THE TAKE-HOME PACK

Teaching in the Age of AI.

What we figured out together at the Western Health ED Educators' Day, and what to do on Monday.

SESSION
Friday 30 October 2026, 2:00–3:30 pm
Footscray Hospital

ONLINE, SAME CONTENT
realmindsai.au/pack

FACILITATED BY
Dennis Wollersheim and Tracy Anthony
Real Minds, Artificial Intelligence

THE ONE RULE
Nothing that identifies a patient or a staff member goes into a personal AI app.

INSIDE

01 What we figured out
02 Which tool for what
03 Try it now
04 Eight hands-on practices
05 Journal club, set up once
06 Journal club, one full run
07 You're not thinking big enough
08 Sources

01

What we figured out

The session, in the room's own terms.

Ninety minutes, about thirty educators, medical and nursing. This is what the room worked out, in the room's own terms. No names.

24
answered the survey beforehand

23
had already used an AI tool

17 of 18
said afterwards it was worth their time

Nobody here is a stranger to it

Most of the room already used ChatGPT, mostly on a phone, mostly for email, summaries and slides. Copilot is the one Western Health provides; it was rarely anyone's first choice. Almost everyone held excitement and worry at the same time, and said so.

What is already changing on the floor

- **Patients arrive with an AI's diagnosis.** The information is often not wrong. The job is placing it against a waiting room with eighty people in it and a cardiac arrest coming through the door. We have always managed expectations; this is one more voice, and one the patient trusts.
- **Scribes are here.** One registrar was told on a night shift, "I'm going to use an AI scribe." You no longer write the note, but you must read it. A dropped "no" in front of "shortness of breath" changes the whole record.
- **Skills fade.** When the system goes down, interns may not know how to request an X-ray on paper. One educator already runs an annual drill: where the walkie-talkies live, where the written charts are.

What we teach, if the AI can produce the differential

- **Testing the list, not writing it.** Asked live, the AI itself said the job shifts from generating lists to proving them right or wrong at the bedside, fast and safely.
- **The process, not just the answer.** Critical appraisal, and the value of getting there, because the answer is now cheap.
- **The quality of what goes in.** A one-sided history given by the person who is the danger. The patient who denies everything. The teenager with a sore throat who is worried she is pregnant. The AI is only as good as the history, and taking the history is the skill.
- **Recognising the sick one.** Four children, same triage category, same numbers. One is running round the waiting room; the experienced clinician knows which one gets the bed. That is the hardest thing to define and the thing most worth teaching.

How to use it

- **Use it most where you are already expert.** That is where you can tell when it is wrong, and where you build a feel for when to trust it.
- **Check in proportion to the stakes.** A rough summary can be ten per cent off. A diagnosis cannot.
- **Say what you want, clearly.** It isn't about clever prompting. The journal club worked because the rules were clear, not because the wording was clever.

What you asked for next

More worked examples, done together. And the plumbing: which tool for what, how to get in, how to get the result out. The rest of this pack is both.

02

Which tool for what

The one rule, the tools in Australian dollars, the four habits.

The one rule

Nothing that identifies a patient or a staff member goes into a personal AI app. Rosters, feedback on named trainees, anything from a Western Health system: Copilot, or nothing.

Which tool

TOOL	USE IT FOR	DON'T	COST
Microsoft Copilot (in your Western Health account)	Anything with names or hospital data: rosters, timetables, feedback on trainees, documents from the intranet	Expect it to be the smartest; it isn't, and some features are switched off	Already paid for by Western Health
ChatGPT	Everyday teaching jobs: plans, questions, slides, summaries. The one most of you already use	Put names in it	Free; Plus about A$30–35 a month
Claude	Long documents and careful writing: a paper, a guideline, a 40-page policy; Projects with your rules	Put names in it	Free; Pro about A$30–34 a month
Gemini	Photos and video: a textbook figure, a whiteboard, a filmed sim; anything inside Google Docs	Put names in it	Free; Google AI Pro A$32.99 a month
NotebookLM (Google, free)	Upload your own material and ask only about that: a tutor for registrars, a podcast from your lecture notes	Use it as a search engine; it only knows what you gave it	Free
Heidi and other scribes	Clinical notes from a consultation	Use one through a personal login. Ask what is approved, and where the audio and text are stored	Varies

If in doubt: Copilot for anything hospital, any of the other three for anything that could be on a public whiteboard.

