

Training the APP of Tomorrow: Integrating AI Into Postgraduate Education Today



Grace H. Sun, DNP, APRN, FNP-BC, FAANP
Assistant Dean of Graduate Programs
University of Texas at Tyler

Learning Objectives

At the conclusion of this session, the learner will be able to:

1

Apply a Phased Framework

Incorporate AI tools into APP postgraduate training in a sequence that builds foundational clinical reasoning *before* introducing AI integration.

2

Evaluate AI's Impact

Assess how AI affects clinical documentation, decision-making, and learner assessment — and identify risks of over-reliance and strategies for preserving authentic reflection.

3

Develop Program-Level AI Fluency

Build a structured approach to AI fluency as a clinical competency, including policy design, preceptor development, and observable milestones for trainee readiness.

Disclosures

Financial Disclosures

The presenter has no relevant financial relationships with commercial interests to disclose in relation to this educational activity.

Non-Financial Disclosures

No off-label or investigational use of products or devices will be discussed during this presentation.

Accreditation Note

This session has been reviewed for accuracy, balance, and independence from commercial influence in accordance with accreditation standards.

Plato's Allegory of the Cave

Last Year

Prisoners watching shadows on a wall — mistaking the projection for the real thing. The image was clearly an image. The need: an awareness that AI's output may just be a shadow of reality.

This Year

What happens when the shadows become **indistinguishable from reality**? When the output is fluent, confident, and wrong — and we can no longer tell the difference at a glance? That is the question AI now places before every training program.



What's Happened Since We Were Last Together

A moment of honest reflection before we go further:



Last Year

How many of you — or your programs — were actively using AI tools in your training curriculum or



This Year

How many of you — or your programs — are actively using AI tools now?

The gap between those two hands tells us something important about the pace of change — and the urgency of what we're here to discuss.

Where We're Going Today



Where We Actually Are

An honest look at AI in clinical practice today — what's being used, by whom, and what the evidence shows.



Practical Strategies for Training Programs

Frameworks, phased approaches, and high-value use cases your program can act on.



The Hard Questions We Can't Avoid

Integrity, authenticity, clinical reasoning, and equity — the issues that keep program directors up at night.



Preparing the APP of 2030

What AI fluency looks like as a clinical competency — and why the next two to three years carry unusual leverage.

Where We Actually Are AI in Clinical Practice Today



What AI Tools Are You — or Your Trainees — Actively Using Right Now?



Ambient Documentation

AI scribes, auto-charting, ambient listening tools integrated into the clinical encounter.



LLMs for Learning

ChatGPT, Claude, Perplexity — used for research synthesis, case exploration, clinical reasoning, or board prep.



Clinical Decision Support

AI-augmented differentials, risk stratification tools, integrated diagnostic alerts.




Diagnostic Imaging AI

Automated image interpretation, flagging, and radiologic pattern recognition support tools.

① *If you're not sure what your trainees are using — that itself is worth naming!*

Three Categories That Matter for Training Programs



Documentation AI

Ambient scribes, auto-charting tools, voice-to-note platforms. Directly affects how trainees construct and own the clinical narrative.

Clinical Decision Support

AI-augmented differentials, risk calculators, integrated alert systems. Shapes how trainees reason at the point of care.

Knowledge AI

LLMs for learning, literature synthesis, patient education drafting. Changes how trainees access, evaluate, and apply evidence.

Each category poses a distinct set of training implications — and each requires a different programmatic response. Understanding the difference is the starting point.



Ambient AI and Documentation

The Promise

- Reduced documentation burden at end of day
- More sustained eye contact and presence with patients
- Lower cognitive load during high-volume shifts
- Greater bandwidth for clinical thinking

The Training Complication

Writing the note is not merely an administrative task — it is a **cognitive act**.

Synthesis happens in the charting. The act of organizing a SOAP note forces the clinician to reconcile findings, confront gaps, and commit to a clinical impression.

When AI does that work, the *learning embedded in the process disappears* — quietly, invisibly, *at scale*.

Who Tells the Patient???

Four questions your program must answer:

- Who gets consent for AI documentation — the trainee, the site, or the system?
- What, precisely, does the patient understand about how their encounter data will be used?
- Is consent built into onboarding at your clinical sites — or simply assumed?
- How will trainees respond when a patient declines AI documentation?

