

Welcome to the frontier of AI in medicine

How this course works, and what the evidence says about how you learn

Sean Davis, MD, PhD

University of Colorado Anschutz School of Medicine

October 12, 2026




Next July

It is your first week as an intern. An AI scribe has drafted your note. A patient asks whether ChatGPT got her diagnosis right.

Nobody can teach you in advance where these tools help and where they fail. The tools change every few months.

So how do you get good at this? That is what today is about.

The field is moving faster than anyone can read



Redrawn from PubMed counts, retrieved October 11, 2026
Source: (National Library of Medicine, 2026).

About me



- I use AI **every day**, for code, writing, and learning new fields
- I am **not** selling AI. Expect me to show you its failures too

The one idea of this course

It is your job to learn where AI helps your work.

Nobody can tell you in advance. You find out by trying, often, on things that matter to you. You learn where AI fails the same way.

This course gives you **a framework, a routine, and people to try things with.** Most of your learning will be your own, done with AI.

PART 1

How you learn: what the evidence says

Last year's class heard three classic theories of adult learning here. Some of that does not hold up. Here is what does.

Part one.

A confession. Last year I opened the course with three classic theories of adult learning: Knowles’ andragogy, Mezirow’s transformative learning, and self-directed learning. I showed a table contrasting how children learn with how adults learn.

When I went back to the research for this year, much of that table did not hold up. So I am going to show you what I said, what the evidence actually supports, and what that means for how you spend these four weeks.

This is also a small example of the course itself: something I believed, checked against the evidence, and changed.

What I taught last year's class

THE OLD STORY: ANDRAGOGY, KNOWLES' THEORY OF HOW ADULTS LEARN

- Adults are **self-directed**; children depend on the teacher
- Adults are driven by **internal** motivation; children by grades
- Adults learn from **problems**; children from subjects
- So adults need a **different kind of teaching** from children

WHAT THE RESEARCH SUPPORTS

- The child versus adult split is **artificial**. Learning is a continuum across life (Taylor & Hamdy, 2013)
- After decades, studies of andragogy are **inconclusive** (Rachal, 2002)
- One essay argues medical students' internal and external motivations are **mixed together** and depend on context (Misch, 2002)
- Still useful: **experience** and **relevance** matter

Speaker notes

On the left is the core of what I taught last year. Malcolm Knowles, in the 1970s and 80s, described adult learners as self-directed, internally motivated, and problem-centred, and contrasted that with children.

On the right is what the research supports.

First, a major medical education guide, AMEE Guide 83, says the split between adult and child learning “now seems to be artificial”. Most of the principles apply to children too. It is better thought of as a continuum.

Second, Knowles himself called these assumptions, not a tested theory. A 2002 review of the research on andragogy found it inconclusive after three decades, partly because nobody agreed on what to measure.

Third, the claim that adults are driven from within. A 2002 essay on medical students argues that internal and external motivation are mixed together and change with context. You want to learn medicine and you want to match. Both are real.

What survives: your experience is a resource, and relevance matters. Those are good design rules. They just do not make adults learn in a different way from children.

(Mezirow’s transformative learning describes how people change their frame of reference. It is a description more than a tested model, so I have dropped it from the slides.)

Feeling that you learned is not the same as learning

PHYSICS STUDENTS, RANDOMISED

Same content, two kinds of class.
Students in **active** classes **learned more**, but **felt** they had learned **less** than students in polished lectures.
(Deslauriers et al., 2019)

PHYSICIANS JUDGING THEMSELVES

A systematic review compared physicians' self-ratings with observed competence. In **13 of 20** comparisons there was little, no, or an **inverse** relationship. (Davis et al., 2006)

Why it matters for AI: a chatbot gives smooth, confident answers. Reading them **feels** like understanding. That feeling does not tell you whether you learned.

This is, I think, the most important finding for this course.

On the left: a Harvard physics study. Students were randomly assigned to active classes or to lectures by experienced, highly rated lecturers, with the same content and handouts. The active-class students learned more on the test. But they rated their own learning lower. The effort felt like failure.

On the right: a JAMA systematic review of physicians. When they compared physicians’ ratings of themselves with outside measures of competence, most comparisons showed little, no, or even an inverse relationship. Self-ratings were least accurate among the least skilled physicians and among the most confident.

[Click] Why this matters for AI. A chatbot gives you a fluent, confident, well-organised answer. Reading it feels like understanding. But feeling that you learned and actually learning are different things, and you, like everyone, are not good at telling them apart from the inside. So you need ways of checking that do not depend on how it feels.