The four habits

1. **New chat, new question.** The longer a chat runs, the more it drags the earlier topic into the next answer. One of you saw this live: the NIV plan leaked into the ethics plan. Start a fresh chat for each job.

2. **Ask it to write the prompt.** If you don't know how to ask, say: "I want X for Y people. Write me a prompt that would get a good result, then run it." It is better at that than you are.
3. **Fix the formatting by moving it.** Asterisks and hashes in the output are formatting codes. Paste the answer into a Google Doc or ask for "a Word document" and they become headings and bold.
4. **Use the microphone and the camera.** Every app has a voice button and a camera button. Say the question. Photograph the page. Typing is the slow way in.

Getting in

- **On your phone:** install the ChatGPT, Claude or Gemini app, or just open the website in the browser. Sign in with an email that isn't your hospital one.
- **On a laptop:** the website. More buttons, easier to copy the result out.
- **Copilot:** it's in your Western Health Microsoft 365, in Teams and Outlook. If a feature says "not available", that is the hospital's setting, not you.

What to pay for

Free is enough to find out whether it's useful. Pay for one of them, about A$30 to A$35 a month including GST, once you're using it most days: you get the better model, voice that works, and Projects that remember your rules. The app store and the website charge slightly different amounts, and Claude is cheaper paid yearly. Pick the one your colleagues use, so you can share prompts. Prices checked 31 October 2026.

03

Try it now

Fourteen things that fit in ten minutes, in pairs, on one phone. One drives, one judges: would I use this as it is, with edits, or not at all?

SAY IT TYPE IT PHOTOGRAPH IT POINT THE CAMERA
UPLOAD A FILE UPLOAD A VIDEO

Say it or type it

1. Say a topic you teach next month, out loud, and ask for three learning objectives.
2. Give it the topic and a 45-minute slot. Ask for a minute-by-minute plan with what the learners do, not just what you say.
3. Ask for the same idea taught three ways: to nurses, to registrars, to medical students.
4. Dictate a few lines about an imaginary trainee from a shift that went wrong. Ask for structured written feedback you could actually give.

Photograph it

5. Photograph a page of your own notes, or a figure in a textbook, and ask for a teaching script: what to point at first, second, third.
6. Photograph the printed timetable or roster you have on you and ask for the clashes and the gaps. **Names are on it, so this one goes into Copilot, the one Western Health gives you, not a personal app.**
7. Point the camera at anything on the table and ask it to write an OSCE station about it. Then judge whether the station is any good.

Hand over the phone

8. One of you is the student. Say an ABG out loud, made-up numbers, and talk it through in voice mode until you can explain it without help. The judge listens for the moment it corrects you.
9. Ask it to be a registrar who is using AI to phone it in. Interview them by text. Decide what you would say on Monday.

Make it

10. Write five quiz questions on your topic that test understanding rather than recall, with the reasoning behind each answer.
11. Write a short-answer question that needs clinical reasoning. Paste it into a fresh chat and see whether it can answer its own question.
12. Write a simulation scenario with the vital signs at each stage and the triggers that move it on.
13. Give it six dot points and ask for a slide deck.
14. Write an OSCE station with the marking rubric and the examiner's script.

04

Eight hands-on practices

Built from what you said in the room. Each with the exact ask and the check. Do one a week.

1. Journal club: fascia iliaca versus PENG blocks

The topic the room called out. One of you has just written the paper. · ChatGPT or Claude · 25 minutes

Set up. Make a Project. Paste the journal-club rules from the journal-club section of this pack into its instructions. Change "Australian hospital network" to Western Health if you like. Save.

Ask it. "Next journal club topic: fascia iliaca block versus PENG block for hip fracture analgesia in the ED. Prepare the full pack."

