who is liable if a patient is harmed by AI in healthcare?

Liability is a complex issue with several potential actors: the doctor, the healthcare institution (hospital/clinic), the AI developer, and even (controversially) the patient. Currently, existing tort law dictates that doctors are responsible if their actions fall below the standard of care, regardless of whether AI was used. Hospitals may also be liable under vicarious liability if doctors are hospital employees. The liability of AI developers is less clear, but is something that is beginning to garner more attention. It is likely to increase in prominence in coming years. The implementer, who connects the system, validates the data, and turns things on, also carries some shared liability.

5. What steps can healthcare professionals take to mitigate liability risks associated with AI?

As of right now, clinicians must take responsibility for medical decisions even when relying on AI, but in the future this may not always be the case. AI setting the standards of care could shift liability elsewhere. Clinical judgment is paramount. Careful documentation of AI's role in decision-making is crucial. Hospitals should form subcommittees to establish internal principles, guidelines, and standards for AI implementation and documentation.

6. What aspects of AI in healthcare should be regulated, and why?

Regulation is needed across AI's entire life cycle, from research and development to deployment and post-market surveillance.

Areas for regulation include:

• Research and Development: Ensuring developers adhere to best practices, such as engaging with patient groups and using appropriate data sources to mitigate against biases.

• Market Approval: Establishing clear requirements for regulators to assess AI devices, not just regarding risk, but also issues like medical liability and informed consent.

• Clinical Practice: Creating mechanisms to ensure patients can exercise their medical law rights (informed consent, liability) when AI is used in their care.

7. What are the potential risks of AI to the physician-patient relationship?

AI could potentially dehumanize healthcare by leading to over-reliance on technology and a loss of human touch. Doctors may become overly reliant on AI, exhibiting confirmation bias and potentially overlooking crucial information. There is concern AI may lead doctors to become "lazy" and simply rely on AI decisions without question. It is critical to consciously guard against that.

8. How can patient data be secured and ethically used in AI systems?

Data bias, privacy concerns, and security vulnerabilities must be addressed.

• Data Bias: Careful data selection in the research phase is critical to avoid building AI systems that are ineffective or discriminatory.

• Data Privacy: Existing data protection laws (e.g., GDPR) provide a framework, but local authorities should supplement these with healthcare-specific guidelines.

• Data Security: Technologies like blockchain can enhance data security by giving patients greater control over who accesses their data. The issue of intellectual property rights surrounding algorithms can complicate data privacy regulations.

27 March 2025

Notes and Takeaways from WMA Webinar: CURRENT AND FUTURE APPLICATIONS OF ARTIFICIAL INTELLIGENCE IN MEDICINE

Overview

This webinar explored the current and emerging applications of Artificial Intelligence (AI) in medicine, including diagnostics, decision support, patient engagement, drug development, and data governance. The conversation also addressed regulatory and ethical considerations surrounding AI adoption in healthcare.

The session provided both a technical and strategic look at how AI is transforming clinical workflows, supporting physicians, and potentially redefining care delivery through automation, personalization, and augmentation of medical decision-making.

Current Medical Applications

1. Diagnostic Support and Accuracy

- AI systems can now emulate expert reasoning, such as that of a seasoned radiologist.
- FDA-approved AI tools are widely available to assist physicians in radiology, dermatology, and pathology, with proven gains in diagnostic accuracy when combined with human oversight.
- Example: In rheumatology, AI classified capillary images with performance nearly equal to expert rheumatologists.

2. Predictive Analytics and Clinical Decision Support

- AI models developed in Zurich predict clinical outcomes such as ICU delirium, lung function decline, and the need for medication adjustments in chronic conditions.
- LLMs (Large Language Models) like ChatGPT have demonstrated >90% accuracy on US medical licensing exam questions, surpassing average student performance.

3. Conversational Agents in Clinical Workflows

- AI-powered assistants can support physicians in creating radiology reports, translating text, and navigating imaging data, increasing efficiency and reducing screen time.
- Example: A radiologist assistant prototype allows interactive queries such as "show me pathology" or "write a report."

4. Operational and Logistical Applications

- AI tools help optimize workforce allocation (e.g., predicting nursing demand) and reduce no-show rates (e.g., MRI appointments), although human behavior remains a limiting factor.