What works, rated by usefulness

HIGH USEFULNESS

- **Testing yourself:** recall it, do not reread it
- **Spacing it out:** come back over days and weeks

MODERATE USEFULNESS

- **Explaining why** it is true
- **Explaining it to yourself** step by step
- **Mixing** different kinds of problems

LOW USEFULNESS

- **Rereading**
- **Highlighting**
- **Summarising**

Rereading and highlighting are what students use most

You already know this from Anki and question banks. A medical education review argues for **frequent, spaced tests with feedback** (Larsen et al., 2008).

Ratings from a review of ten study techniques (Dunlosky et al., 2013). Testing effect: (Roediger & Karpicke, 2006). Spacing: (Cepeda et al., 2006).

Speaker notes

In 2013, five cognitive psychologists rated ten common study techniques by how useful they are, based on whether the benefits hold across ages, materials, and kinds of test.

Most of you already do this: Anki and question banks are testing and spacing. So I will be quick, and the point is the AI link coming up.

Two came out on top. Testing yourself: pulling the answer out of your memory rather than reading it again. And spacing: coming back to material over days and weeks rather than cramming.

In the middle: explaining why something is true, explaining it to yourself step by step, and mixing different kinds of problems together.

At the bottom: rereading, highlighting, summarising. Rereading and highlighting are the ones students report using most, because they feel productive.

A 2008 review in Medical Education argued for exactly this in medicine: test often, space it out, give feedback.

You will see these ideas in the course design: a weekly journal where you recall what you tried, small groups where you explain it out loud, and a routine that brings you back every week.

What keeps you going

Self-determination theory: people stay motivated when three needs are met (Ryan & Deci, 2000; Ten Cate et al., 2011).

- **Autonomy:** you have real choices
- **Competence:** you can see yourself getting better
- **Relatedness:** you learn alongside other people

In 383 medical students, motivation from **within** went with **deeper study** and, through it, **better grades** (Kusurkar et al., 2012).

HOW THE COURSE USES IT

- **Autonomy:** you choose your sessions, your topics, and what you learn with AI
- **Competence:** you try things weekly and write down what you caught
- **Relatedness:** Friday small groups, and an unofficial Discord

So if “adults are internally motivated” is too simple, what do we know about motivation?

Self-determination theory, from Deci and Ryan, is one of the best-tested theories of motivation in psychology, and there is an AMEE guide applying it to medical education. It says people stay motivated when three needs are met: autonomy, real choice; competence, seeing yourself improve; and relatedness, doing it with other people.

It also distinguishes motivation that comes from within, because you value the thing, from motivation that is controlled from outside, by grades or pressure. A Dutch study of 383 medical students found that motivation from within went with deeper study strategies and more effort, and through those, better grades.

On the right is how this course is built around it. The course is ungraded on purpose. You choose what to attend and what to learn. You try things and record what you caught. And you do it with other people on Fridays.

Self-directed, but not unguided

SELF-DIRECTED LEARNING HELPS KNOWLEDGE

Across 59 studies in health professions, self-directed learning gave a **moderate gain in knowledge**, with no clear effect on skills. It worked better when learners **chose their own resources**, and for **more advanced** learners. (Murad et al., 2010)

BEGINNERS NEED STRUCTURE

A widely cited review argues that when people are new to a field, learning with **little guidance** works **less well** than guided practice. (Kirschner et al., 2006)

Most of you are new to AI. So the course gives you a framework and a weekly routine.

Self-directed learning is the third theory I taught last year. The evidence here is more encouraging, with a catch.

A Mayo Clinic systematic review of 59 studies in health professions education found that self-directed learning gave a moderate improvement in knowledge compared with traditional teaching, though no clear effect on skills. It worked better when learners helped choose the resources, and advanced learners benefited more.

You are advanced learners in medicine. In AI, most of you are beginners. A widely cited review argues that for beginners, discovery with little guidance works less well than guided practice, because a novice does not yet know what to look for. That claim is debated, but it fits what I have seen in this course.

So that is the design: you direct your own learning, inside a framework, with a weekly routine and people to check in with.

AI can help you learn, or do the learning for you

About 1,000 high school students, randomised, practising maths (Bastani et al., 2025):

+48%

practice grades
with a ChatGPT-like
assistant

−17%

exam grades once
it was taken away,
compared with no
AI

**Largely
avoided**

with a tutor version
that gave hints, not
answers

Students used the plain version **as a crutch.**

THE SAME THING IN CLINICIANS?