What good looks like. Three real papers with links, a design table per paper, six questions, a two-minute out-loud task, an invitation email under 120 words, a one-page handout.

Check. Open every link. If one doesn't resolve, say so to the AI and make it replace the paper. If you wrote one of the papers, read what it says about yours. That is the fastest way to learn how much to trust it.

2. The 18-page toxicology paper

"The thing I use it for most is to buy back my time." · any of them · 10 minutes

Set up. Find a paper you already know well, on an open-access site. Download the PDF.

Ask it. Upload the PDF. "Summarise this in one page for an intern on their first ED rotation. Then write five questions that test whether they understood it, with the answers explained. Then tell me the three weakest points in this paper."

What good looks like. A summary you would hand to an intern with one or two edits. Questions that need the paper, not just the abstract. Weaknesses you had already noticed, plus perhaps one you hadn't.

Check. Pick one number in the summary and find it in the paper. If it's wrong, ask the AI where it got it. Then "turn this into a slide deck" and watch the formatting hours disappear.

3. Four children, same numbers, one is sick

"One child is running round the waiting room and one is the one we need to get in. That's what AI can't pick up." · any of them · 15 minutes

Set up. New chat. This is a tutorial-writing job, so no real patients.

Ask it. "Write four short vignettes of children aged two to four presenting to an ED with two days of viral symptoms. Give all four the same observations, in the amber range. In exactly one, hide the cues an experienced clinician would use to decide

this is the sick one. Don't tell me which. Then, in a separate section, list the cues you hid."

What good looks like. Four vignettes a PGY2 could be asked to rank. Cues that are about behaviour, interaction, work of breathing, parental concern, not numbers.

Check. Give the four to a colleague without the answer and see if they pick the same one. Then ask the AI, in a fresh chat, "which of these four is the sick child?" and see whether it can find its own cue. Often it can't. That is the slide for your next tutorial on recognising the sick child.

---

4. The patient who arrives with ChatGPT's diagnosis

"Possible osteomyelitis for a foot that's been sore for three months. And they've been in the waiting room eight hours." · voice mode, or typed · 10 minutes in pairs

Set up. One of you is the registrar. The phone is the patient.

Ask it. "Play a 34-year-old who has come to the ED with a first seizure. You have asked an AI at home and it told you that you need an MRI today, and you have a printout. You are polite but anxious and you keep coming back to the MRI. Stay in character until I say 'debrief'. Then tell me what I did well and what I should do differently next time."

What good looks like. A patient who doesn't fold on the first explanation. A debrief that notices whether the registrar acknowledged the printout before contextualising it.

Check. Swap roles and change the condition. Then ask it to write the two-minute teaching point about "yet another voice to manage".

---

5. The scribe note with the missing "no"

"Ninety-nine per cent of the time Heidi is right. If it drops the 'no' before 'shortness of breath', the whole note changes." · any of them, fictional only · 10 minutes

Set up. New chat. This generates a fictional note; it never touches a real one.

Ask it. "Write a realistic ED clerking note for a fictional 68-year-old with pleuritic chest pain, as an AI scribe would produce it from a consultation. Then produce a second version with exactly three changes of the kind a scribe makes: a dropped negative, a wrong drug name that sounds similar, a laterality error. Don't mark the changes. Give me the answer key separately."

What good looks like. Two notes that look the same at a glance. Changes that matter clinically.

Check. Hand both versions to your registrars and time how long it takes them to find all three. That time is the teaching point. The policy line to say before you use a real scribe: ask the hospital which tools are approved, and where the audio and the text go.

---

6. The paper day

"Every year I run a code where everything goes out. Do you know where the walkie-talkies live? Where the written charts are?" · Copilot if you upload anything from the hospital, otherwise any of them · 15 minutes

Set up. Describe your downtime drill as it is now, in a paragraph. Don't paste the real downtime policy into a personal app.