5. Education and Support Tools for Physicians

- AI is being integrated into medical education curricula to support clinical reasoning, though only after traditional learning to ensure foundational knowledge is established.

Future Applications & Trends

1. Personalized and Precision Medicine

- AI enables personalized care by identifying patient cohorts with similar genetic or clinical profiles.
- Example: In oncology, AI helps match patients beyond standard treatment guidelines using data-driven similarity searches.

2. Drug Development and Discovery

- Tools like AlphaFold are revolutionizing protein structure prediction, expediting target identification in drug development.
- AI also assists in designing CRISPR-based gene editors tailored to individual mutations.

3. Closed-loop Systems and Medication Delivery

- AI-controlled drug delivery systems (e.g., for anesthesia or hypertension) have potential, though regulatory and trust barriers remain high.
  Voice-based diagnostics (e.g., for schizophrenia) and wearable-guided medication

management are emerging.

4. Patient Engagement and Remote Monitoring

- AI-powered virtual assistants and chatbots offer triage, education, and behavioral nudges, though long-term engagement remains challenging.
- Embedded passive sensors may offer a solution for compliance without active user interaction.

5. Digital Twins and Future Care Models

- While digital twins are conceptually valuable, they are currently limited by data availability and system integration challenges.
- The future may include AI-guided care for low-risk pathways, with human oversight reserved for complex cases.

Key Quotes or Insights

- “We’re trying to emulate expert thinking and put that into computers.”
- “AI and physicians both make mistakes. Together, they can reduce them.”
- “We expect AI to be unbiased, but we forget that humans are biased too.”
- “AI will support, not replace physicians—especially by taking over simpler, repetitive tasks.”
- “Patients are gaining access to expert-level knowledge through AI—this changes everything.”
- “Trust and transparency are critical. Physicians must know when AI is being used and why.”

Summary of Takeaways

- AI Augments, Not Replaces: AI excels at supporting physicians in diagnostics, documentation, and decision-making but does not eliminate the need for human judgment.
- LLMs Show Promise: Tools like ChatGPT achieve high accuracy on medical exams and can support clinical reasoning if properly validated.
- Workload Relief: Automation of repetitive tasks improves efficiency and may help address workforce shortages and burnout.
- Precision Medicine at Scale: AI facilitates personalized treatment by analyzing large patient datasets and predicting outcomes.
- Operational Value: AI improves hospital efficiency through tools for staffing, scheduling, and predictive logistics.
- Behavioral Engagement Remains a Challenge: AI tools are less effective for long-term behavior change and must be designed for seamless integration into daily life.
- Digital Ethics and Trust: Transparency, bias mitigation, and human oversight are key to ethical AI deployment.
- Cybersecurity and Data Privacy: Strong safeguards, anonymization, and secure platforms are essential, especially when using external cloud services.
- Digital Twins and Closed-loop Systems: These are on the horizon but require significant advances in data quality and regulatory clarity.
- AI in Education: Integrating AI into curricula must be balanced to preserve critical thinking skills and clinical reasoning.

Q&A from the Webinar:

Question:
Do LLMs have the capacity of quantifying the degree of certainty in the answers they give?
Answer:
Yes, certain methods and semantic-embedding techniques enable LLMs to estimate confidence scores by analyzing response variation or semantic coherence, helping identify unreliable answers. These techniques are being refined to improve trust and interpretability in medical contexts.

Question:
Will AlphaFold and new models help find cures for genetic diseases like Alpha-1 antitrypsin deficiency?
Answer:
Yes, models like AlphaFold and AlphaMissense are advancing the identification of pathogenic mutations, while CRISPR-based methods have shown potential in correcting DNA mutations in alpha-1 antitrypsin deficiency, accelerating the development of curative therapies.

Question:
Could future models trained on AI notes instead of physician notes eliminate bias?
Answer:
Unclear based on available sources. While AI-generated notes might reduce individual human biases, they risk reproducing and amplifying systemic biases unless rigorously audited and corrected during training.

Question:
Will understanding AI's "black box" improve diagnostic learning?
Answer:
Yes, interpretable AI approaches like semantic uncertainty quantification can provide insights into reasoning patterns, potentially enhancing clinician learning and diagnostic decision-making frameworks.

