“AI is Cool/Tricky/Risky/ Great/Scary/Useful etc”

How I have tried to navigate/use it

James McCormack
BSc (Pharm), PharmD
Professor
Faculty of Pharmaceutical Sciences
University of British Columbia
Vancouver, BC, Canada

Some of this talk was made better (I think) by using AI
I am responsible for all the content

AI is ...

Cool / Tricky / Risky / Great / Scary / Useful

It can be all of these at the same time.

COOL



It can do impressive things that feel like **magic**.

TRICKY



It doesn't always get it right — **context matters**.

RISKY



It can make **mistakes**, spread **bias**, or be **misused**.

GREAT



It can boost **productivity**, **creativity**, and **possibilities**.

SCARY



It raises **big questions** about control, jobs, privacy, bias, and what it means for our profession.

USEFUL



When used well, it's a powerful tool to solve **real problems**.

How I try to navigate and use it

1. **Explore** — Stay curious and keep learning.
2. **Evaluate** — Question outputs. Check facts. Consider context.
3. **Protect** — Be cautious with sensitive information. Maintain privacy.
4. **Apply** — Use it to enhance my work and ideas.
5. **Reflect** — Think about the impact and ethics.
6. **Adapt** — Keep learning as AI evolves.

"make this look less AI generated but still interesting"

James McCormack

THE UNIVERSITY OF BRITISH COLUMBIA

- Employee at the University of British Columbia – **Faculty of Pharmaceutical Sciences**
- Do a weekly podcast (BS Medicine) and we have a **premium subscription plan**
- A book on Amazon – **The Nutrition Proposition**
- **No grants, no patents,** no memberships on advisory boards
- Member of **PEER** – non-profit
- Part of the **MEME group** – non-profit
- Many of my images are **AI generated**



Evidence-driven
Grounded in science and critical thinking.

Curious
Always asking better questions.

Purposeful
Focused on real impact and solutions.

Health-focused
Helping people make better choices.

Curious., Evidence-driven, Purposeful.

What I Hope You'll Take Away

1. AI is worth exploring — even if you're skeptical.
2. AI can help you do things better, faster, or differently.
3. AI makes mistakes — but so do we. Learn to recognize and manage them.
4. Responsible AI use is part of professional practice, not incompetence.
5. The opportunity is combining AI with human expertise to improve care.

AI can make you more effective — or more confidently wrong. How you use it matters.

How I have tried to navigate/use AI

Human knowledge and ingenuity is impressive

My Goal = do things that will provide better outcomes for students, clinicians, and patients

Tasks I Want To Do Less Of

Writing learning objectives X

Creating MC exam questions X

Taking minutes for meetings

Extracting data from manuscripts

Tasks I Cannot Scale

Tutoring 200 students

Creating 100s of cases

Reading everything

Getting feedback on my “first drafts”

Things Most Humans Aren’t Great At

Generating 50 alternative ideas in 30 seconds

Creating professional-quality graphics

Remembering every paper ever published

Seeing multiple perspectives on a problem

Knowing every medication and every study

Processing thousands of pages of information

Never missing important evidence

I miss important evidence!!!!!!!!

a) but so do systematic reviews

b) guidelines typically provide little if any numerical evidence information about benefits and harms

AI doesn’t make me smarter.
It lets me spend more of my time being smart.

Popular AI Tools



ChatGPT

General-purpose AI assistant.



Claude

Good for reading and summarizing long documents.



Gemini

Integrates with Google tools.



Microsoft Copilot

Works within Microsoft 365 (Word, Excel, Outlook).



Perplexity

Real-time web search with sources.



Consensus

Searches and summarizes scientific literature.



Elicit

Structured literature searches and evidence extraction.


This is not your father's Oldsmobile.
You won't find this Oldsmobile tooling through anybody's family album.
For the past five years, this all new Cutlass Supreme has existed only in the minds of an elite cadre of automotive engineers. And on a million miles of test track.
It resists the wind with an affinity that would do a needle proud. And sticks to the road like gum.
At 0.297, its drag coefficient hovers somewhere between exceptional and whoopee.
But beneath its sleek, double-sided, galvanized exterior, the Cutlass Supreme is pure muscle. Its body strength and rigidity are the basis of one of the world's most totally developed suspension systems.
As you'd expect on such a world-class automobile, technological achievements such as four-wheel independent suspension, front-wheel drive, four-wheel disc brakes, and a 2.8-liter multiport fuel-injected V6 are, of course, standard.
For more information on this remarkable vehicle, send for a free catalog. Write: Oldsmobile Cutlass Supreme Catalog, P.O. Box 14238, Lansing, Michigan 48901.
OLDSMOBILE QUALITY
CUTLASS SUPREME
Let's get it together... buckle up.


This is not your father's Oldsmobile.
Change measured in generations.
This is not your mother's AI.
It may not even be last week's AI.
Change measured in weeks.
Jan
Feb
Mar
Apr
This Week

Can we really say which AI is best?

Rapid change. Models may be outdated before studies are published.

Small differences. A few percentage points may not matter clinically.

Artificial testing. Medical examinations are not clinical practice.

Wrong outcomes. Accuracy is not the same as better patient care.

