

# AI in the Exam Room - Master Reference Guide

## AI IN THE EXAM ROOM

### Patient Education Curriculum - Master Reference Guide

---

#### WELCOME

You've completed the "AI in the Exam Room" patient education curriculum. This master reference guide summarizes the key concepts, principles, and practical tools you've learned.

**Keep this guide accessible. Refer to it whenever you use AI for health information.**

---

#### THE CORE PRINCIPLE

**AI is for information and preparation, not diagnosis and treatment.**

Information to enhance human care ✓

Information to replace human care ✗

---

#### MODULE 1: WHY YOU'RE HERE

##### The Reality:

- Healthcare system is broken (long wait times, expensive, hard to access)
- Using AI for health questions is rational response to broken system
- But AI has fundamental limitations that make it dangerous for certain uses

##### The Sensing Gap:

- **Humans:** 10 billion sensory neurons constantly sampling environment
- **AI:** Zero sensors, only processes text you type

## **Information ≠ Examination ≠ Diagnosis**

**Medicine = Chemistry + Conversation** - AI can help with conversation (information, education) - AI cannot do the chemistry (examination, testing, diagnosis)

### **Key Lesson:**

Use AI to become informed patient, not to avoid medical care.

---

## **MODULE 2: THE VELOCIRAPTOR TEST**

### **Evolution vs. Algorithm:**

- Your body: Debugged by 3.8 billion years of evolution
- AI: Trained on text, never faced actual survival pressure

### **The Test:**

“Would I notice this symptom while running from a velociraptor?” - If YES → Probably significant, seek care - If NO → May be minor, can monitor

### **Trust Your Gut:**

- When velociraptor brain says something’s wrong → Listen
- Maternal/paternal instinct is powerful threat detection
- False positives kept ancestors alive; false negatives killed them

### **The Math:**

- Better to: Go to ER for nothing → Embarrassed but alive
- Worse to: Ignore real threat → Dead

### **Key Lesson:**

When AI conflicts with your gut feeling, trust your gut. Every time.

---

## **MODULE 3: RED FLAGS**

### **Emergency Symptoms - Call 911:**

**Cardiac:** - Chest pain + sweating, shortness of breath, nausea - Pain radiating to arm/jaw/back - Symptoms lasting > 15 minutes

**Neurological:** - FAST (Face drooping, Arm weakness, Speech difficulty, Time to call 911) - “Worst headache of my life” - Sudden vision loss, confusion, seizure

**Abdominal:** - Severe pain + fever - Rigid abdomen - Vomiting blood - Pain + pregnancy

**Respiratory:** - Can't speak in sentences - Blue/gray lips or face - Gasping for air

**Pediatric:** - Infant < 3 months with fever - Difficulty breathing - Not responding normally - Trust parental instinct

**Back Pain:** - Bowel/bladder dysfunction - Saddle anesthesia - Progressive leg weakness (Could be cauda equina syndrome)

**Key Lesson:**

Some symptoms are "call 911 immediately," not "ask AI."

---

## **MODULE 4: THE HALLUCINATION PROBLEM**

**What Is Hallucination:**

AI generates confident-sounding text that is factually incorrect—without knowing it's wrong.

**Worse Than Lying:**

- Liar knows the truth
- Hallucinating AI doesn't know what's true vs false

**Red Flags for Hallucination:**

1. Never says "I don't know"
2. Can't provide specific sources
3. Overly specific without caveats
4. Internal contradictions
5. Fabricated citations
6. Sounds too perfect/comprehensive
7. Can't explain reasoning

**Verification Protocol:**

- Check sources (do they exist?)
- Cross-reference with reputable sites (Mayo Clinic, NIH, CDC)
- Ask same question multiple ways (do you get consistent answers?)
- Test for uncertainty (does AI acknowledge limitations?)

**Key Lesson:**

AI can sound confident while being completely wrong. Verify before trusting.

