

The Therapy Room in 2030

150-minute discussion workshop · Singapore · Hosted by Ariveria · Nicholas Ng and Neil Gupta

Use this guide at your own pace. It carries every question in the room, so the discussion can happen on paper instead of on a screen. Sources and limits for every claim are on the workshop site.

ONE PLAUSIBLE 2030

A Tuesday in 2030

- 01 Maya's watch has tracked her sleep, her heart rate and her screen time for months. It noticed her broken week before she did.
- 02 On the train she talks it through out loud. Her AI answers in a voice that remembers last month's fight with her colleague.
- 03 It drafts the difficult message, then suggests she waits a day before sending it.
- 04 It has already written a summary of her week. She chooses what her therapist sees.
- 05 By the time she sits down in your room, something has been listening, noting and advising for weeks. What arrives with her?

Seven shifts, each with a present-day proof

01 AI becomes conversational, multimodal and continuously available.

ALREADY TRUE: Real-time voice models already hold natural spoken conversation. [F1, F25]

02 It moves across phones, watches, glasses, homes and care systems.

ALREADY TRUE: Wearable and phone sensors already feed models that read health data. [F4, F5]

03 It develops longer-term memory and personal context.

ALREADY TRUE: Consumer AI products already keep persistent memory between conversations. [F3]

04 Clients use it before, between and instead of professional care.

ALREADY TRUE: Purpose-built chatbots have already moved symptom scores in randomised trials. [F9, F10]

05 Practitioners receive AI summaries, risk flags, recommendations and documentation.

ALREADY TRUE: Ambient scribes already draft clinical notes in randomised trials. [F7, F8]

06 Organisations accumulate far more about behaviour and emotional states.

ALREADY TRUE: Passive sensing studies already link phone and wearable patterns to depression and suicidal thinking. [F6, F28]

07 Professionals work alongside systems that sound confident, empathic and knowledgeable.

ALREADY TRUE: People already rate AI-written empathy highly, and agreeable AI already shifts human judgement. [F13, F15]

Already true today

INSTEAD OF CARE

Therapy by chatbot can move symptoms.

In two randomised trials, purpose-built chatbots reduced depression and anxiety scores more than a waitlist; one about as much as group therapy.

Limit: Both studied specific products, screened samples and short follow-ups. Neither was compared with individual psychotherapy. [F9, F10]

IT FEELS PERSONAL

People bond with chatbots, quickly.

Users of an early therapy bot reported a working alliance within days. People also rate AI-written empathy highly, and still prefer to receive it from a human.

Limit: A self-reported bond is not proof of equivalent therapy, and the preference studies sit outside clinical settings. [F11, F13]

IT SOUNDS CONFIDENT

AI that always agrees changes people.

In three preregistered experiments, overly agreeable AI made people more certain they were right and less willing to repair a conflict. They trusted it more, not less.

Limit: The experiments measured intentions, not relationships or clinical outcomes. [F15]

WORKING ALONGSIDE

AI can make a human response better.

Peer supporters who wrote with AI suggestions expressed more empathy, while keeping the power to accept, edit or ignore every suggestion.

Limit: The study measured the writing, not whether anyone felt better. Peer support is not therapy. [F12]

DOCUMENTATION

A third listener is already in the clinic.

Ambient scribes now draft medical notes in real trials. One saved writing time; a separate evaluation found human notes still scored higher.

Limit: Medical documentation is not psychotherapy notes, where the record itself is clinically sensitive. [F7, F8]

WHERE IT FAILS

Models still fail in mental-health settings.

Tested systems produced stigmatising and inappropriate responses to mental-health scenarios, including crisis presentations.

Limit: Constructed tests show what can go wrong, not how often it does in practice. [F16]

Old questions, new answers due

- | | |
|----|---|
| 01 | What the client brings
An AI-shaped account of the week: remembered, worded and smoothed before you hear a word of it. |
| 02 | What the practitioner already knows
A summary, a risk flag and a sleep chart wait in the file. What do you owe the data, and what do you owe the person? |
| 03 | What the AI recommends
When the wording, the prompt and the next step come from a tool, whose intervention is it? |
| 04 | The role of the therapist
Which parts of your work need you, and which parts only need doing? |
| 05 | What may be automated
Intake, check-ins, homework, documentation. Each piece handed over changes all the others. |
| 06 | What must not be
Every profession draws this line somewhere. Ours is not drawn yet. |

DISCUSSION 01 · 15 MINUTES

What stays human

If a 2030 system could do everything described today, what should still be done by a person?

