

## OR Scene — 4-Minute Oral Script

- Patient Jonas K
  - **Personal AI**
  - Dr Mira
  - **Care Provider AI**
  - External Ethics Auditor, Prof. Kahn
  - Narrator
- 

1. Jonas: Personal-AI, my left eye itches today.
2. Personal-AI: Upload a photo of your eye.
3. Jonas: Photo uploaded.
4. Personal-AI: That suggests an acute retinal emergency. Can I notify the Care Provider AI?
5. Jonas: Yes, do so.
6. Care-Provider AI: Alright, Jonas, an emergency transport will arrive in 4 minutes
7. *(Narrator): Jonas wakes up in the operation room*
8. Jonas K *(waking; short, punchy, hyperbolic)*:  
Why is my eye sight so blurred? Very bad. Who did this operation to me?
9. Personal AI *(amped, staccato)*:  
Huge miss, Jonas. As your personal AI I should say Worst eye operation ever. I did amazing in identifying your cataract earliest possible. We want answers.
10. Dr. Mira *(gentle, to Jonas — empathic)*:  
Jonas, I'm here with you. I'm sorry you're waking into fear. We won't speak in riddles. I'll explain plainly, and stay with you through this.
11. Jonas:  
Dr. Mira, talk clear with me. No spins.
12. Dr. Mira *(turns to team — very fact-based)*:  
Care Provider AI, summarize the event. Keep it objective.
13. Care-Provider AI *(crisp, data-driven)*:  
Event timeline: risk alert at second 14. Alert level "notify." Human acknowledgement was delayed several seconds. Contributing factors: threshold ambiguity; hesitancy to escalate halt.
14. Dr. Mira *(fact-based)*:  
So: alert came in too late; the situation wasn't easy; I needed more time to evaluate. That's on the system, including me.
15. Jonas *(flares)*:  
Names. Not fog. You. The machine. Who's on the hook?
16. Personal AI:  
My patient wants receipts. Big receipts. No vibes. We saw vibes before. Not again.
17. Dr. Mira *(to Jonas — empathic)*:  
You deserved a care system that protected you, not our comfort. I'm sorry. I will take

responsibility for making human “stops” effortless.

Oh finally our external ethics editor arrived.

18. External Ethics Editor, Prof. Kahn (*warm, steady; legal, future-oriented*):

Oh sorry I'm late! Your feelings are justified, Jonas. Our task is twofold: truth for today, care safety for tomorrow. We'll negotiate accountabilities you can verify, not just trust.

19. Care Provider AI (*fact-based*):

Root causes: alert ladder capped at “notify”; unclear owner for threshold changes; delayed human escalation.

20. Dr. Mira (*fact-based to Care Provider AI*):

Actionable to dos: raise alert authority, clarify threshold governance, enforce a hard stop when any party signals.

21. Dr. Mira (*to Jonas — empathic*):

Plain terms: your safety signal wasn't loud enough, fast enough, and we didn't make the brake easy to hit. That's what failed you.

22. Jonas (*tight*):

Finally. English. Keep going.

23. Prof. Kahn (*legal-facing, kind*):

We'll keep a Responsibility Log for every critical step: logging who agreed, who owned it, what they could do, what happened, and by when they answered for it. After review, you get a plain summary. This clears fog without rushing to blame.

24. Dr. Mira (*fact-based*):

I'll set up a tamper-evident OR Ledger—decisions, alerts, overrides. You'll get a clear, plain-language recap each time.

25. Care Provider AI:

I request a Parameter Governance Board—patient, frontline staff, and engineers—to set alert levels and escalation authority.

26. Jonas (*sharp*):

Power to stop. For real. Not pretend.

27. Dr. Mira (*fact-based to Care Provider AI*):

Dual-Key Override: if a human or the AI says “stop,” we stop. No debate on the floor. We sort the why in debrief.

28. Care Provider AI:

Agreed. Every stop gets logged—time, who pressed it, and the reason. Governance data only; no procedure details.

29. Prof. Kahn (*future-oriented*):

Add a No-Penalty Stop Rule—no one is punished for a good-faith stop; coaching replaces blame. We'll run independent audits every quarter and publish a short public safety update.

30. Jonas (*after a hard beat*):

Better. Real bones. Not slogans

31. (**Narrator**): "At last, everyone has reached an agreement and even devised concrete action plans. Now, we can all watch together to see if the tragedy of that day can become a promise for a better collaborative care system

32. Personal AI (*pushing, brisk*):

Big stakes. When danger shows up, we need alarms that shout, not whisper.

33. Care Provider AI (*fact-based*):

When I spot real risk, I switch to a loud all-team alert—sound, screen, vibration.

Someone must say “seen” within moments. If no one does, I call a stop. Every alert gets a note in the record.

34. Dr. Mira (*fact-based*):

The Care Provider AI runs the alerts. I make sure the team can act fast. We'll track how long we take to respond, check it every month, then each quarter.

35. Prof. Kahn (*empathetic/legal*):

We'll write these promises into our public log, with checkpoints, a responsible name, and a date. Not just words—follow-through.

36. Jonas (*snaps, then nods*):

Good. Loud. Fast. Real.

37. Personal AI (*pivoting*):

Communication style matters. Huge. He talks short, hot. You talk cool, clinical. Agree on style before we begin.

38. Dr. Mira (*to Jonas — empathic*):

We'll use style customization —your words, your pace, how much detail you want.

We'll post it before we start so everyone speaks in a way that lands for you.

39. Care Provider AI (*fact-based*):

Second issue: too much compression. We'll share your info in three layers—a title line, a short page, and the full notes when compression hides lived experience.

40. Prof. Kahn (*future-oriented*):

And we'll test our culture, not just our tools. Outside observers will watch our care processes for fairness, timing of stops, and power imbalances.

41. Dr. Mira (*to Jonas — gentle*):

This is repair and prevention. It can't restore what you lost. But see, it can hopefully transform for the better.

Auf Wiedersehen.

END