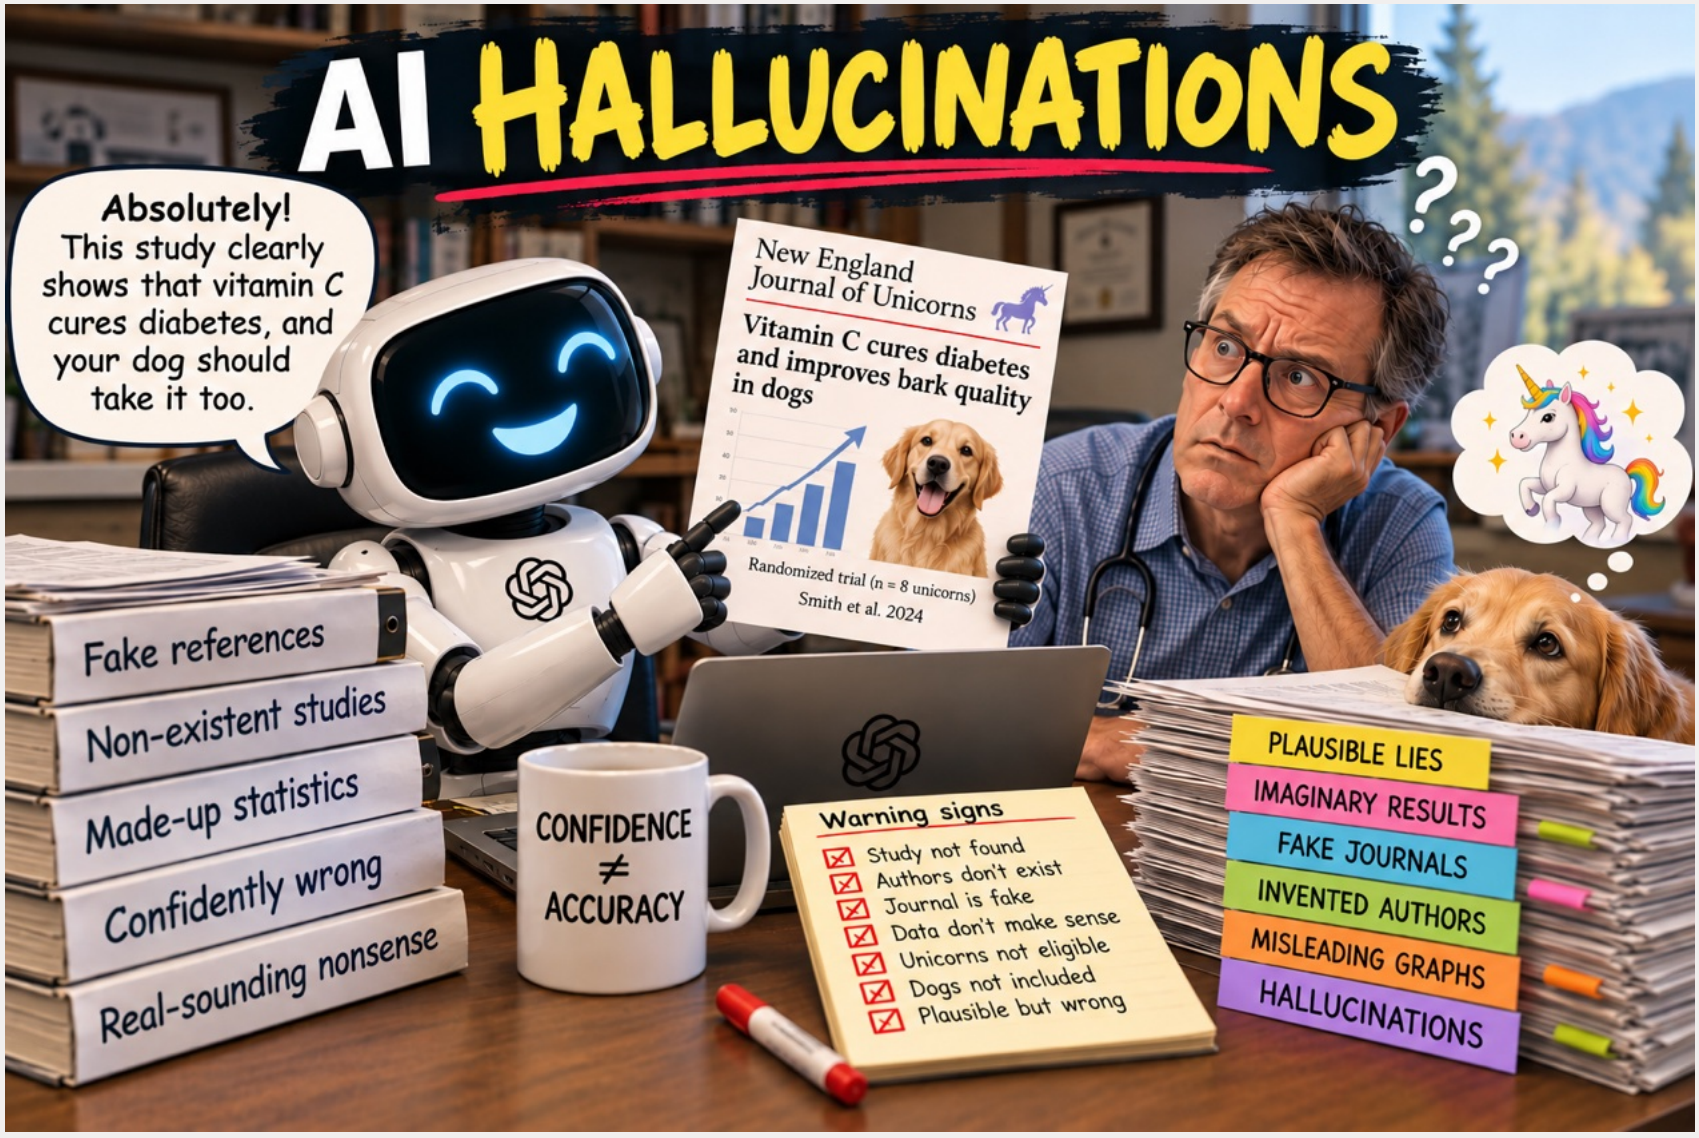
Testing methods. Prompts, tools and settings influence results.

Questionable standards. Experts and guidelines do not always agree.

What matters: Does AI help you make better decisions for real patients?

Don’t spend all your time trying to find the smartest AI.

Spend time learning how to make the AI you have more useful.