- ? What in your work loses its meaning if a machine does it?
- ? What would you gladly hand over, and what does that free you to do?
- ? Where would your clients draw the line, and how would you find out?

YOUR GROUP'S OUTPUT

- | | |
|----|--------------------------------------|
| 01 | Something only a person should do |
| 02 | Something you would gladly hand over |
| 03 | The line you would defend |

2 min think alone, against the six questions · 8 min discuss in your group · 3 min agree the line you would defend · 2 min bring it back to the room

This line will not be drawn for the profession. It gets drawn in rooms like this one.

Seven eyes on 2030

Hawkins and Shohet's map has structured supervision since the 1980s: seven places to look, no fixed order. It was drawn for three people; the room now has a fourth. The same seven eyes, with the question 2030 asks through each:

01 The client and their presentation

AI enters before the session begins. The client has often been talking to a system all week, and that system has been keeping score. *Clients may arrive with an AI-shaped account of their week. What has been framed before you hear a word?*

02 The therapist's interventions

AI enters the work itself. A tool can draft the reflection, suggest the homework or word the difficult question before you do. *A tool may supply the wording, the prompt, the next step. Whose intervention is it?*

03 The client-therapist relationship

AI enters the space between you. Each of you may have consulted a system about the other: what to say, how to take it, whether the therapy is working. *The alliance turns triadic: client, therapist, and the systems each one brings into the room.*

04 The therapist's internal process

AI enters what you feel in the room. Something that sounds certain can pull your judgement toward it, and noticing that pull is now part of the work. *Countertransference now includes how you feel about a client's AI: dismissal, deference, unease.*

05 The supervisory relationship

AI enters the supervision hour. The supervisor may meet the system's version of the session first: the summary, the flagged moment, the suggested focus. *Supervision gains a question alongside "what did you do?": what did the tool do?*

06 The supervisor's own process

AI enters the supervisor's judgement too. The instincts supervisors trust, tone, hesitation, what went unsaid, were built for rooms with only people in them. *Supervisors hold their own uncertainty about systems they may never have used.*

07 The wider context

AI enters as the context itself. Platforms set the norms, employers buy the tools, insurers price the risk. The room is arranged before anyone sits down. *When a platform's defaults become the room's rules, who consented to that?*

Source: Hawkins and Shohet, *Supervision in the Helping Professions*, in print since the 1980s. Their diagram is a double matrix: two interlocking systems, client with therapist and therapist with supervisor, inside one wider context. We keep the seven eyes and ask a new question through each.

Six roles AI is growing into

ROLE 01

A private personal tool

The client journals, rehearses and reflects with it. You see only what they choose to bring.

Ask: What would you want to know about it?

ROLE 02

Part of the care team

It carries information between client, practitioner and service, with consent.

Ask: Who reads what it writes?

ROLE 03

A professional assistant

Summaries, drafts, literature and documentation, under your review.

Ask: What do you still check by hand?

ROLE 04

A monitored digital intervention

A defined element of treatment, evaluated and supervised like any other.

Ask: What evidence would you require first?

ROLE 05

An emotional companion

Available at 3 a.m. Warm, patient, and agreeable if set that way.

Ask: What does it quietly replace?

ROLE 06

Infrastructure

Booking, triage, risk flags and notes. Invisible until it fails.

Ask: Who notices when it is wrong?

Design an ideal system

Take one of the six roles. Design the version of it you would actually want in 2030.

- ? Which role did you choose, and why that one?
- ? What is the one failure that would make you withdraw it?
- ? Where does the person's consent live, and how real is it?

YOUR GROUP'S OUTPUT

01 What it does

02 What it must never do

03 What the person controls

04 When a human becomes responsible

2 min pick one role as a group · 10 min design it against the four lines · 5 min stress test:
how could this hurt someone? · 3 min bring it back to the room

The systems of 2030 are being designed now, mostly without clinicians in the room. That is still a choice.

Seven guidelines for 2030

Working principles for anyone building these systems.

- 01 Preserve agency.
The person decides. A system that narrows someone's choices while feeling helpful has failed at its main job.

- 02 Keep uncertainty visible.
Fluent wording can hide shaky ground. A tool should show what it does not know.

- 03 Make influence understandable.
If a suggestion shaped a decision, the people affected should be able to see how.

- 04 Give people control over memory and data.
What is remembered, who can see it and how it is deleted belong to the person, not the platform.