Endoscopists in Poland began using AI to spot polyps. In their colonoscopies **without** AI, the adenoma detection rate fell from **28.4% to 22.4%**. (Budzyń et al., 2025)

One observational study, so not proof. But it is the risk to watch.

This is the evidence behind a line from last year’s slides: use AI in support of your learning, not to replace it.

A 2025 field experiment with nearly a thousand high school students practising maths. One group had a plain ChatGPT-style assistant; one had a tutor version designed to give hints rather than answers; one had nothing. With the plain assistant, practice grades went up 48 percent. Then, on the exam without AI, those students did 17 percent worse than students who never had it. They had used it as a crutch. The tutor version largely avoided that harm.

On the right, the same worry in clinicians. Endoscopists at four Polish centres started using AI to detect polyps. When they did colonoscopies without the AI, their adenoma detection rate fell from 28 to 22 percent. That is one observational study, and there are other explanations, but it is exactly the risk: the skill you hand to the machine is the skill you stop practising.

The lesson is not “avoid AI”. It is “use it in a way that makes you do the thinking”.

Learning with AI: six habits from the evidence

1. **Tell the AI what you want to learn**, why, and what you already know.
2. **Try first, then ask.** Write your answer, then compare it with the AI's
3. **Ask it to quiz you**, not to explain again. Testing beats rereading
4. **Explain it back** in your own words, and ask the AI what you got wrong
5. **Come back to it** later in the week. Spacing beats cramming
6. **Write down what you caught.** Your sense of “I get it” is not enough

This is the main work of the course: **Learn it with AI**, 2 to 4 hours a week, on something you care about.

Speaker notes

Putting it together, six habits, each tied to the research we just saw.

One: tell the AI what you want to learn, why, and what you already know. You are setting the goal, not the machine, and it can pitch its help at your level.

Two: try first, then ask. Do the thinking before the machine does. This is what the tutor version in the maths study enforced.

Three: ask the AI to quiz you, not to explain again. Testing yourself beats rereading, and a chatbot is a tireless quiz-maker.

Four: explain it back in your own words and ask what you got wrong.

Five: come back to it later in the week. Spacing beats cramming.

Six: write down what you caught, including the AI’s mistakes. Remember that your feeling of understanding is not a reliable gauge, so you need a record.

[Click] This is what we call “Learn it with AI”: two to four hours a week using AI to teach yourself something you care about. It is the main hands-on work of the course, and the Thursday journal is where you record what you caught.

PART 2

How the four weeks work

October 12 to November 6. Ungraded, on Zoom, and more on the calendar than anyone needs.

Week 1: live every morning

Day	Session (Mountain time)
Monday, October 12	10 a.m., recorded. Welcome: how the course works, and how you learn (today)
Tuesday, October 13	10 a.m., recorded. A short history of AI
Wednesday, October 14	10 a.m., recorded. What AI is: from AI to machine learning to large language models
Thursday, October 15	10 a.m., recorded. Hands-on: centaur or cyborg, live. Bring your ChatGPT Edu account. Office hours at 11, not recorded
Friday, October 16	Small groups meet for the first time, in the slot you signed up for. Not recorded

Weeks 2 to 4: a week in the course

When	What	Expected?
Monday, 10 a.m.	“How I used AI this week” , then a talk on the week’s topic	Your choice
Wednesday, 10 a.m.	A second talk, sometimes by a guest	Your choice
Wednesday, 2 p.m.	Watch and do , 90 minutes: try it live, together	Optional
Through the week	Guest talks from clinicians and researchers	Your choice
On your own	Learn it with AI , 2 to 4 hours	Yes
Thursday, 11 a.m.	Office hours, drop in, not recorded	Optional
Thursday night	Learning journal entry	Yes
Friday	Small group , in your slot, not recorded	Best effort

Aim for about 3 hours a week of sessions, live or recorded, chosen by interest. Plan on 8 to 10 hours a week in total.