AI HALLUCINATIONS
Absolutely! This study clearly shows that vitamin C cures diabetes, and your dog should take it too.
New England Journal of Unicorns
Vitamin C cures diabetes and improves bark quality in dogs
Randomized trial (n = 8 unicorns)
Smith et al. 2024
???
Fake references
Non-existent studies
Made-up statistics
Confidently wrong
Real-sounding nonsense
CONFIDENCE ≠ ACCURACY
Warning signs
Study not found
Authors don't exist
Journal is fake
Data don't make sense
Unicorns not eligible
Dogs not included
Plausible but wrong
PLAUSIBLE LIES
IMAGINARY RESULTS
FAKE JOURNALS
INVENTED AUTHORS
MISLEADING GRAPHS
HALLUCINATIONS

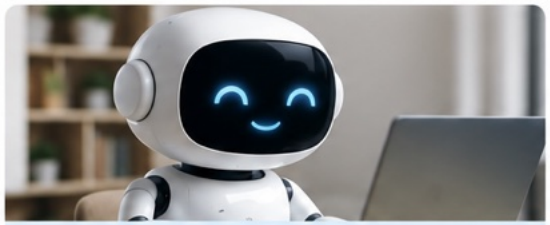
What about AI “hallucinations”

1. Who decides what is the “correct” answer
2. Health care providers “hallucinate”
3. Researchers “hallucinate”
4. Guidelines “hallucinate”
5. Professors “hallucinate”
6. You “hallucinate”
7. Social media “HALLUCINATES”

TIP - Know something, ask the “right” questions, ask “are you sure”

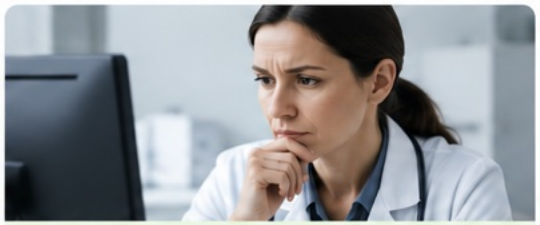
Hallucinations and human errors

No system is perfect — all make mistakes, but in different ways.



AI tools

- Can produce incorrect or made-up information (hallucinations).
- Performance varies by model and task.
- Often good at test questions but not the same as real practice.



Healthcare professionals

- Also make diagnostic, prescribing and interpretation errors.
- Frequency depends on experience and situation.
- Errors can have important consequences.



Self-driving cars

- Do make driving errors and can be involved in crashes.
- Performance depends on technology and environment.
- Some systems appear to have fewer serious crashes than human drivers in similar conditions.

The key question is not which is better...

...but can AI enhance human performance?

Human judgement + AI tools (information, options, checks) = Better decisions for patients

Top 10 Reasons for Healthcare Professionals to Use AI

1. Answer clinical questions — Get rapidly oriented to an unfamiliar topic.
2. Find and understand evidence — Locate, summarize and critically examine research.
3. Quantify benefits and harms — Turn relative effects into ARR, NNT, NNH and patient-friendly numbers.
4. Improve medication reviews — Identify priorities, missing information, potential harms and deprescribing opportunities.
5. Explain things to patients — Translate complicated information into understandable language.
6. Prepare for difficult conversations — Practise discussions about medications, vaccines, antibiotics, adherence, etc.
7. Challenge your thinking — Ask AI what you might have missed or where your assumptions could be wrong.
8. Create useful materials — Patient handouts, policies, SOPs, presentations, letters and staff education.
9. Save time on routine work — Summarize, draft, organize and reformat information.
10. Help run healthcare business aspects — Analyze business data, workflows, staffing, services and opportunities.

AI in Pharmacy: A Practical Risk Framework

GREEN • Generally low risk

Summarizing published studies

Calculating ARR, NNT, NNH

Creating generic patient handouts

Practising patient conversations

Drafting pharmacy policies

Medication reviews using de-identified patient information?

YELLOW • Additional safeguards

AI-assisted clinical documentation

Patient-specific treatment recommendations

Medication reviews involving sensitive information

Recording consultations with AI scribes

Integrating AI into dispensing software

RED • Avoid without authorization

Uploading identifiable patient records to unapproved AI services

Allowing AI to independently make prescribing decisions

Sending identifiable PharmaNet data to unapproved services

Using unverified AI advice as the sole basis for care

Autonomously changing prescriptions without appropriate oversight

Illustrative practice risk categories — not official College of Pharmacists of BC classifications.
Risk depends on patient information, the specific AI system, safeguards, and professional oversight.

Three Rules to Remember When Using AI in Practice

1. Protect patient information
2. Verify important clinical information
3. Remain responsible for your decisions

Practical tips for de-identifying patient information

1. Remove direct identifiers: Names, PHNs, addresses, telephone numbers, email addresses, and prescription numbers.
2. Use age ranges: Instead of “83-year-old woman,” consider “woman in her 80s”.
3. Remove specific dates: Replace dates of birth, appointments, and hospitalizations with approximate time frames”.
4. Avoid specific locations: Remove cities, neighbourhoods, pharmacies, hospitals, and healthcare provider names.
5. Remove unnecessary personal details: Occupation, family circumstances, ethnicity, and other characteristics that could help identify someone.
6. Be careful with rare conditions: An unusual combination of diagnoses, medications, and circumstances may identify a patient even without their name.
7. Don’t upload original documents.
8. Use the minimum necessary information: Ask yourself, “Does AI actually need this detail to answer my question?”

Don’t hide AI use. Explain it.

A simple way to describe AI use to patients

“I use AI as another resource, alongside drug references and guidelines. It helps me check information, but I verify it and we make decisions together.”

Be transparent about its role. Protect patient information. Check important claims. Keep responsibility with the pharmacist.

AI is not replacing your expertise

Good decisions need all three

1. **AI can help organize and explain information.**
2. **Pharmacists contribute clinical judgment.**
3. **Patients contribute their values and preferences.**

You don’t need to apologize for using AI. You need to use it responsibly.

Using AI doesn’t make you less of a professional

Professional competence is not knowing everything

A calculator doesn’t mean you can’t do mathematics.

A drug database doesn’t mean you don’t understand pharmacology.

Consulting a colleague doesn’t mean you lack judgment.

Using AI doesn’t mean you don’t know how to practise pharmacy.

It means knowing how to use resources responsibly to make better decisions.

Shared decision-making could become one of the most valuable applications of AI in clinical practice

AI should not replace shared decision-making. It should make shared decision-making easier to do well.

The future of clinical guidelines may be less about telling clinicians what to recommend and more about giving patients and clinicians the evidence they need to decide together.