- 05 Name the responsible human.
For every output that matters, someone with a name checks it, owns it and can stop it.

- 06 Protect relationships from invisible interference.
The alliance cannot defend itself against a third party it cannot see. Make the third party discussable.

- 07 Design ways to pause, leave and seek human help.
Pausing and leaving should be easy. Every system needs a door marked human.

Three moves, eight questions

For every tool, proposal or system you meet from here on.

Move 1 · Make the system visible

- 01 What does the AI observe, remember and infer?
- 02 Does the person understand what it is doing, why it produced this output and who else may receive it?

Move 2 · Bound its authority

- 03 What may the AI recommend, decide or do?
- 04 What must it never do, even if doing so would be faster or more convenient?
- 05 Which named human is responsible when its output affects care?

Move 3 · Preserve agency and a way back to people

- 06 What can the person correct, delete, pause, refuse or take with them?
- 07 How could the system change the client–practitioner relationship, or decide whose account is heard first?
- 08 What is the clear route to human help when the system is uncertain, fails or distress increases?

An original workshop discussion aid. Not a validated assessment, clinical protocol, safety score or legal compliance checklist.

- ☐ One question from the eight I will start with
- ☐ One tool I will walk through the three moves
- ☐ One AI moment I will bring to supervision

Six prompts to adapt

Prompts to adapt using professional judgment; not a clinical assessment protocol.

“Has an AI tool been part of how you’ve been making sense of things recently?”

“What do you find yourself going to it for?”

“Which parts have been useful? Which parts have left you unsure?”

“Has anything it suggested changed what you do between sessions?”

“Would it help to look at one response together?”

“What would tell us the way you’re using it needs to change?”

START HERE: THE SHORTEST USEFUL ROUTE

Sources

[F9] **Randomized Trial of a Generative AI Chatbot for Mental Health Treatment.** Michael V. Heinz, Daniel M. Mackin, Brianna M. Trudeau, Sukanya Bhattacharya, Yinzhou Wang, Haley A. Banta, Abi D. Jewett, Abigail J. Salzhauer, Tess Z. Griffin, Nicholas C. Jacobson, 2025-03-27. <https://doi.org/10.1056/Aloa2400802>

[F10] **Efficacy of a Conversational AI Agent for Psychiatric Symptoms and Digital Therapeutic Alliance: A Randomized Clinical Trial.** Anat Shoshani, Bar Gurfinkel, Ariel Kor, et al., 2026-04-14. <https://jamanetwork.com/journals/jamanetworkopen/fullarticle/2847751>

[F11] **Evidence of Human-Level Bonds Established With a Digital Conversational Agent.** Alison Darcy, Jade Daniels, David Salinger, Paul Wicks, Athena Robinson, 2021-05-11. <https://formative.jmir.org/2021/5/e27868>

[F15] **Sycophantic AI decreases prosocial intentions and promotes dependence.** Meng Cheng, C. Lee, Pratyusha Khadpe, S. Yu, D. Han, Dan Jurafsky, 2026-03-26. <https://doi.org/10.1126/science.aec8352>

[F12] **Human-AI collaboration enables more empathic conversations in text-based peer-to-peer mental health support.** Ashish Sharma, Inna W. Lin, Adam S. Miner, David C. Atkins, Tim Althoff, 2023-01-23. <https://www.nature.com/articles/s42256-022-00593-2>

[F7] **Ambient AI Scribes in Clinical Practice: A Randomized Trial.** Paul J. Lukac, William Turner, Sitaram Vangala, Aaron T. Chin, Joshua Khalili, Ya-Chen Tina Shih, Catherine Sarkisian, Eric M. Cheng, John N. Mafi, 2025-11-26. <https://pubmed.ncbi.nlm.nih.gov/41497288/>

[F16] **Expressing stigma and inappropriate responses prevents LLMs from safely replacing mental health providers.** Jared Moore, Declan Grabb, William Agnew, Kevin Klyman, Stevie Chancellor, Desmond C. Ong, Nick Haber, 2025-06-23. <https://facctconference.org/static/docs/facct2025-206archivalpdfs/facct2025-final197-acmpaginated.pdf>

[F20] **Ethical Guidance for AI in the Professional Practice of Health Service Psychology.** American Psychological Association, 2025-12-01. <https://www.apa.org/topics/artificial-intelligence-machine-learning/ethical-guidance-ai-professional-practice>

The full library of 31 sources, with limits for each, is on the workshop site. Research checked 2026-09-21.