Question:
How can AI integration address LMIC-specific medical challenges?
Answer:
AI in LMICs must be adapted for limited resources, using low-compute models, mobile-first tools, and localized datasets to improve diagnostics and access. Shared benefits include better triaging and telemedicine support, but infrastructure gaps remain a key barrier.

Question:
What hidden biases affect AI fairness in medicine?
Answer:
Biases can stem from non-representative training data, clinician labeling errors, and failure to

include social determinants of health. WMA can help by promoting global standards for AI fairness, transparent auditing, and equity-centered model development.

Question:
Do health insurers use AI to override physicians and deny coverage?
Answer:
Yes, AI tools are increasingly used by insurers for prior authorization, with reports indicating denial rates significantly higher than manual reviews and lawsuits alleging AI-led systematic coverage denials.

30 April 2025

Best Practices in Medical AI Development

WMA Educational Webinar Series on Artificial Intelligence in Medicine April 30, 2025

1. From Concept to Clinic: Understanding the AI Development Journey

- Start with real clinical problems, not technology hype. Clinical problems are similar worldwide, though workflows and administrative processes differ between institutions. Use design thinking: interview and observe stakeholders to thoroughly understand the real problem before developing solutions.

- Validate the problem before seeking technological solutions. Example: What appeared to be a hospitalist-physician communication issue was discovered to be specifically about discharge processes. Creating follow-up appointments from the hospital to outpatient services led to 93% patient attendance and 12% reduction in readmissions.

- Consider the appropriate technology approach. Early AI applications focused on imaging as data was already digital. Before generative AI (pre-2023), machine learning required thousands of data points to prove a single algorithm. Generative AI has transformed healthcare technology implementation, requiring less data and offering more flexibility.

- Build minimum viable products by adapting existing technologies when possible. Example: Defense industry image analysis technology designed to detect changes in maps was successfully adapted for mammogram analysis to identify year-over-year changes, flagging differences for physician review rather than attempting specific diagnoses.

- Test thoroughly before implementation. Always obtain IRB approval and conduct comparative studies with and without the technology. Successful integration depends on embedding tools directly into clinical workflow, such as AI algorithms for brain hemorrhage detection that prioritize abnormal CT scans with color-coding.

- Focus on user adoption and workflow integration. Technology should integrate smoothly into existing clinical processes. Success is evident when clinicians actively request the technology, as happened when radiologists wanted the AI tool installed on their personal computers after seeing its value.

2. The Clinician's Role in a Health Tech Team

- Physicians are essential for meaningful healthcare technology development. Technology specialists often don't understand the complexity of medical care, ethical considerations, and patient relationships. Example: A CEO without healthcare background assumed a UTI consultation could be reduced to under one minute, not understanding that even simple cases involve complex patient contexts.

- Clinicians serve multiple critical roles in technology development: domain experts for clinical functionality, gatekeepers for patient safety, evaluators of technology fit, and advocates for ethical implementation. Their insight is crucial from concept through deployment for successful integration.

- Create effective multidisciplinary teams by bringing together medical specialists, data scientists, engineers, and IT experts. Develop structured evaluation systems to assess technologies from multiple perspectives (clinical value, integration feasibility, workflow impact).

- Bridge the gap between startup speed and healthcare institution caution by identifying and modifying bureaucratic barriers. Example: Reducing contract length from 60+ pages to 3 pages facilitated startup relationships and reduced legal costs. Provide project managers to help companies navigate hospital systems and regulatory processes.

- Clinicians must be assertive about their value in technology development. Companies that fail to incorporate physician perspectives often struggle to create successful healthcare solutions. Physician involvement is essential for understanding clinical thinking and workflow integration.

3. Ensuring Safety, Ethics, and Trust

- Patient safety is non-negotiable. Always follow proper regulatory processes (IRB, FDA) for any technology involving patient care. Run comparative studies to demonstrate effectiveness and safety. Example: AI tools for detecting lung nodules in ER chest X-rays reduced missed diagnoses and related lawsuits by automatically flagging concerning images.