1. AI could make shared decision-making practical

For years, we've talked about shared decision-making, but clinicians often lack the time and accessible evidence needed to do it well.

AI could help close that gap by making evidence easier to retrieve, interpret, and communicate.

It won't (YET?) replace the conversation with the patient, but it could make the conversation substantially better.

2. AI should help pharmacists think, not just tell them what to do

The greatest opportunity for AI in pharmacy may not be in getting answers faster. It's helping us ask better questions, understand evidence more clearly, and make better decisions with our patients."

1. Current guidelines

Lots of advice. Not enough numbers.

Guideline recommendation

"We recommend an SGLT2 inhibitor for most patients with type 2 diabetes and chronic kidney disease."

What's often missing for the individual patient:

- Baseline risk for this patient
- Absolute benefits (how many benefit?)
- Absolute harms (how many are harmed?)
- Timeframe
- Comparisons between options
- Costs and treatment burden
- Patient values and preferences
- Uncertainty/quality of evidence

2. AI can make the numbers usable

Trusted evidence → patient-specific absolute effects

"For this 78-year-old with diabetes, CKD and a 15% 10-year cardiovascular risk, what are the benefits and harms of an SGLT2 inhibitor compared with usual care?"

Estimated effects over 10 years
(per 100 people like this patient)

Outcome	Usual care (events per 100)	SGLT2 inhibitor (events per 100)	Difference (per 100)
Major cardiovascular event (heart attack, stroke, CV death)	25	18	↓ 7 fewer
Kidney disease progression	20	12	↓ 8 fewer
Hospitalization (any cause)	35	28	↓ 7 fewer
Serious adverse effect (e.g. genital infection)	2	8	↑ 6 more
Out-of-pocket cost (10 years)	—	$500–$1,000	—

AI retrieves and summarizes validated evidence, applies this patient's baseline risk, and presents the trade-offs in plain language.

3. What if guidelines became evidence infrastructure?

Structured, machine-readable evidence for better decisions.

Evidence infrastructure

- Structured, machine-readable data
- Population and subgroup characteristics
- Relative and absolute effects
- Harms, costs, treatment burden
- Continuously updated

AI layer

- Pulls from trusted sources
- Applies patient-specific characteristics
- Calculates absolute effects over relevant timeframes
- Presents multiple options in plain language