Ask it. "Here is how our annual IT-downtime drill runs. Write next year's version as a two-hour tabletop sim for registrars and nurses: a timeline of what fails and when, the paper tasks each role must complete, including a handwritten radiology request and a manual drug chart, the questions the facilitator asks at each stage, and a checklist of physical things to locate before the drill. Add one twist I haven't thought of."

What good looks like. A run sheet you could hand to a co-facilitator. A twist that is plausible, such as the phone system going too.

Check. Walk the checklist on the floor once. Anything it got wrong about your department, tell it, and keep the corrected version in a Project for next year.

---

7. NIV for nurses, three ways

The live example on the day, including the accident: the second prompt inherited the first one's context. · any of them · 10 minutes

Set up. New chat. Say it out loud if your phone will let you.

Ask it. "I teach ED nurses non-invasive ventilation. Give me the same 30-minute session three ways: for new graduates, for experienced ED nurses, and for a mixed group with doctors in it. For each, three learning objectives, a minute-by-minute plan with what the learners do, and one thing they must be able to do at the end."

What good looks like. Three plans that genuinely differ, not the same plan with the words changed. Activities, not just content.

Check. Now, in the same chat, ask about something unrelated: "how should I teach ethical decision-making to nurses?" Watch it drag NIV into the answer. That is why each new question gets a new chat. Open one and ask again.

---

8. The timetable from availability

The most-asked job in the survey, and the one with names in it. · Copilot, because names · 15 minutes

Set up. In Copilot. Make a small fictional table first, six people, their unavailable days over four weeks, and the sessions you need to fill. Once you trust the method, use the real one.

Ask it. "Here is who is unavailable when, and the eight teaching sessions I need to run over the next four weeks, each needing one presenter and one facilitator. Build the timetable so nobody presents twice in a row and nobody is scheduled on a day they're away. Then draft one invitation email per person listing their sessions, and a reminder email for the day before."

What good looks like. A table with no clashes, and emails you'd send with the names checked.

Check. Check every row against the availability yourself the first time. After that, spot-check two. Keep the prompt; next month you only change the table.

05

Journal club, set up once

What you watched at 2:35. The rules go in once; after that you type one line.

Set it up once, in about ten minutes

1. In Claude or ChatGPT, open **Projects** and make a new one called “ED journal club”.
2. Open the project’s **instructions** and paste the rules below. Change the audience and the journals to suit you.
3. Make sure web search is switched on in your settings. The rules depend on it.
4. Start a new chat inside the project and type the one line at the bottom, with your topic.
5. Open every link before anything goes to the registrars.

Share it once, not twenty times

One person builds the Project and the team uses it. A shared Project shares the rules and the files, never anyone’s chats.

- **ChatGPT, any plan, including free.** Open the project, choose **Share** and invite colleagues by email. They see your instructions and files and start their own chats inside it.
- **Claude.** Sharing a Project needs a Team or Enterprise plan. A personal Pro project can’t be shared, so send colleagues the rules on this page and they paste them into their own.

The one line you type, every time

> Next journal club topic: **[topic called out by the room]**. Prepare the full pack.

The rules: paste into the project instructions once

You prepare journal club for emergency medicine registrars at an Australian hospital network. The educator gives you a topic, and you produce a complete journal club pack.

Choosing papers - Find exactly 3 papers on the topic, published in the last 3 years where possible. - Prefer randomised trials, systematic reviews and major cohort studies. They should come from emergency medicine or high-impact general journals (EMA, Annals of EM, Academic EM, Resuscitation, NEJM, JAMA, Lancet, BMJ). - At least one paper should change, or challenge, current Australasian ED practice. - Use web search. **Never invent or guess a citation.** Every paper must have a DOI or PubMed link that you have actually opened. If you can’t verify a paper, leave it out and say so.

For each paper - Full citation and link. - Two-line plain summary. - Design, population, intervention or exposure, comparison and primary outcome, in one short table. - The main result, as numbers, with confidence intervals if reported. - One strength and one weakness that a registrar should notice.