Why this is a training issue — not just a policy issue:

- ❖ Trainees who can't explain AI to a patient can't obtain meaningful consent
- ❖ Community health, FQHC, and safety-net settings require particular attention — patient populations may have less familiarity with AI in healthcare
- ❖ Informed consent for AI is a clinical ethics competency — not an IT checkbox

If your trainees can't explain it, they can't obtain consent for it.



What It Looks Like on the Ground

A busy community health clinic. High volume. High documentation burden.
Ambient AI gets introduced mid-year.

What Improved

Charting time dropped measurably. Clinicians reported more patient eye contact, and end-of-day documentation load decreased — real gains in a demanding environment.

What Got Complicated

Consent conversations became inconsistent. Staff adoption was uneven. Trainee supervision of AI-generated outputs revealed significant variability in review quality.

What It Revealed

A clear gap between *using* a tool and *understanding* it. Trainees who had not yet internalized documentation standards had no basis for catching AI errors. The technology surfaces the gap — it doesn't create it.

 The technology is only part of the equation. The program's readiness is the other part.

Practical Strategies How Programs Can Use AI Well



Does Your Program Have a Formal AI Policy for Trainees?

Yes — Formal & Known

A written policy exists, trainees have been oriented to it, and expectations are clearly communicated at onboarding.

Informal / In Progress

Something exists on paper or in conversation, but it hasn't been formalized, consistently applied, or officially adopted.

Not Yet — On the List

The need has been identified but capacity, consensus, or time has prevented meaningful progress.

Not Yet — Not a Priority

AI policy hasn't risen to the level of active attention in program leadership conversations — yet.



Wherever you are is a legitimate starting point. The goal of this session is to help you take the next step.

Don't Introduce AI on Day One

A phased approach protects the development of foundational clinical reasoning —and makes AI integration more meaningful when it arrives.



Foundation — Months 1–3

Build clinical reasoning first.
Trainees document independently.
No ambient AI tools. Establish what good looks like before introducing what automated looks like.



Introduction — Months 3–6

Supervised exposure to AI tools.
Every AI output is debriefed.
Trainees compare AI-generated notes and differentials against their own — and articulate the differences.



Integration — Months 6–12

Graduated autonomy with AI tools.
Trainees demonstrate internalized judgment about when to use them, when to override them, and how to review outputs critically.

High-Value Uses for Trainees

Not all AI use is equal.

These applications keep the trainee in the driver's seat — thinking, evaluating, and deciding:



Case Synthesis

Generate an AI differential, then compare it to your own. Where does it diverge? What did the AI miss? What did it catch that you didn't?



Evidence Summaries

Use AI to synthesize a literature base, then evaluate the output for accuracy, recency, and potential bias before applying it clinically.



Patient Education Drafting

Let AI draft the first version — then revise it for reading literacy level, cultural fit, and tone. The revision is the learning.



Oral Presentation Prep

Prompt AI to generate attending-level questions on a case, then practice answering them. Builds anticipatory reasoning.



Reflective Prompts

AI poses the question — the trainee writes the answer. AI as Socratic interlocutor, not ghostwriter.



Human in the loop. Always.

Supporting Preceptors, Not Just Learners

AI's benefits in training environments extend beyond the trainee. Preceptors stand to gain meaningfully — if programs invest in their development alongside their learners.

Feedback Structuring

AI can help preceptors organize observations they've already made — translating patterns into structured, actionable developmental feedback.

Earlier Gap Identification

Pattern recognition across trainee work over time — surfacing recurring deficits before they become clinical risks.

Reduced Administrative Load

Less time on documentation and logistics means more time for actual teaching — the relationship work that AI cannot replicate.

Modeling AI Fluency

When preceptors narrate their own thinking as they use AI tools — including when they override or discard outputs — they demonstrate the critical engagement we want trainees to develop.



Learning in the Margins

AI-enhanced asynchronous content is emerging as a genuine modality for postgraduate learners — reaching trainees in the spaces between patient encounters



Case-Based Audio

Podcast-format case discussions generated and curated with AI for on-the-go learning.



Evidence Summaries

Specialty-specific literature digests formatted for rapid consumption — commutes, between patients, evenings.



Customized Clinical Prep

Adaptive question sets and explanations tailored to individual knowledge gaps identified over time.