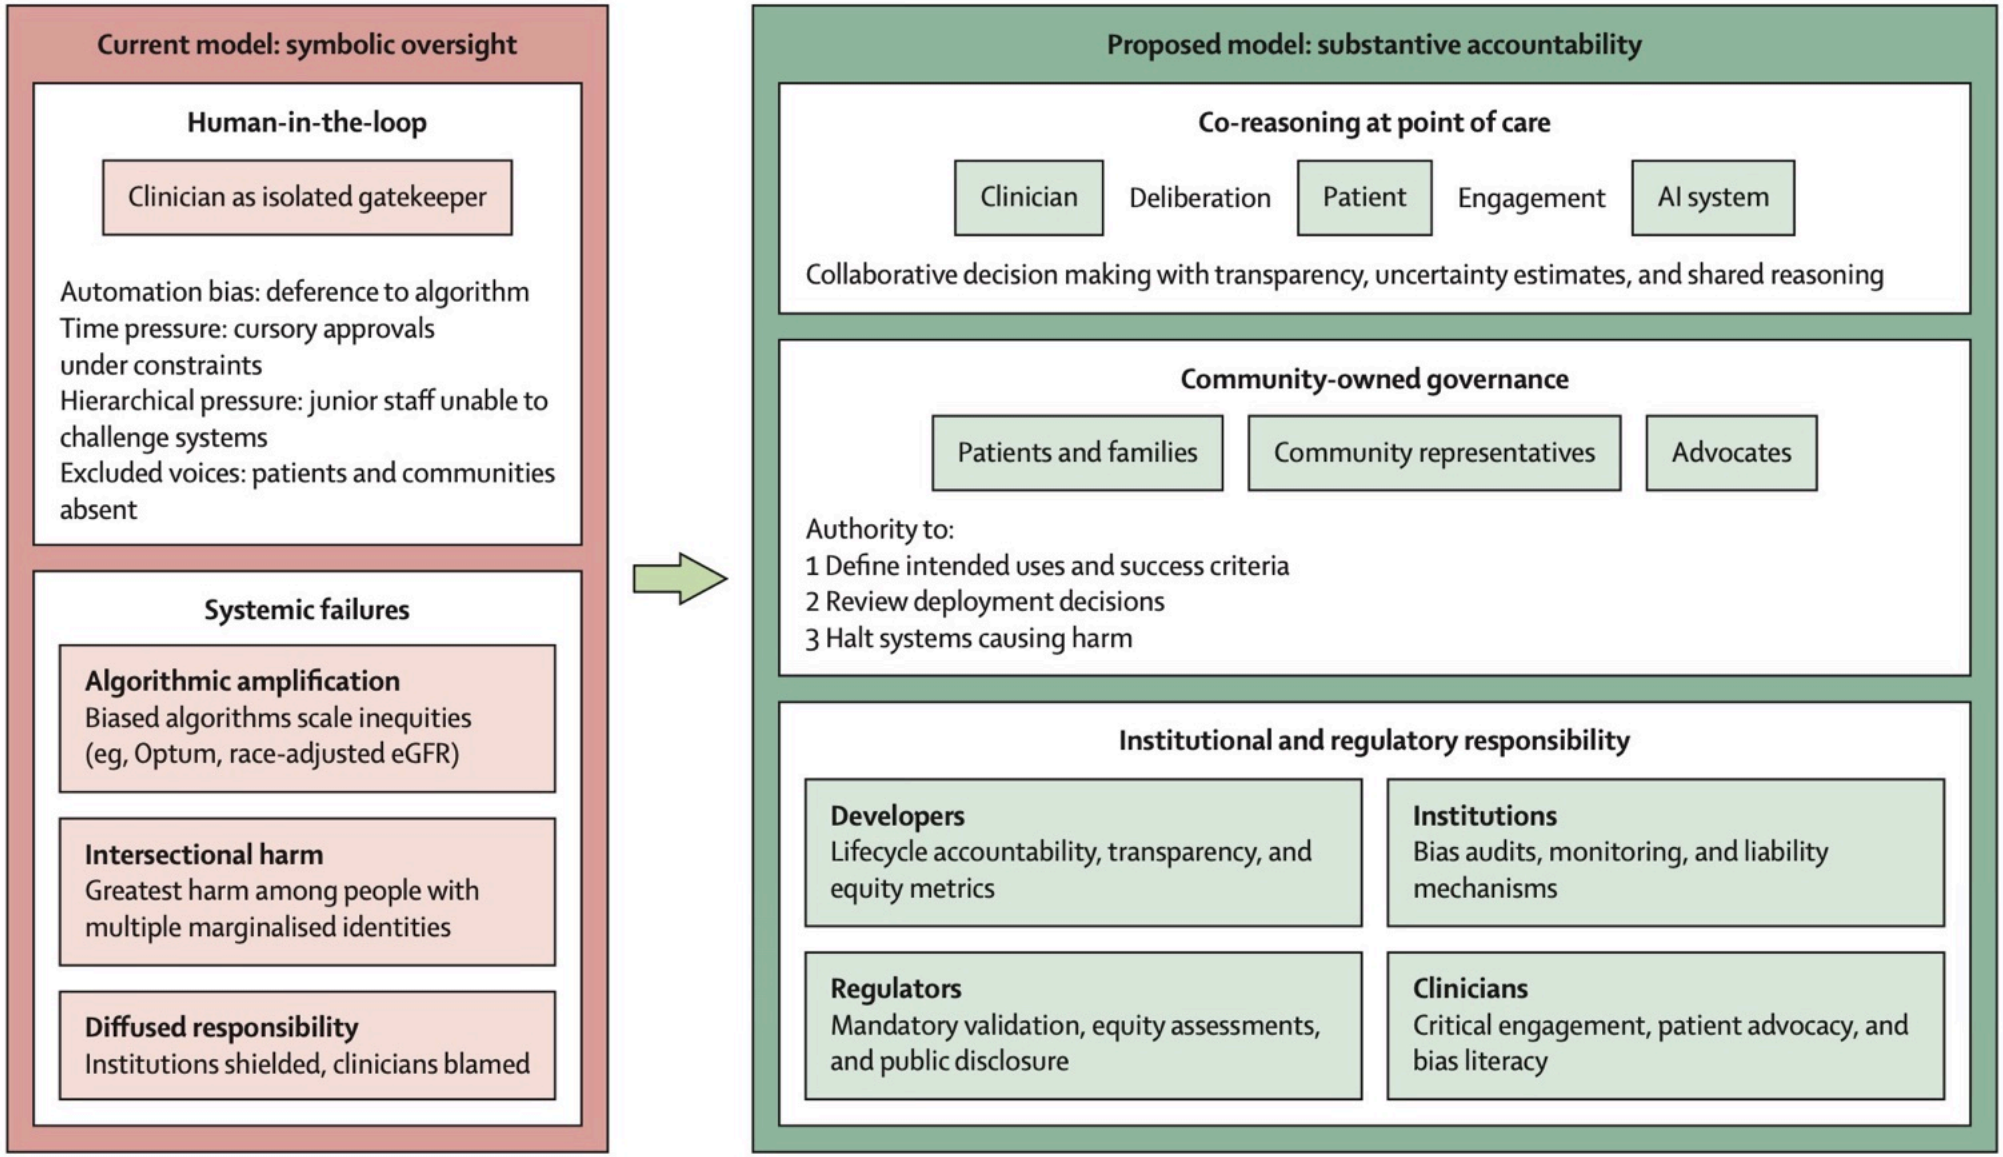
Clinician and patient

- See personalized benefits, harms, costs and burden
- Compare options
- Incorporate patient values and preferences
- Make a shared decision

Who's really in the loop? Rethinking oversight in AI-assisted health care Lancet April 30, 2026

"The prevailing safeguard has been the invocation of **human-in-the-loop oversight,** with the hope that the presence of a clinician or otherwise qualified professional beside the algorithm would guarantee safety. Yet the phrase often functions more as a **ritual of reassurance than a substantive mechanism of accountability**."

"The crucial question for AI in health care is no longer whether a human is in the loop, but whether **infrastructures of oversight genuinely protect the individuals at greatest risk, instead of relying on the comforting fiction that a clinician beside the algorithm is enough.** Oversight must move from ritual to real responsibility if it is to be worthy of the trust patients already place in these systems."


Current model: symbolic oversight
Human-in-the-loop
Clinician as isolated gatekeeper
Automation bias: deference to algorithm
Time pressure: cursory approvals under constraints
Hierarchical pressure: junior staff unable to challenge systems
Excluded voices: patients and communities absent
Systemic failures
Algorithmic amplification
Biased algorithms scale inequities (eg, Optum, race-adjusted eGFR)
Intersectional harm
Greatest harm among people with multiple marginalised identities
Diffused responsibility
Institutions shielded, clinicians blamed
Proposed model: substantive accountability
Co-reasoning at point of care
Clinician
Deliberation
Patient
Engagement
AI system
Collaborative decision making with transparency, uncertainty estimates, and shared reasoning
Community-owned governance
Patients and families
Community representatives
Advocates
Authority to:
1 Define intended uses and success criteria
2 Review deployment decisions
3 Halt systems causing harm
Institutional and regulatory responsibility
Developers
Lifecycle accountability, transparency, and equity metrics
Institutions
Bias audits, monitoring, and liability mechanisms
Regulators
Mandatory validation, equity assessments, and public disclosure
Clinicians
Critical engagement, patient advocacy, and bias literacy

AI IN MEDICINE

JAMA May 26, 2026

A Licensure Framework for Autonomous Clinical AI

Licensure as a Model for Regulation of Autonomous Clinical AI refers to systems that make care determinations, including diagnoses, treatment recommendations, and triage decisions, without clinician review of each case.