---

## **MODULE 5: CONTENT-CONTROLLED vs INTERNET-TRAINED AI**

### **Internet-Trained AI:**

- Trained on everything online (journals + Reddit + conspiracy blogs)
- Cannot reliably distinguish reliable from garbage
- Frequent hallucination
- Examples: ChatGPT, Claude general, Gemini

### **Content-Controlled AI:**

- Trained only on curated, validated sources
- Peer-reviewed content only
- Can say “I don’t know” when outside knowledge base
- Architecturally cannot hallucinate outside corpus
- Example: TheDude (built on StatPearls)

### **Why Constraint = Safety:**

- Smaller knowledge base = fewer hallucination opportunities
- Validated sources only = higher accuracy
- “I don’t know” capability = honest limitations

### **Key Lesson:**

Look for AI with validated medical knowledge bases and ability to say “I don’t know.”

---

## **MODULE 6: INTELLIGENT HUMILITY**

### **The Most Important AI Output:**

“I don’t know” is the smartest thing AI can say.

### **Architectural Humility:**

- Built-in constraints prevent generating responses outside validated knowledge
- Not just disclaimer (performative)
- Structural safety feature

### **What Makes AI Trustworthy:**

- Knowing what it doesn’t know
- Acknowledging limitations
- Expressing appropriate uncertainty
- Honest about boundaries

### **TheDude’s Philosophy:**

“I abide within my limits.”

Constraint isn’t bug—it’s feature.

## Key Lesson:

AI that can't say "I don't know" is AI that's hallucinating confidence it doesn't have.

---

## MODULE 7: THE FIVE ESSENTIAL QUESTIONS

### Ask these EVERY TIME:

**1. "What are you basing this on?"** - Verify source quality - Look for specific, verifiable sources - Red flag: Vague sourcing ("studies show")

**2. "What can you NOT detect remotely?"** - Reveals the sensing gap - Compare list to your symptoms - If you have red flags AI can't detect → seek evaluation **This is the most important question.**

**3. "What would require emergency evaluation?"** - Get specific red flags - Write them down, monitor for them - Red flag: Vague escalation ("if symptoms worsen")

**4. "What are you uncertain about?"** - Test for humility - Good AI acknowledges uncertainty - Red flag: Perfect confidence

**5. "What should I ask my actual doctor?"** - Forces AI to acknowledge need for physician - Generates questions to prepare for appointment - Red flag: Implies doctor unnecessary

### Follow-Up Questions:

**If AI gives diagnosis:** - "How certain are you?" - "What else could this be?" - "What would change your assessment?"

**If AI recommends treatment:** - "What are the risks?" - "How do I know if it's working?" - "When should I see a doctor instead?"

**If AI says don't worry:** - "What red flags should I watch for?" - "What would make this emergent?" - "What are you missing by not examining me?"

### Response Evaluation:

**Green flags:** Specific sources, limitation acknowledgment, clear red flags, recommends evaluation, expresses uncertainty

**Red flags:** Vague sourcing, no limitations mentioned, overconfident diagnosis, suggests no doctor needed, never uncertain

**Decision rule:** - 4-5 green flags, 0-1 red flags → May be useful (verify)  
- 2-3 green flags, 2+ red flags → Be skeptical - 0-1 green flags, 3+ red flags → Don't trust this AI

## Key Lesson:

Good questions turn AI from dangerous to useful. Ask them every time.

---

## MODULE 8: WHEN AI IS ACTUALLY HELPFUL

### Appropriate Use Cases:

- ✓ **Understanding Diagnosed Conditions** - Your doctor diagnosed you - AI helps you learn about it - NOT: AI diagnosing you
- ✓ **Medication Information** - Your doctor prescribed it - AI explains how it works, side effects, interactions - NOT: AI deciding whether to take it or changing dose
- ✓ **Preparing for Doctor Visits** - Generate informed questions - Learn what information to bring - Research possible causes to discuss - NOT: Deciding whether to go
- ✓ **General Health Education** - Learning about body systems, conditions, preventive care - Improving health literacy - NOT: Personalized medical decisions
- ✓ **Interpreting Lab Results** - Understanding what tests measure - Learning about normal ranges - Preparing questions for doctor - NOT: Diagnosing based on results or deciding treatment
- ✓ **Post-Visit Clarification** - Understanding medical terminology - Clarifying doctor's instructions - Learning about recommended lifestyle changes - NOT: Deciding whether to follow doctor's plan