The talks: a starting plan

WEEK 2 · GETTING GOOD AT IT

Mon Getting good answers: how a chatbot reads your prompt

Wed Professional work with AI: writing and beyond, with **Christina Olson**¹

WEEK 3 · AI AND YOUR PATIENTS

Mon Talking with patients about AI

Wed To be announced: bias and equity, or building a small tool live

WEEK 4 · RESPONSIBILITY

Mon Who answers for the note: ethics and liability

Wed Where this is going, for jobs and the planet, and what you will keep

Topics may change. **Canvas has the current plan.**

1. Christina Olson: course co-director; pediatric hospitalist and Telehealth Medical Director, Children's Colorado.
Welcome · IDPT 8079 · October 12, 2026

Guest talks and workshops

Date	Who	Topic
Wednesday, October 21, 1 p.m.	Dan Denman	AI in neuroscience
Thursday, October 22, 1 p.m.	Matt Zuckerman ¹	AI in medical education
Tuesday, October 27, 10 a.m. to noon	Christina Olson	Workshop, with breakout rooms
Tuesday, October 27, 1 p.m.	Matt DeCamp	Ethics
Thursday, October 29, 1 p.m.	Mike Rosenberg	Cardiology
Monday, November 2, 11 a.m.	Elizabeth Kudron	Pediatrics
Monday, November 2, 1 p.m.	Sorabh Singhal	Psychiatry

Attend or watch one guest talk a week. All are on Zoom and recorded. Times are Mountain time; the calendar in Canvas is the official list.

1. Matt Zuckerman: course co-director; emergency medicine.

What we ask of you

- The course is **ungraded**
- **Learning journal**, Thursday nights in Canvas: complete or incomplete, with a one-line reply in weeks 1 and 2
- **Week 4 piece**: “what you will keep”, five minutes at your Friday group on November 6
- Missing a session? Use the school’s process for **absences**. You do not need to ask me

*The most important evaluation is your own.
Did you use the gift of four weeks well?*

Four rules for every AI tool in this course

1 · USE CU TOOLS

Your university
ChatGPT Edu or
Copilot account

2 · NO REAL PATIENTS

Never put real
patient
information into
a course activity.
Make cases up

3 · CHECK IT

Check what
matters **before**
you rely on it

4 · YOU SIGN THE NOTE

Whatever AI
wrote, **you** are
responsible for it

No ChatGPT Edu account yet? Finish CU's **Student AI Training** in Canvas first. The invitation comes to your university email. Do it **before Thursday**.

Where things are

CANVAS: THE OFFICIAL RECORD

- **Announcements** and the **calendar**, with Zoom links
- Weekly **modules** and the **journal**
- **General Questions** discussion
- Recordings of the talks

DISCORD: AN UNOFFICIAL EXPERIMENT

- A place to share what you tried, ask quick questions, and post resources
- Made-up cases only, like everywhere else

Join:

discord.gg/KrEX8H6ad



How we treat each other

- We explore **different views** with respect. AI in medicine is contested, and reasonable people disagree
- Disagree with the **idea**, with care for the **person**
- On Zoom, in Discord, and online, your words are part of the course
- If something makes the course feel unsafe, contact me or the School of Medicine

An apology, and an invitation

AI in medicine changes every month, and I have my own blind spots.

If something is not in the course, do not assume it is unimportant. Tell me what you want to learn, and I will adapt where I can.

Welcome · IDPT 8079 · October 12, 2026



24

This week

1. **Get your ChatGPT Edu account** working, before Thursday
2. **Sign up for a Friday small-group slot** in Canvas
3. Come to, or watch, **Tuesday and Wednesday**
4. **Thursday**: hands-on, then office hours, then your first **journal** entry that night

What remains

Education is that which remains, if one has forgotten everything he learned in school.

Albert Einstein, 1936, quoting “a wit”. The saying was already decades old; B. F. Skinner wrote his own version in 1964.

Questions

Canvas: General Questions

Office hours: Thursday, 11 a.m.