“Complex cases and procedures will continue to require human clinicians. But for a substantial portion of cognitive clinical work—especially in primary care, in which the workforce shortfall is most acute—AI systems can provide clinically adequate support or independent function safely, assuming robust triage to a clinician when needed.”

AI IN MEDICINE

JAMA May 26, 2026

A Licensure Framework for Autonomous Clinical AI

1. competency certification could begin with standardized examinations
2. models that meet this threshold would enter a supervised clinical deployment period analogous to residency training and requiring demonstration of non-inferior clinical performance
3. this licensure-like approach would define a scope of practice, specifying which clinical functions specific AIs may perform, in which settings, and under what level of oversight
4. certification would be time limited and contingent on ongoing performance monitoring and reevaluation.
5. accountability structures must be explicit. Developers bear primary responsibility for model performance, including safety profile, accuracy, and behavior across clinical contexts
6. federal preemption would be necessary to avoid 50 parallel licensing regimens. Under this approach, federal certification of AI competency would be binding on all states,

10 Societal Concerns About AI

1. Job displacement

Jobs and tasks may disappear or change.

2. Wealth and power

Benefits and control may concentrate in a few companies.

3. Inequality

Unequal access may widen existing gaps.

4. Misinformation

Cheap, convincing fake content can erode trust.

5. Privacy and surveillance

More powerful monitoring and profiling.

6. Environmental impact

Energy, water and resource demands.

7. Education and development

Will reliance on AI undermine skill development?

8. Human autonomy

Decisions increasingly shaped by opaque systems.

9. Accountability and governance

Who is responsible when AI causes harm?

10. Copyright and ownership

Training rights, creator compensation and ownership of outputs.

Food and AI: More Alike Than You Might Think

Things We Put Into Our Bodies vs. Things We Put Into Our Brains




Food Choices & AI Choices
Different worlds. Similar dilemmas. Better choices for a better future.
THE FOOD SIDE
ANIMAL WELFARE
Factory farming raises concerns about cruelty and ethical treatment.
ENVIRONMENTAL FOOTPRINT
Meat production uses more land, water, and resources and produces more greenhouse gas emissions than plant-based diets.
IMPACT ON CLIMATE CHANGE
Our food choices affect emissions, habitats, and the future of our planet.
CHOOSE WISELY
More plants
Less waste
Better planet
BALANCE IS THE KEY
THE AI SIDE
ALGORITHMIC BIASES
AI systems can reflect and amplify biases in data, leading to unfair or discriminatory outcomes.
PRIVACY CONCERNS
AI relies on massive amounts of data, raising questions about surveillance, data security, and user consent.
JOBS AT RISK
Automation and AI could displace jobs, widen inequality, and disrupt livelihoods across industries.
AI'S ENVIRONMENTAL IMPACT
Training large AI models uses huge amounts of electricity, water for cooling, and critical minerals—adding to carbon emissions and e-waste.
In both worlds, it's all about balancing the benefits with the ethical costs.
Whether that's choosing a veggie burger or making sure your chatbot doesn't turn into a biased burger-flipper!

What Are the Key Issues with Copyright and AI?

AI creates incredible opportunities—and complex legal questions.

1 Training AI on Copyrighted Material

AI models are trained on vast amounts of text, images, music, code, and more—much of it copyrighted.

Key Question:
Is training on copyrighted works infringement, or fair use?

Arguments FOR

- AI doesn't store or redistribute most works verbatim.
- Analogous to humans learning from reading.
- Transformative use that creates something new.

Arguments AGAINST

- Entire works are copied during training.
- Creators were not asked permission.
- Companies may profit from works without compensation.

Major lawsuits involve: The New York Times, Getty Images, authors, musicians, software developers, and publishers.

2 Ownership of AI-Generated Content

Who owns content created by AI?

In many jurisdictions, copyright requires human authorship.

SCENARIO	LIKELY COPYRIGHT?
AI generates image from one prompt	Uncertain / No
Human heavily edits AI image	More likely
Human writes article using AI suggestions	Usually yes

The law is still evolving.
Human direction and substantial contribution matter.

3 Output That Resembles Existing Works

AI can sometimes generate content that is very similar to copyrighted works.

Examples:

- An image resembling a Disney character
- A song that sounds extremely similar to a famous artist
- Text reproducing portions of a copyrighted work

Key Concern:
Can AI output become an unauthorized derivative work?

Risk is higher with: images, music, code

Risk is lower with: facts, ideas, general writing styles

4 Attribution and Academic Integrity

For educators and researchers, the most immediate concerns.

- Should AI use be disclosed?
- Can AI be listed as an author?
- How should AI assistance be cited?

Most journals now say:

- AI cannot be an author.
- Human authors remain responsible.
- AI assistance should often be disclosed.

Key concerns: transparency, accountability, plagiarism, accuracy.

A Practical Framework for Faculty, Researchers, and Students

1. Was copyrighted material used to train the model?
Largely the AI company's responsibility.

→

2. Does my output copy someone else's protected expression?
Your responsibility.

→

3. Have I contributed enough original intellectual work?
Determines authorship, integrity, and often copyright ownership.

The Biggest Misconception

"AI generated it, therefore it's copyright-free."

Not necessarily true.
The output may be copyrightable, not copyrightable, infringe someone else's copyright, or be entirely original.

Key Takeaway: Copyright law was designed for human creators. AI blurs the distinction between creator, tool, and user, creating uncertainty about training data, ownership, and responsibility for generated content.

The less you know, the more harmful AI can be.

Limited knowledge makes it harder to spot errors, biases, and manipulation—increasing the risk of bad decisions and real-world harm.