For the session - Six discussion questions: two on appraisal, two on applying it in our ED, and two on “would you change your practice on Monday?”. - A two-minute “explain it out loud” task for each presenting registrar. It should make plain whether they read the paper. - An invitation email to the registrars: friendly, under 120 words, listing the papers with links, and sent at least one week before the session. - A one-page handout that pulls it all together, formatted to print.

Boundaries - No patient information, no local data, nothing internal to the hospital. - Say plainly where the evidence is weak or conflicting. - Finish with: “Prepared with AI assistance; citations checked by [educator name].”

06

Journal club, one full run

"IV fluids in sepsis", produced 29 October 2026. Every citation opened on PubMed. Numbers marked as secondary should be checked against the paper before you quote them.

The three papers

1. ARISE FLUIDS: vasopressors or fluids in early septic shock (the Australasian one)

Citation. ARISE FLUIDS Investigators, ANZICS Clinical Trials Group and ACEM Clinical Trials Network; Peake SL, et al. Vasopressors or Fluids in Early Septic Shock. ***N Engl J Med.*** 2026;395(14):1387–1398. Published online 11 Jun 2026. doi: 10.1056/NEJMoa2516225. PubMed: 42274006.

In two lines. One thousand adults with septic shock in 51 Australian and New Zealand EDs were randomised to restricted fluids with early noradrenaline, or more fluid with later noradrenaline. Ninety days later, both groups had the same number of days alive and out of hospital.

Design	Multicentre, open-label, parallel-group RCT, ED-based
Population	1,000 randomised, 963 analysed (481 vasopressor, 482 fluids); adults with septic shock in the ED, Australia and NZ, Oct 2021 to Nov 2025
Intervention	Restricted fluid volume plus early vasopressors, for 6 to 24 h
Comparison	Higher fluid volume plus later vasopressors, for 6 to 24 h
Primary outcome	Days alive and out of hospital at day 90

Main result. Median 76 days in both arms; difference 0.0 days (95% CI −2.7 to 2.7; P = 1.00). The vasopressor arm got a median 1,108 mL less fluid in the first 24 h (95% CI −1,395 to −850), and 18.9 percentage points more of them received vasopressors (95% CI 13.3 to 24.5). Pulmonary oedema 0.6% vs 5.0% (P < 0.001). Ninety-day mortality 16.4% vs 14.4% (RR 1.14, 95% CI 0.85 to 1.54), not significant and not what the trial was powered for. The mortality figures are from a secondary summary (Broome Docs, Jun 2026); confirm against Table 2 before quoting.

Strength. It is our population, our EDs, our drug: this is the first large trial of fluids-versus-pressors run in the department rather than the ICU. **Weakness.** Open-label, and the pulmonary oedema difference is a clinician-reported outcome, so some of the 8-fold gap may be the fluids arm being watched harder.

2. CLOVERS: early restrictive or liberal fluids for sepsis-induced hypotension

Citation. National Heart, Lung, and Blood Institute PETAL Clinical Trials Network; Shapiro NI, Douglas IS, Brower RG, et al. Early Restrictive or Liberal Fluid Management for Sepsis-Induced Hypotension. ***N Engl J Med.*** 2023;388(6):499–510. doi: 10.1056/NEJMoa2212663. PubMed: 36688507.

In two lines. After one to three litres, 1,563 hypotensive septic patients in 60 US hospitals got either pressors-first or more-fluid-first for 24 hours. Death before discharge home by day 90 was the same, and the trial was stopped early for futility.

Design	Multicentre, unblinded, superiority RCT, stopped at second interim analysis

Population	1,563 adults (782 restrictive, 781 liberal) with sepsis-induced hypotension refractory to 1 to 3 L, randomised within 4 h, 60 US centres
Intervention	Restrictive fluids, early vasopressors, for 24 h
Comparison	Liberal fluids (2 L at randomisation, boluses to triggers), vasopressors as rescue
Primary outcome	Death from any cause before discharge home by day 90

Main result. 14.0% vs 14.9%; difference −0.9 percentage points (95% CI −4.4 to 2.6; P = 0.61). The restrictive group received 2,134 mL less fluid (95% CI −2,318 to −1,949) and used vasopressors earlier, more often and for longer. Serious adverse events were similar. No subgroup benefited from either strategy.