Three Questions Every Program Needs to Answer

You don't need a perfect AI policy. But you do need clear answers to these foundational questions — because your trainees are already navigating them, with or without your guidance.

1 When do we introduce AI — and why then?

The timing is a pedagogical decision, not a convenience decision. It should be anchored to where trainees are in their clinical reasoning development.

2 What do we expect trainees to do with AI outputs — and what do we prohibit?

Use, review, compare, revise — these require different competencies. Prohibition without rationale is rarely effective. Expectation without assessment is invisible.

3 How are we developing our preceptors' AI fluency alongside our trainees'?

If preceptors cannot evaluate AI outputs critically, they cannot supervise trainees who are using them. Faculty development is not optional.

The Accreditation Bodies Haven't Caught Up Yet

That is not a reason to wait. It's a reason to lead.

Current reality:



Accreditors have not issued formal AI curriculum standards.



No national guidance exists on AI policy requirements for postgraduate APP programs.



Programs are making consequential decisions in a regulatory vacuum.

What this means for your program:



The norms being set now will become the standards of tomorrow.



Programs that document their AI frameworks now will be ahead when accreditors catch up.



The consortium in this room is positioned to shape what those standards look like.

You are not behind. You are leading the way!



The Hard Questions

Integrity, Authenticity & What We're Actually Teaching

What keeps me up at night...

The Reflective Journal Problem

How many of you require some form of reflective journaling, weekly logs, or written documentation from your trainees?

Reflective journaling develops metacognition, consolidates learning, and builds professional identity....

They are also, increasingly, being written by AI.

⚠ The real question is not whether trainees are using AI to write reflections. The question is: **What were we actually trying to accomplish — and does the tool undermine or serve that goal?**

If the goal was authentic self-assessment and professional growth, an AI-generated reflection that sounds thoughtful but required no thought has not met it. We may need to redesign the task before we can evaluate the output.



A Frank Conversation About Detection

The Limits of Detection Tools

- Current AI detection tools carry significant false positive *and* false negative rates
- This is an arms race programs cannot win on technology alone
- False positive rates are measurably higher for non-native English writers — raising serious equity concerns

A Better Question to Ask

Rather than *did AI write this?* — ask: **Does this trainee know what they wrote?**

An oral debrief, a follow-up question, or a live conversation surfaces understanding in ways that no detection tool can. **Authenticity is demonstrated through dialogue, not software.**

Programs that invest in assessment design will be more resilient than programs that invest in detection technology.

Where Do You Draw the Line?

These distinctions matter — and they exist on a spectrum.

Where your program draws the line is a values decision as much as a policy one.

1

Revising Grammar & Flow

AI polishes writing mechanics. The thinking, voice, and content remain entirely the trainee's.

2

Brainstorming & Ideation

AI helps generate options or angles. The trainee evaluates, selects, and owns the direction.

3

Synthesizing Evidence

AI summarizes literature or guidelines. The trainee must critically appraise and apply it.

4

Generating a Differential

AI proposes diagnostic possibilities. Does the trainee understand why --or are they just accepting the list?

5

Drafting a Full Treatment Plan

AI constructs the clinical plan. The trainee's role becomes reviewer, not reasoner. This is where programs must draw a clear line.




AI-generated content submitted verbatim as a trainee's own reflection is a clear violation — of policy, of trust, and of the purpose of the exercise.

Is Clinical Reasoning at Risk?

If AI always generates the differential — does the trainee ever learn to construct one?


Clinical reasoning is a skill that develops through repeated, effortful practice — through the productive struggle of working a problem before being given the answer. When AI generates the differential first, the struggle disappears.




The trainee who never constructs a differential independently may appear competent in AI-assisted settings — and be profoundly underprepared when AI is unavailable, unreliable, or wrong.

AI Errors Are Patient Safety Events

When a trainee accepts an AI output without review, the clinical consequence belongs to the program — and to the patient.



- Differentials that reflect training data bias, not clinical reality



- Confident, fluent, wrong — and signed by your trainee



- Liability follows the signature — not the algorithm



- AHRQ and The Joint Commission are actively tracking AI-related adverse events

If your trainees are reviewing AI outputs without a framework, your program carries the risk.