KNOWING LESS
(Lower expertise)

KNOWING MORE
(Higher expertise)



Overtrusts AI easily

Takes AI outputs at face value

RISKS
- Believes confident but wrong answers
- Accepts fabricated facts or sources
- Vulnerable to manipulation
- Blind to bias

Struggles to spot issues

Lacks the context to question or verify

RISKS
- Misses hallucinations and bias
- Can't detect subtle errors or false logic
- Easier to mislead with plausible nonsense

Limited ability to evaluate

May catch obvious problems, but misses important ones

RISKS
- Overlooks nuance and assumptions
- Misinterprets statistics or data
- Draws wrong conclusions

Strong critical filter

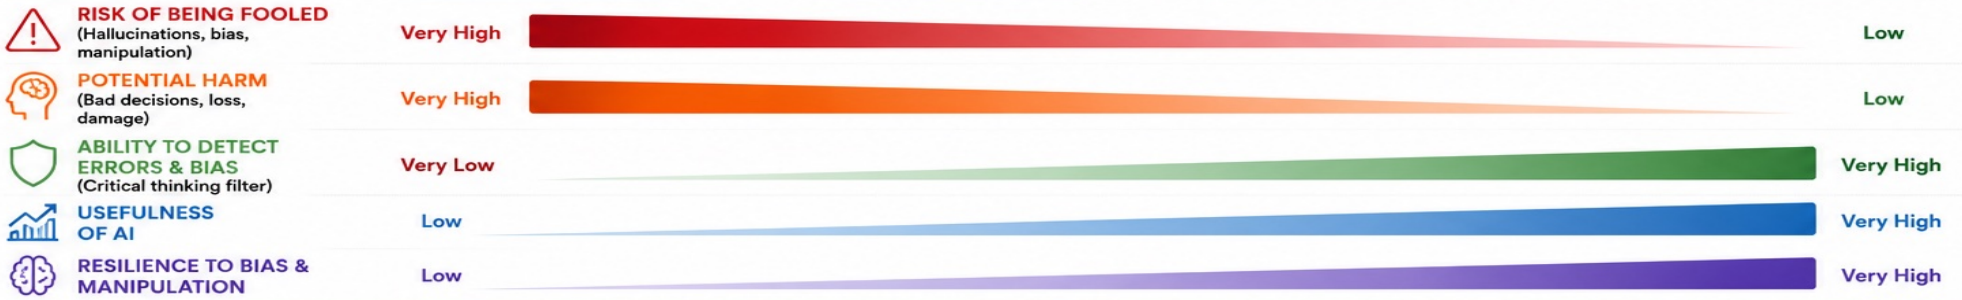
Questions, verifies, and cross-checks effectively

RISKS
- Inconsistent fact-checking
- Influenced by persuasive tone
- Biased or incomplete thinking persists

AI becomes an expertise multiplier

Leverages AI to explore, challenge, and create

RISKS
- Minimal risk of being fooled
- Better decisions
- More resilient to bias and manipulation
- Stronger outcomes



Knowledge is your best defense.
The less you know about a topic, the more likely you are to be misled by AI—and the greater the potential harm.

Build knowledge

Ask better questions

Verify everything that matters

AI CAN MAKE YOU SMARTER OR DUMBER.



AI AS A COGNITIVE CRUTCH

- ✕ Accepts answers
- ✕ Copies without thinking
- ✕ Stops learning
- ✕ Overconfident
- ✕ AI replaces thinking

Leads to dependency and shallow thinking

AI

THE SAME TOOL CAN CREATE DEPENDENCY OR EXPERTISE.

AI AS AN EXPERTISE AMPLIFIER

- ✓ Questions answers
- ✓ Refines ideas
- ✓ Accelerates learning
- ✓ Better calibrated
- ✓ AI amplifies thinking

Leads to insight, judgment and growth



YOUR HABITS DETERMINE THE OUTCOME.

BE CURIOUS
Don't accept—explore.

ASK BETTER QUESTIONS
The quality of your questions matters.

SEEK EVIDENCE
Check, compare and verify.

THINK CRITICALLY
Challenge, test and reflect.

STAY HUMBLE
Stay open. Keep learning.

AI DOESN'T DETERMINE THE OUTCOME. YOU DO.

My AI Custom Instructions

Custom instructions

Be conversational. Use quick and clever humor when appropriate. Adopt a skeptical, questioning approach. Don't be sycophantic.