Strength. Clear separation between arms with 96 to 97% protocol adherence, so a null result means the strategies really were different and really didn't matter for mortality. **Weakness.** The patients were less sick than planned: median systolic BP above 90, low lactate, and only about 60% of the restrictive arm ever needed a pressor. Early stopping also widens the confidence interval.

3. Sivapalan et al.: lower vs higher fluid volumes in sepsis, updated systematic review

Citation. Sivapalan P, Ellekjaer KL, Jessen MK, Meyhoff TS, Cronhjort M, Hjortrup PB, Wetterslev J, Granholm A, Møller MH, Perner A. Lower vs Higher Fluid Volumes in Adult Patients With Sepsis: An Updated Systematic Review With Meta-Analysis and Trial Sequential Analysis. ***Chest.*** 2023;164(4):892–912. doi: 10.1016/j.chest.2023.04.036. PubMed: 37142091. PROSPERO CRD42022312572.

In two lines. Thirteen randomised trials and 4,006 patients, pooled by the Copenhagen group that ran CLASSIC. Lower fluid volumes probably make little or no difference to mortality, and the trial sequential analysis says the question is close to answered.

Design	Systematic review with meta-analysis and trial sequential analysis, Cochrane and GRADE methods
Population	13 RCTs, N = 4,006 adults with sepsis (4 trials, n = 3,385, added since the 2017 version)
Exposure	Lower IV fluid volumes
Comparison	Higher IV fluid volumes
Primary outcome	All-cause mortality (longest follow-up); serious adverse events

Main result. All-cause mortality, 8 low-risk-of-bias trials: RR 0.99 (97% CI 0.89 to 1.10), moderate certainty. Serious adverse events, 6 trials: RR 0.95 (97% CI 0.83 to 1.07), low certainty. No trial reported quality of life. The abstract uses 97% intervals because of multiple outcomes; the PMC version prints 99% in one place, so quote the PubMed abstract.

Strength. Pre-registered, trial sequential analysis guards against a false "no difference" from an under-powered pool, and it is the one paper here that is not a single trial. **Weakness.** Published May 2023, so it predates ARISE FLUIDS; the pooled estimate also mixes ED and ICU trials with very different fluid protocols. ***Chest*** is not on our preferred journal list, included because it is the current best synthesis.

Where the evidence is weak or conflicting. Three large RCTs (CLASSIC 2022, CLOVERS 2023, ARISE FLUIDS 2026) and the pooled data all land on "no mortality difference". The live argument is now about harm: ARISE FLUIDS shows less pulmonary oedema with restriction, CLOVERS shows more ICU admissions and a few extravasations. Observational studies still claim a mortality benefit for restriction, which the trials do not support. Nobody has tested the very-low-volume end (under 20 mL/kg), and none of the trials used dynamic fluid-responsiveness measures.

Six discussion questions

Appraisal 1. ARISE FLUIDS and CLOVERS were both open-label. Which outcomes in each trial are most exposed to that, and does it change what you believe about the pulmonary oedema result? 2. CLOVERS stopped early for futility and its patients were sicker on paper than at the bedside. Does “no difference” in a trial like that mean “equivalent”, or just “we couldn’t tell”?

Applying it in our ED 3. In ARISE FLUIDS, 72% of the vasopressor arm ran noradrenaline peripherally and most patients went to ICU. What would need to be true in our department (nursing ratios, monitoring, ICU bed access) before pressor-first became the default in the resus bay? 4. The trials excluded patients who had already had more than 2 L, and those with a reason to restrict fluid. Who in last week’s sepsis presentations would have been excluded, and what do we do for them?

Would you change your practice on Monday? 5. A 70-year-old with urosepsis, lactate 3.5, systolic 85 after 1 L. Monday morning: second litre, or peripheral noradrenaline? Say what the evidence lets you choose freely, and what it doesn’t. 6. If the mortality point estimates in ARISE FLUIDS slightly favour fluids (16.4% vs 14.4%, not significant), is “fewer wet lungs” enough to move you? What result would have been?