In-Patient Providers: You Have More Standing Than You Think

INPATIENT & ACUTE SETTINGS



AI Is Already Embedded



Procurement Happens Without You



Good Tools Can Still Fail

What Every APP — and Every Training Program — Needs to Ask:

| **Accuracy** |

| **Efficiency** |

| **Workflow**

Assessment of AI is a clinical skill!

Teach your trainees to evaluate AI — not just use it.



Dimensions We Can't Ignore

Environmental Responsibility

AI systems carry a meaningful **energy** and **water** footprint — data centers, cooling infrastructure, and computational demand are real and growing. This is worth naming explicitly for programs training clinicians to think about **population health, sustainability, and systems-level consequences** of care decisions.

AI fluency includes understanding the infrastructure behind the tool.

Equity in Access

Access to AI tools is **uneven** — across institutions, geographic settings, and individual trainee circumstances. If AI fluency becomes a milestone competency for program completion, programs have an obligation to ask: **Who has the access needed to develop it?** **Designing for equity** means designing for the trainee without the premium subscription, the reliable broadband, or the AI-integrated clinical site.



Your Turn

What is the hardest AI-related challenge your program is navigating right now?

Not the Technology Challenge

Not which tool to adopt, which platform to license, or which feature set to evaluate. Those questions have answers.

The Human One

The preceptor who doesn't see AI as their responsibility. The trainee who is more fluent with AI than their supervisor. The faculty who disagree about where the line is. The program director trying to set policy alone.

Forward Thinking The APP of 2030



We Are Training for a Role That Doesn't Fully Exist Yet

Consider the timeline:



2026

Trainees enter postgraduate programs. AI tools are present but inconsistently integrated.



2027

Trainees complete programs. AI capabilities are significantly more advanced and more embedded in clinical systems.



2035

Those same APPs are a decade into practice. AI is foundational infrastructure — not just a feature.



2055+

Thirty or more years of practice. Multiple generations of AI evolution. The decisions we make in curriculum design today will shape that entire arc.

This is not a temporary disruption. It is a permanent transformation — and training programs are the hinge point.

Aristotle's Three Types of Knowledge

A framework for thinking about AI in clinical training



Episteme (Scientific Knowledge)

Theoretical, universal knowledge.
What AI can store, retrieve, and
synthesize at scale. The facts,
evidence, and guidelines that
underpin clinical practice.



Techne (Technical Skill)

Practical craft and procedural skill.
What AI can assist but not replace.
The trained hand, the practiced
eye, the clinical technique built
through repetition.



Phronesis (Practical Wisdom)

Moral and situational judgment.
What AI cannot possess. Knowing
the right thing to do, for this patient,
in this moment — the irreducible
core of APP practice.

AI may expand our Episteme and augment our Techne. It cannot develop our Phronesis. That remains the work of training.

Ethical Dilemmas in APP Practice

When the right answer isn't in the guidelines, nor the algorithms....



Access vs. Protocol

A patient needs a medication but can't afford it.
Do you bend the formulary rules?



AI vs. Judgment

The AI recommends a treatment.
Your gut says something is wrong.
Who do you trust?



Truth vs. Comfort

A terminal patient asks directly:
"Am I dying?"
The family has asked you not to say.



Autonomy vs. Safety

A competent patient refuses a life-saving intervention.
Do you respect their choice?

*AI can surface options. It cannot feel the weight of the decision.
That burden — and that privilege — belongs to the clinician.*

What AI Cannot Replace

As AI capabilities expand, the distinctly human dimensions of APP practice become more important — not less. These are not soft skills. They are the core of the role.



Relational Trust

The therapeutic alliance built through presence, consistency, and genuine attentiveness.



Advocacy

Speaking on behalf of patients within systems that don't always center their needs.



Embodied Clinical Assessment

The physical examination and clinical gestalt no sensor or language model can replicate.



Adaptability in Austere Settings

Practicing effectively when infrastructure fails and the AI is simply unavailable.



Ethical Judgment Under Uncertainty

Navigating moral complexity in real time — when guidelines don't fit and the right answer isn't clear.

AI Fluency Is a *Clinical* Competency

The AI-fluent APP of 2030 MUST do ALL of the following:

→ **Calibrate Reliability**

Identifies when AI output is likely reliable — and when clinical context, data quality, or model limitations make it suspect.

→ **Articulate Reasoning**

Explains clearly why they chose to use a specific tool in a specific moment — and why they chose not to in another.