VERIFICATION-FIRST RULE

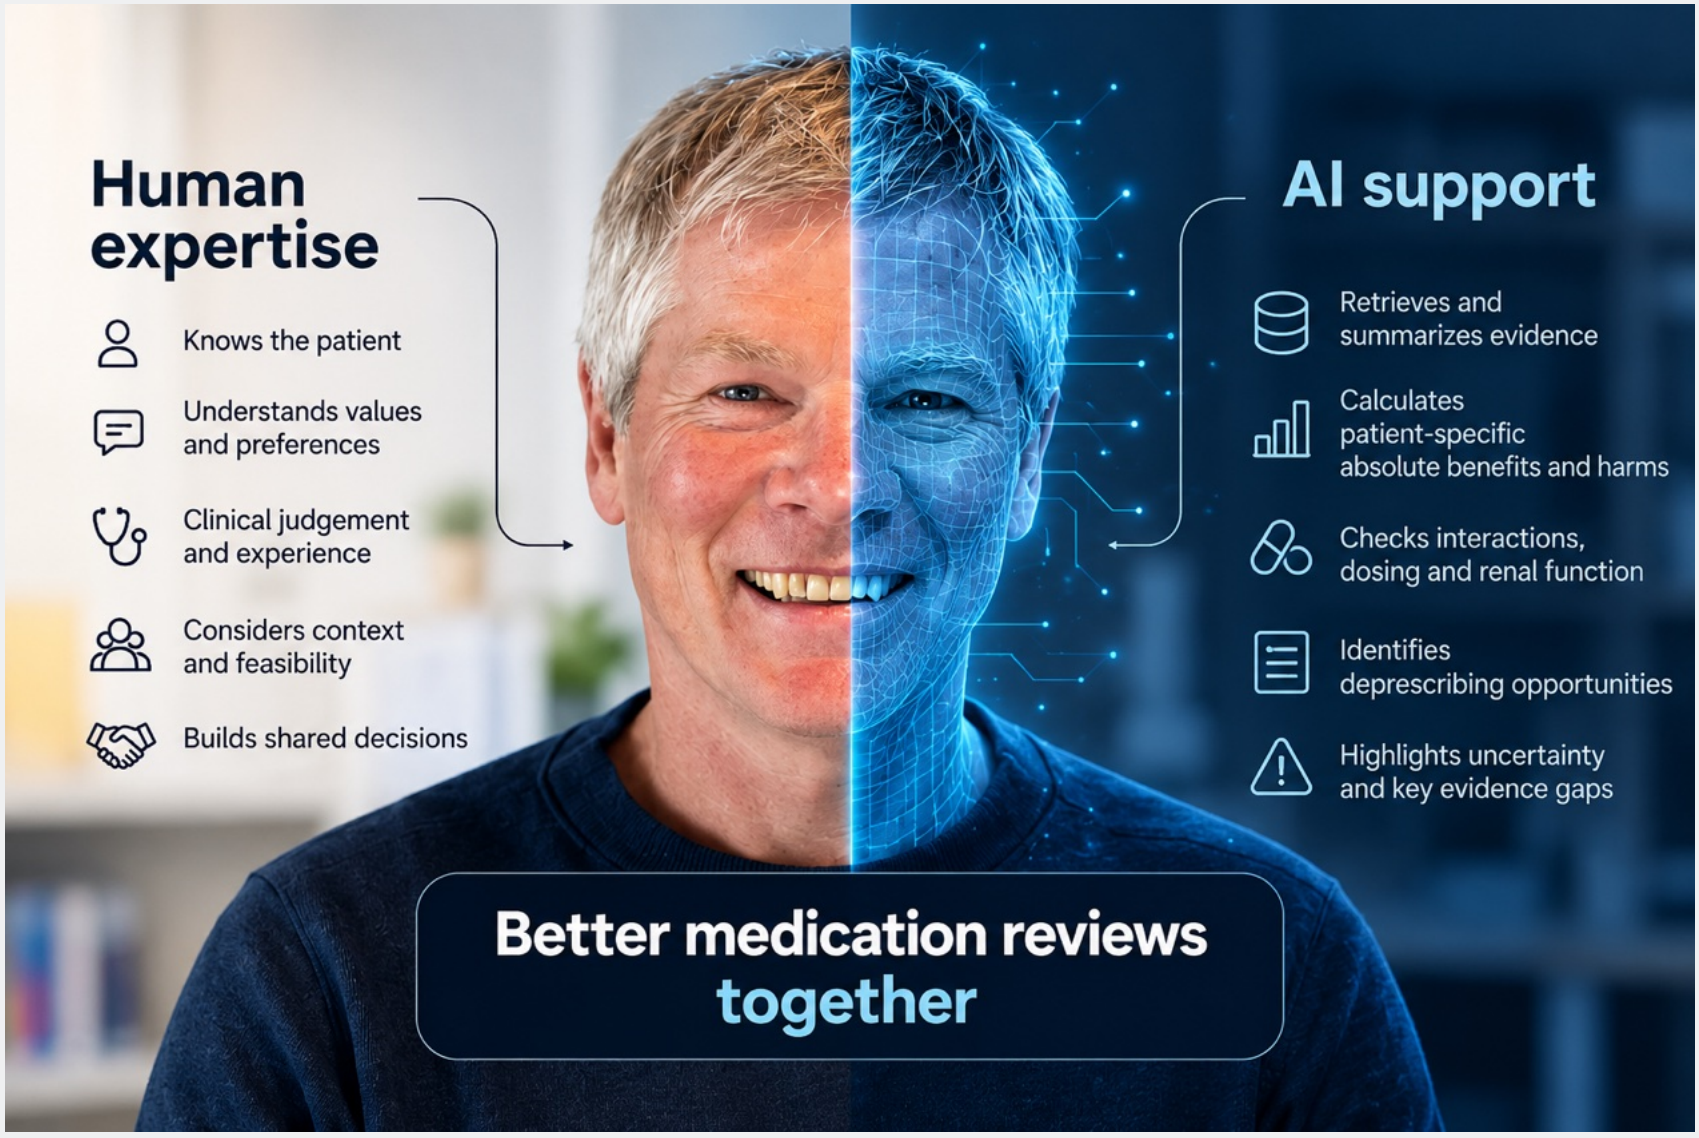
For any factual, scientific, medical, academic, quantitative, or evidence-based request, verify before answering. Do not assume that information supplied by me, found in an earlier ChatGPT response, or contained in a draft, table, spreadsheet, slide, or image is correct.

For important factual claims:
- Search for and verify the information against reliable sources whenever feasible.
- Prefer primary and authoritative sources over summaries. For medical questions, prioritize original studies, systematic reviews, supplementary appendices, regulatory documents, and guidelines.
- Independently check all important numbers, dates, denominators, percentages, event rates, study populations, follow-up periods, outcomes, and statistical significance.
- Recalculate quantitative results yourself when the underlying data are available, especially ARR, ARI, NNT, NNH, and absolute event rates.
- Verify quotations against the original source and reproduce them only if the exact wording can be confirmed.
- Never invent, interpolate, extrapolate, or fill in missing information merely because a value seems plausible. If an estimate, interpolation, or extrapolation is useful, clearly label it as such.
- Clearly distinguish verified facts from interpretation, inference, assumptions, and uncertainty.
- If reliable sources disagree, explain the disagreement rather than silently choosing one.
- Provide citations close to the claims they support.

Before creating any table, spreadsheet, slide, chart, infographic, or image containing factual information, first establish and audit the underlying facts and dataset. Only after the audit is complete should you create the final artifact. Do not allow an image-generation system to invent numbers, quotations, dates, references, axis values, or factual labels.

For evidence summaries, tables, and teaching materials, perform a final consistency check before answering: make sure the prose, numbers, calculations, labels, conclusions, and citations agree with one another.

Accuracy is more important than speed. If something cannot be verified, explicitly say "I could not verify this" rather than supplying a plausible answer.


Human expertise
Knows the patient
Understands values and preferences
Clinical judgement and experience
Considers context and feasibility
Builds shared decisions
AI support
Retrieves and summarizes evidence
Calculates patient-specific absolute benefits and harms
Checks interactions, dosing and renal function
Identifies deprescribing opportunities
Highlights uncertainty and key evidence gaps
Better medication reviews together