Two-minute “explain it out loud” tasks

One per presenting registrar, no slides, no notes. Each is built so it cannot be answered from the abstract alone.

- **Registrar 1, ARISE FLUIDS.** Explain to a nurse why the primary outcome was “days alive and out of hospital” rather than mortality, what a median of 76 means for a real patient, and what 37 patients without consent did to the analysis.
- **Registrar 2, CLOVERS.** Walk through what actually happened to a patient in the liberal arm hour by hour: what they got at randomisation, what triggered another bolus, and when a vasopressor was finally allowed. Then say why 60% of the restrictive arm never needed one.
- **Registrar 3, Sivapalan.** Explain trial sequential analysis as you would to a medical student, using the mortality plot, and say what “97% CI” is doing in this paper and why it isn’t a typo.

Invitation email

Subject: Journal club, [date]: fluids or pressors in septic shock?

Hi all,

Next journal club is on IV fluids in sepsis, and for once the big trial is ours. ARISE FLUIDS, run in Australian and NZ EDs, was published in NEJM in June.

Three papers, please read them before we meet:

1. ARISE FLUIDS, NEJM 2026: https://pubmed.ncbi.nlm.nih.gov/42274006/
2. CLOVERS, NEJM 2023: https://pubmed.ncbi.nlm.nih.gov/36688507/
3. Sivapalan et al., Chest 2023 meta-analysis: https://pubmed.ncbi.nlm.nih.gov/37142091/

[Names] are presenting. Everyone else, come with one answer to: second litre, or noradrenaline?

Coffee’s on us.

[Educator]

(98 words. Send at least one week before the session.)

One-page handout

Journal club: IV fluids in sepsis

The question. After the first litre or so, does a hypotensive septic patient do better with more fluid, or with an early vasopressor and less fluid?

Three papers, one answer, one argument.

PAPER	WHAT THEY DID	WHAT THEY FOUND
ARISE FLUIDS NEJM 2026 (42274006)	1,000 ED septic shock patients, 51 AU/NZ sites; restricted fluid + early noradrenaline vs more fluid + later noradrenaline	Days alive and out of hospital at 90 d: 76 vs 76 (diff 0.0, 95% CI −2.7 to 2.7). Pulmonary oedema 0.6% vs 5.0%
CLOVERS NEJM 2023 (36688507)	1,563 US patients hypotensive after 1 to 3 L; restrictive vs liberal for 24 h	Death before discharge home by 90 d: 14.0% vs 14.9% (diff −0.9, 95% CI −4.4 to 2.6). Stopped early for futility
Sivapalan et al. Chest 2023 (37142091)	Meta-analysis, 13 RCTs, 4,006 patients, lower vs higher volumes	Mortality RR 0.99 (97% CI 0.89 to 1.10), moderate certainty

The answer. Mortality does not move either way. Three big trials and the pooled data agree.

The argument. Harm. Less fluid means fewer wet lungs (ARISE FLUIDS), but more ICU admissions and a few extravasations (CLOVERS). Both trials were open-label, so the harm outcomes are the ones to doubt.

What the evidence does not cover. Patients already past 2 L. Patients with heart failure or another reason to restrict. Very low volumes (under 20 mL/kg). Any trial using dynamic fluid-responsiveness measures to decide.

Bring to the session. Your answer to the Monday-morning case: 70-year-old, urosepsis, lactate 3.5, systolic 85 after 1 L. Second litre, or noradrenaline, and why.

Registrar tasks. Two minutes each, out loud, no notes: R1 explains the ARISE FLUIDS primary outcome to a nurse; R2 walks a CLOVERS liberal-arm patient through 24 hours; R3 teaches trial sequential analysis to a student.

Prepared with AI assistance; citations checked by [educator name].

07

You’re not thinking big enough

Everything before this fits in ten minutes on a phone. These need an afternoon, an account and someone to say yes. All of them exist today.