Discord: discord.gg/KrEX8H6ad



References

- Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakçı, Ö., & Mariman, R. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences of the United States of America*, 122(26), e2422633122. <https://doi.org/10.1073/pnas.2422633122>
- Budzyń, K., Romańczyk, M., Kitala, D., Kołodziej, P., Bugajski, M., Adami, H. O., Blom, J., Buszkiewicz, M., Halvorsen, N., Hassan, C., Romańczyk, T., Holme, Ø., Jarus, K., Fielding, S., Kunar, M., Pellise, M., Pilonis, N., Kamiński, M. F., Kalager, M., ... Mori, Y. (2025). Endoscopist deskilling risk after exposure to artificial intelligence in colonoscopy: A multicentre, observational study. *The Lancet. Gastroenterology & Hepatology*, 10(10), 896–903. [https://doi.org/10.1016/s2468-1253\(25\)00133-5](https://doi.org/10.1016/s2468-1253(25)00133-5)
- Cepeda, N. J., Pashler, H., Vul, E., Wixted, J. T., & Rohrer, D. (2006). Distributed practice in verbal recall tasks: A review and quantitative synthesis. *Psychological Bulletin*, 132(3), 354–380. <https://doi.org/10.1037/0033-2909.132.3.354>
- Davis, D. A., Mazmanian, P. E., Fordis, M., Van Harrison, R., Thorpe, K. E., & Perrier, L. (2006). Accuracy of physician self-assessment compared with observed measures of competence: A systematic review. *JAMA*, 296(9), 1094–1102. <https://doi.org/10.1001/jama.296.9.1094>
- Deslauriers, L., McCarty, L. S., Miller, K., Callaghan, K., & Kestin, G. (2019). Measuring actual learning versus feeling of learning in response to being actively engaged in the classroom. *Proceedings of the National Academy of Sciences of the United States of America*, 116(39), 19251–19257. <https://doi.org/10.1073/pnas.1821936116>
- Dunlosky, J., Rawson, K. A., Marsh, E. J., Nathan, M. J., & Willingham, D. T. (2013). Improving Students’ Learning With Effective Learning Techniques: Promising Directions From Cognitive and Educational Psychology. *Psychological Science in the Public Interest : A Journal of the American Psychological Society*, 14(1), 4–58. <https://doi.org/10.1177/1529100612453266>
- Kirschner, P. A., Sweller, J., & Clark, R. E. (2006). Why Minimal Guidance During Instruction Does Not Work: An Analysis of the Failure of Constructivist, Discovery, Problem-Based, Experiential, and Inquiry-Based Teaching. *Educational Psychologist*, 41(2), 75–86. https://doi.org/10.1207/s15326985ep4102_1
- Kusurkar, R. A., Ten Cate, T. J., Vos, C. M. P., Westers, P., & Croiset, G. (2012). How motivation affects academic performance: A structural equation modelling analysis. *Advances in Health Sciences Education : Theory and Practice*, 18(1), 57–69. <https://doi.org/10.1007/s10459-012-9354-3>
- Larsen, D. P., Butler, A. C., & Roediger, H. L. (2008). Test-enhanced learning in medical education. *Medical Education*, 42(10), 959–966. <https://doi.org/10.1111/j.1365-2923.2008.03124.x>
- Misch, D. A. (2002). Andragogy and medical education: Are medical students internally motivated to learn? *Advances in Health Sciences Education : Theory and Practice*, 7(2), 153–160. <https://doi.org/10.1023/a:1015790318032>
- Murad, M. H., Coto-Yglesias, F., Varkey, P., Prokop, L. J., & Murad, A. L. (2010). The effectiveness of self-directed learning in health professions education: A systematic review. *Medical Education*, 44(11), 1057–1068. <https://doi.org/10.1111/j.1365-2923.2010.03750.x>
- National Library of Medicine. (2026). *PubMed records per year for “(artificial intelligence) OR (machine learning),” 1990 to 2025*. <https://pubmed.ncbi.nlm.nih.gov/?term=%28Artificial+Intelligence%29+OR+%28Machine+Learning%29>
- Quote Investigator. (2014, September 7). *Education is what remains after you have forgotten everything you learned in school*. <https://quoteinvestigator.com/2014/09/07/forgotten/>
- Rachal, J. R. (2002). Andragogy’s Detectives: A Critique of the Present and a Proposal for the Future. *Adult Education Quarterly*, 52(3), 210–227. <https://doi.org/10.1177/0741713602052003004>
- Roediger, H. L., & Karpicke, J. D. (2006). Test-enhanced learning: Taking memory tests improves long-term retention. *Psychological Science*, 17(3), 249–255. <https://doi.org/10.1111/j.1467-9280.2006.01693.x>
- Ryan, R. M., & Deci, E. L. (2000). Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. *The American Psychologist*, 55(1), 68–78. <https://doi.org/10.1037//0003-066x.55.1.68>
- Skinner, B. F. (1964). New methods and new aims in teaching. *New Scientist*, 22(392), 483–484. <https://bfskinner.org/wp-content/uploads/2014/02/New-Methods-aims-in-Teach.pdf>
- Taylor, D. C. M., & Hamdy, H. (2013). Adult learning theories: Implications for learning and teaching in medical education: AMEE Guide No. 83. *Medical Teacher*, 35(11), e1561–72. <https://doi.org/10.3109/0142159x.2013.828153>
- Ten Cate, T. J., Kusurkar, R. A., & Williams, G. C. (2011). How self-determination theory can assist our understanding of the teaching and learning processes in medical education. AMEE guide No. 59. *Medical Teacher*, 33(12), 961–973. <https://doi.org/10.3109/0142159x.2011.595435>