- Maintain the human-in-the-loop approach where AI makes suggestions but clinicians maintain control and oversight. Physicians must remain responsible for final decisions in all patient care situations. Key evaluation question: "Would I use this AI tool for a family

member? If not, it's not ready for implementation."

- Ensure transparency and explainability in AI systems. Understand what data was used to train algorithms and be alert to "black box" solutions that can't be explained. Technology must be continuously monitored for algorithmic drift (when algorithms change behavior over time) and bias.

- Address accountability considerations. AI can both reduce liability (by catching issues humans might miss) and create new concerns. Technology must be continuously monitored, and institutions should establish clear lines of responsibility for AI-assisted decisions.

- Design for equity and ethical implementation by vigilantly monitoring for biases in data and algorithms, ensuring technology works for diverse populations, and implementing proper data handling and privacy protections.

Future Directions

- Medical education must evolve to include technology evaluation skills for future healthcare professionals. The goal is not to teach every technology (impossible given rapid changes) but to develop critical thinking about technology implementation.

- AI can enhance patient-centered care by reducing administrative documentation burden, transcribing and summarizing information, enabling pre-visit symptom checking, and supporting personalized medicine approaches through better data analysis.

Conclusion

AI in medicine is not something being done to clinicians but must be built with their active involvement. Healthcare professional expertise, caution, and vision are critical for ensuring these technologies enhance patient care rather than compromise it. The successful integration of AI into healthcare requires ongoing clinician involvement, careful validation, and a commitment to maintaining human judgment and oversight throughout the process.

22 May 2025

WMA Webinar #5 – AI for Health Equity: Bridging the Global Divide - Key Notes and Takeaways

Overview

This session focused on AI for global health and equity, exploring how technology can bridge the gap between high and low-resource settings. The discussion, led by Professor Annie Hartley, highlighted the use of open-source, participatory frameworks to make clinical AI inclusive, adaptable, and globally relevant.

Current Challenges & Opportunities

- The AI Opportunity in Low-Resource Settings: While often seen as difficult environments for technology, low-resource settings are arguably the best places to introduce AI because there are often no alternatives for care. The goal is to provide access where it currently does not exist.
- Data Inequity: Global data is inequitably available; for example, less than 3% of PubMed content is pertinent to Africa. This leads to models that are inaccurate for underrepresented populations.
- Imperfect Models: We cannot wait for perfect data digitization; instead, we must accept imperfect models and adapt them through continuous use and validation.

Core Solutions: Meditron & MOOVE

- Meditron (The Model): An evolving suite of open-source medical Large Language Models (LLMs) adapted specifically for low-resource settings.
    - These models are designed to be efficient, with smaller versions capable of running on local hardware (e.g., phones) without internet or heavy infrastructure.
    - Despite their smaller size, they perform competitively on medical licensing exams.
- The MOOVE Platform (The Method): Stands for Massive Open Online Validation and Evaluation.
    - It allows clinicians to "nudge" or tune global models to their local context by rating answers and providing feedback.
    - This process promotes local ownership and co-development, ensuring the AI reflects local medical standards and language.

Implementation & Ethical Considerations

- Vigilance over Replacement: The narrative of AI replacing doctors is irrelevant in many global health contexts where millions never see a doctor at all. Training should focus on "vigilance"—teaching physicians to know when the tool works and when it fails.
- Privacy and Sovereignty: Local ownership is key to privacy. By using open-source models hosted locally, institutions can maintain data sovereignty and do not need to send patient data to external commercial servers.
- Validation Approach: Rather than waiting for new global regulations, AI deployment should follow existing Good Clinical Practice (GCP) standards, treating deployment like a clinical trial to monitor efficacy and safety.

Summary of Takeaways

- Bridging the Divide: AI can provide essential medical knowledge to regions with severe physician shortages, serving as a critical access tool rather than just an efficiency tool.
- Local Adaptation: Tools like MOOVE allow clinicians to adapt global models to local contexts without expensive retraining, addressing data bias.
- Open Source Security: Open-source models allow for local hosting, which is superior for patient privacy and prevents dependence on changing commercial APIs.
- Physician Responsibility: Liability remains with the physician, similar to using a medical textbook; the human must interpret the information responsibly.
- Actionable Step: Clinicians are encouraged to participate in validating and tuning models for their specific regions to ensure fair representation.