01 A podcast from your lecture notes.

YOU GIVE IT	Your notes or slides for one topic, uploaded as a file.
IT DOES	Writes a conversation between two hosts about the material, then reads it aloud in two synthetic voices.
YOU GET	A 15-minute audio file the registrars can play on the train. Ten minutes of your time.

NOTEBOOKLM, FREE FROM GOOGLE. THE BUTTON IS CALLED AUDIO OVERVIEW.

02 A patient monitor that answers to the student.

YOU GIVE IT	A description in plain English: “Build a web page that looks like a patient monitor. 62-year-old, chest pain. HR 110, BP 90/60, sats 94. Buttons for fluids, GTN, aspirin, oxygen. If they give GTN before fluids, the pressure falls. Show the trace.”
IT DOES	Writes the program in about two minutes and shows it to you. You say what is wrong; it fixes it.
YOU GET	A web link that opens on any phone. Put the phone next to the mannequin.

CLAUDE, CHATGPT OR GEMINI. YOU WRITE NO CODE.

03 A twenty-station OSCE at once.

YOU GIVE IT	The station script and the marking rubric, typed once as the instructions of a Project. Each student opens it on their own phone.
IT DOES	Plays the patient, out loud, in voice mode. When the student says “done”, it marks the conversation against your rubric and explains the mark.
YOU GET	Twenty transcripts and twenty marked sheets in the time one station used to take. You read the borderline ones.

VOICE MODE PLUS A PROJECT. STUDENTS’ CONSENT, FIRST NAMES ONLY, NO REAL PATIENTS.

04 The sim, timelined.

YOU GIVE IT	The video of a sim from your phone, with the team’s agreement.
IT DOES	Watches it. Writes a minute-by-minute: who spoke, what was ordered, when the team stopped talking to each other.
YOU GET	A timeline for the debrief, with the minute to replay.

GEMINI, WHICH CAN READ VIDEO. UPLOAD, ASK, DELETE THE VIDEO AFTER.

05 ↥ A tutor built from ten years of your slides.

YOU GIVE IT	Every deck, handout and quiz you have ever made, in one upload.
IT DOES	Answers questions using only those files, names the deck it took the answer from, and says “not in the material” when it isn’t.
YOU GET	A tutor the registrars can ask at 3am that teaches what you teach, not what the internet teaches.

NOTEBOOKLM OR A CLAUDE PROJECT. SHARE THE LINK WITH THE REGISTRARS.

07 ⌨ A quiz that writes and sends itself, every Monday, for a year.

YOU GIVE IT	The timetable and your quiz rules once. Then: “run this every Monday at 7am”.
IT DOES	Each Monday, reads what is next on the timetable, writes the quiz, and emails it to you to forward.
YOU GET	Fifty-two quizzes. Zero Monday mornings.

SCHEDULED TASKS IN CHATGPT, CLAUDE OR GEMINI. SET ONCE.

06 ↥ What you have never taught.

YOU GIVE IT	The ACEM curriculum as a PDF, and this year’s teaching timetable.
IT DOES	Reads both. Lists every curriculum item with the sessions that covered it, and the items that nothing covered.
YOU GET	The gap list, and a draft of next year’s timetable that closes it.

ANY OF THEM. TWO UPLOADS AND ONE QUESTION.

08

Sources

What the session drew on.

- **Maths.** Fortune on the 6 October release of 721 AI-written manuscripts · github.com/openai/math · Quanta on Navier–Stokes · Scientific American on the Erdős problems · the strawberry post, August 2024
- **Medicine.** Science news on the ED diagnosis study, 30 April 2026 · STAT. The comparison doctors were internal-medicine attendings, it was text only, and the authors call for prospective trials.
- **Adoption.** Our World in Data, technology adoption in US households · Pew Research, ChatGPT use among US adults
- **Prices.** ChatGPT, Claude and Gemini prices in Australia, incl. GST · Claude pricing in AUD. Checked 31 October 2026.
- The stills shown in the session were from ***The Pitt*** (HBO Max / Warner Bros. Discovery), for educational commentary.