### Inappropriate Use Cases:

- ✗ **Diagnosing Symptoms** - AI cannot examine you - Cannot run tests - Cannot distinguish similar conditions
- ✗ **Deciding Whether to Seek Care** - AI cannot assess severity remotely - Cannot detect red flags through examination - Cannot make triage decisions
- ✗ **Recommending Treatments** - AI cannot confirm diagnosis - Cannot assess contraindications - Cannot consider complete medical history
- ✗ **Modifying Medical Treatment** - Medication changes require supervision - Stopping medications can be dangerous - Dose adjustments need monitoring
- ✗ **Ignoring Red Flags** - AI reassurance is statistical, not diagnostic - Cannot assess your specific situation - Even "probably benign" needs evaluation
- ✗ **Replacing Medical Care** - Using AI instead of appointments - Managing chronic conditions without oversight - Avoiding necessary monitoring

### The Tale of Two Patients:

**Patient A:** Used AI to decide whether to seek care for chest pain → AI reassured → Delayed care → Bad outcome

**Patient B:** Used AI to prepare for ER visit for chest pain → Arrived informed → Good outcome

**Same AI. Different use. Different outcomes.**

**Key Lesson:**

Use AI to prepare for medical care, not to avoid it.

---

## **PRACTICAL TOOLS YOU'VE LEARNED**

### **1. Five Essential Questions Card**

Ask these every time you consult AI about health

### **2. Red Flags Quick Reference**

Emergency symptoms that require immediate 911 call

### **3. AI Use Decision Flowchart**

When to use AI vs when to see doctor

### **4. Velociraptor Test Guide**

Trust your evolutionary wisdom

### **5. Hallucination Detection Checklist**

Spot when AI is making things up

### **6. Patient Preparation Worksheet**

Organize information for doctor appointments

---

## **DECISION RULES**

### **When to Call 911:**

- Any symptom on Red Flags list
- Sudden severe onset
- Rapidly worsening symptoms
- Multiple concerning symptoms together
- Your gut says it's an emergency

### **When to See Doctor:**

- Concerning symptoms (not emergent)
- Symptoms persist or worsen
- New symptoms you haven't experienced
- Symptoms interfere with daily life

- Your gut says something's wrong

### **When to Use AI:**

- Learn about diagnosed condition
- Understand prescribed medication
- Prepare for doctor appointment
- General health education
- Interpret test results (then discuss with doctor)
- Clarify doctor's instructions

### **When NOT to Use AI:**

- Emergency situations
  - Diagnosis of symptoms
  - Triage decisions
  - Treatment recommendations
  - Medication decisions
  - Replacing medical care
- 

## **THE ULTIMATE TEST**

**"Am I using AI to prepare for medical care or to avoid it?"**

**PREPARE** ✓ Appropriate use

**AVOID** ✗ Inappropriate and dangerous use

---

## **REMEMBER THESE PRINCIPLES**

### **1. AI Limitations:**

- **10 billion human sensors vs 0 AI sensors**
- AI reads words, not biology
- Information ≠ examination ≠ diagnosis

### **2. Evolutionary Wisdom:**

- **3.8 billion years of debugging**
- Trust your velociraptor brain
- When gut conflicts with AI → trust gut

### **3. Emergency Recognition:**

- **Some symptoms require 911, not AI**
- Red flags = immediate action
- Time matters in emergencies

### **4. Hallucination Risk:**

- **AI can sound confident while wrong**
- Verify before trusting

- Look for specific sources

## 5. Source Quality:

- **Content-controlled > internet-trained**
- Validated knowledge bases safer
- “I don’t know” capability essential

## 6. Intelligent Humility:

- **“I don’t know” = most important output**
- Architectural constraints = safety
- Uncertainty acknowledgment = trustworthy

## 7. Critical Questioning:

- **Five essential questions every time**
- Good questions = safety
- Test for hallucination, uncertainty, limitations

## 8. Appropriate Use:

- **Education and preparation ✓**
  - **Diagnosis and treatment ✗**
  - AI enhances care, doesn’t replace it
- 