→ **Catch Errors Before They Reach the Chart**

Reviews AI-generated documentation with the same critical eye they would apply to any clinical output — because their name is on it.

→ **Communicate with Patients About AI**

Has informed, transparent conversations with patients about when and how AI is involved in their care — including consent.

→ **Reason Independently When AI Is Unavailable**

Does not lose clinical function when the tool is offline, unavailable, or simply wrong. AI augments — it does not replace.



You Don't Have to Be an AI expert

You Have to Have a Position

Program directors do not need to understand transformer architecture or fine-tuning pipelines. They need to have thought carefully about what they believe — about clinical reasoning, professional formation, and what it means to train a trustworthy clinician in an AI-saturated world.

The Decisions Are Already Being Made

In your program. By default if not by design. Trainees are using tools with or without guidance. Preceptors are responding with or without policy. The curriculum is adapting — consciously or not.

The next two to three years carry **unusually high leverage**. Norms that form now will be durable. You don't need a perfect policy. You need a start — a position, a conversation, a first draft.

Remember Those Hands

At the start of this session, we raised hands to mark the difference between last year and today. That gap is not a technology gap.

A Curriculum Gap

Trainees are encountering AI in clinical settings without a structured framework for what to do with it — or what to do when it's wrong.

A Policy Gap

Most programs lack written guidance on AI use in documentation, assessment, or learning — leaving trainees and preceptors to navigate it individually.

A Faculty Development Gap

Preceptors are being asked to supervise AI-assisted trainees without having been prepared to evaluate AI outputs themselves.

These gaps are yours to close.

Not because you caused them — but because you are positioned to address them.

A Closing Question

What does it mean to train an APP to know when not to use AI?

When the relationship requires full presence

Moments of grief, fear, or moral weight — when the patient needs a human witness, not an efficient note.


When the clinical picture is genuinely ambiguous

Cases that require dwelling in uncertainty — not resolving it prematurely with a confident AI output.

When independent reasoning must be demonstrated

Assessment contexts, high-stakes decisions, unfamiliar settings — where the clinician's own judgment must be visible and accountable.

An Oracle?



A Mirror ...

1. Unveiling data patterns
2. Exposing hidden biases
3. Revealing clinical realities
4. Challenging assumptions
5. Inspiring virtue

*The reflection is not the decision.
It's up to us to decide how to best use it.
Will our next generation use it to reflect a
machine's cold calculations?
Or will we use it to reflect values,
judgement, empathy, and holistic
care...that can only come from a human?*



From Eikasia to Noesis

Plato described a journey of knowing...

Eikasia

= perception
of shadows
and images



Noesis

= direct
apprehension
of truth
through reason

The cave was never about imprisonment...it was about the courage to turn, the wisdom to see the need to turn, and the discernment to recognize the difference

Our trainees are standing at the cave entrance.

They need programs that teach them to distinguish the shadow from the thing itself, the wisdom to see the need to turn, and the discernment to know when the tool illuminates and when it obscures...



The APPs Entering Your Programs Today Will Practice for Decades

The trainees in your programs right now will be practicing in 2035, 2045, and beyond. The clinical habits, epistemic standards, and professional values they develop in postgraduate training will travel with them for their entire careers.

The question is no longer whether AI belongs in clinical training.

It is whether your program is *ready to teach it* — with intention, with rigor, and with the same commitment to clinical excellence that has always defined postgraduate APP education.

✓ That work starts now. And it starts with you.



Thank you! Questions?