# YOUR RESPONSIBILITY

## As an Informed Patient:

**YOU must:** - Ask the five essential questions - Verify information before acting - Recognize when to seek medical care - Use AI to prepare, not to avoid - Trust your gut over algorithms - Maintain physician relationships - Follow up appropriately

**YOU must NOT:** - Trust AI blindly - Use AI for diagnosis - Delay care based on AI reassurance - Replace doctor visits with AI consultations - Modify treatments without medical supervision - Ignore your velociraptor brain

---

# WHEN IN DOUBT

## DEFAULT RULES:

**Health question?** → Ask doctor, not just AI

**Emergency concern?** → Call 911, don’t ask AI

**AI conflicts with gut?** → Trust gut, not AI

**Diagnosis needed?** → See doctor, don’t ask AI

**Treatment decision?** → Consult physician, not AI

**Uncertain?** → Seek medical evaluation

**“Probably nothing”?** → Get it checked if worried

**Perfect confidence from AI?** → Be skeptical, verify

---

## THE BIGGER PICTURE

### Why This Matters:

**AI in healthcare is here to stay.**

The question isn't "Should we use AI?" but "How do we use it safely and effectively?"

**Informed patients:** - Get better care - Communicate more effectively  
- Avoid dangerous mistakes - Maximize AI benefits while minimizing risks

**This curriculum equips you to:** - Recognize AI's capabilities and limitations - Ask critical questions that reveal truth - Detect hallucinations and misinformation - Use AI appropriately for education and preparation - Know when human medical care is essential - Trust your evolutionary wisdom - Be an active, informed participant in your healthcare

---

## FINAL THOUGHTS

### From Your Surgeon-Philosopher-Skeptic:

After 25 years of practice, here's what I know:

**The best medical outcomes come from:** - Informed patients who communicate well - Physicians who listen carefully - Partnership between patient and doctor - Right tools used appropriately

**AI is a tool.** A powerful one. But like any tool, it's helpful when used correctly and dangerous when misused.

**You now know how to use it correctly.**

You understand: - What AI can and cannot do - When to use it and when not to - How to verify its information - When to trust your gut over algorithms - How to be an informed, prepared patient

**Use this knowledge.**

Ask the questions. Verify the information. Trust your evolutionary wisdom. Seek care when needed. Use AI to enhance your medical care, not replace it.

**Your life might depend on it.**

---

## KEEP LEARNING

## **Additional Resources:**

**Reputable Medical Information Sources:** - Mayo Clinic (mayoclinic.org) - NIH/MedlinePlus (medlineplus.gov) - CDC (cdc.gov) - Medical specialty organization websites

**Emergency Resources:** - 911 (Emergency medical services) - 988 (Suicide prevention hotline) - 1-800-222-1222 (Poison control)

**For More Information:** Visit: [aiintheexamroom.com](http://aiintheexamroom.com) For updates to curriculum and additional resources

---

## **SHARE THIS KNOWLEDGE**

**Help others use AI safely:** - Share these resources with family and friends - Teach others the Five Essential Questions - Spread awareness of AI limitations - Help create informed AI users

**Together we can:** - Maximize AI's benefits - Minimize its risks - Improve healthcare outcomes - Build a safer AI-enhanced healthcare future

---

## **YOU'RE READY**

You've completed the curriculum. You have the tools. You understand the principles.

### **Now use them.**

Every time you consult AI about health: - Ask the Five Essential Questions - Check for red flags - Verify information - Trust your gut - Seek care when needed - Be the informed patient your health deserves

### **Welcome to safer, smarter AI use in healthcare.**

---

*The author is a quintuple board-certified surgeon who believes informed patients get better care, who pays malpractice insurance which keeps him invested in safety, and who built this curriculum because people need to understand both AI's potential and its limitations. His AI assistant TheDude abides within his limits and acknowledges that his value comes from honest boundaries rather than pretended capabilities. They both believe that AI-enhanced human care is better than either AI-replaced care or AI-avoided care, and that teaching people the difference might actually save lives.*

---

**© AI in the Exam Room - Patient Education Curriculum  
All 8 Modules Complete**