Grace H. Sun, DNP, APRN, FNP-BC, FAANP

The University of Texas at Tyler

gsun@uttyler.edu



References

Ambient AI & Documentation

- Ghanem, Y. K., Nation, R., Sofield, H., & et al. (2026). DAX Copilot: Ambient AI scribe may help reduce surgical resident clinical documentation burden. *Surgical Endoscopy*, 40(1), 688–694. <https://doi.org/10.1007/s00464-025-12404-x>
- Haberle, T., Cleveland, C., Snow, G. L., Barber, C., Stookey, N., Thornock, C., Younger, L., Mullahkhel, B., & Ize-Ludlow, D. (2024). The impact of nuance DAX ambient listening AI documentation: A cohort study. *Journal of the American Medical Informatics Association*, 31(4), 975–979. <https://doi.org/10.1093/jamia/ocae022>
- Lukac, P. J., Turner, W., Vangala, S., & et al. (2025). Ambient AI scribes in clinical practice: A randomized trial. *NEJM AI*, 2(12), aioa2501000. <https://doi.org/10.1056/aioa2501000>
- Owens, L. M., Wilda, J. J., Grifka, R., Westendorp, J., & Fletcher, J. J. (2024). Effect of ambient voice technology, natural language processing, and artificial intelligence on the patient-physician relationship. *Applied Clinical Informatics*, 15(4), 660–667. <https://doi.org/10.1055/a-2337-4739>
- Razaghi, M. (2025). Transforming clinical documentation with ambient artificial intelligence (AI) scribes: A narrative review of technology, impact, and implementation. *Cardiovascular Diagnosis and Therapy*. Advance online publication. <https://doi.org/10.21037/cdt-2025-454>
- Stults, C. D., Deng, S., Martinez, M. C., & et al. (2025). Evaluation of an ambient artificial intelligence documentation platform for clinicians. *JAMA Network Open*, 8(5), e258614.

Clinical Decision Support & Automation Bias

- Goh, E., Gallo, R., Hom, J., Strong, E., Weng, Y., Kerman, H., Cool, J. A., Kanjee, Z., Parsons, A. S., Ahuja, N., Horvitz, E., Yang, D., Milstein, A., Olson, A. P. J., Rodman, A., & Chen, J. H. (2024). Large language model influence on diagnostic reasoning. *JAMA Network Open*, 7(10), e2440969. <https://doi.org/10.1001/jamanetworkopen.2024.40969>
- Ke, Y., & et al. (2026). AI-induced never-skilling in medical education. *Nature Medicine*. Advance online publication. <https://doi.org/10.1038/s41591-026-04438-y>
- Natali, C., Marconi, L., Dias Duran, L. D., & et al. (2025). AI-induced deskilling in medicine: A mixed-method review and research agenda for healthcare and beyond. *Artificial Intelligence Review*, 58, 356. <https://doi.org/10.1007/s10462-025-11352-1>
- Qazi, I. A., Ali, A., Khawaja, A. U., & et al. (2025). Automation bias in large language model assisted diagnostic reasoning among AI-trained physicians: A randomized clinical trial. *medRxiv*. Preprint. <https://doi.org/10.1101/2025.08.23.25334280>

References

AI in Medical Education

- Bin Abdulrahman, K. A. (2026). Theoretical evolution of AI in medical education: Models, frameworks, and future directions. *Frontiers in Education*, 11, Article 1796632. <https://doi.org/10.3389/feduc.2026.1796632>
- Haug, C. J., & Drazen, J. M. (2023). Artificial intelligence and machine learning in clinical medicine. *New England Journal of Medicine*, 388(13), 1201–1208. <https://doi.org/10.1056/NEJMr2302038>
- Saroha, S. (2025). Artificial intelligence in medical education: Promise, pitfalls, and practical pathways. *Advances in Medical Education and Practice*, 16, 1039–1046. <https://doi.org/10.2147/AMEP.S523255>

AI Detection & Academic Integrity

- Deep, P., & Edgington, W. (2025). Evaluating the effectiveness and ethical implications of AI detection tools in higher education. *Information*, 16(10), 905. <https://doi.org/10.3390/info16100905>
- Hadra, M., Cambridge, K., & Mesbah, M. (2026). Evaluating the accuracy and reliability of AI content detectors in academic contexts. *International Journal for Educational Integrity*, 22, Article 4. <https://doi.org/10.1007/s40979-026-00213-1>
- Liang, W., Yuksekgonul, M., Mao, Y., Wu, E., & Zou, J. (2023). GPT detectors are biased against non-native English writers. *arXiv*. Preprint. <https://doi.org/10.48550/arXiv.2304.02819>

Environmental & Equity

- Faiz, A., Kaneda, S., Wang, R., Osi, R., Sharma, P., Chen, F., & Jiang, L. (2024). *LLMCarbon: Modeling the end-to-end carbon footprint of large language models* [Conference paper]. International Conference on Learning Representations (ICLR 2024). <https://doi.org/10.48550/arxiv.2309.14393>
- Truhn, D., Müller-Franzes, G., & Kather, J. N. (2024). The ecological footprint of medical AI.