Agentic AI Design Canvas

DESIGN REVIEW

designtheagent.com
22 June 2026
Round 1

NORTH STAR

Return scarce ophthalmology capacity to the patients who need it by safely discharging routine cataract patients after surgery.

Overall

LOAD-BEARING GAPS

This is a carefully scoped clinical voice agent with named workflow, measurable bars, bounded tools and real human sign-off — one of the more coherent canvases of its kind. The biggest open question is whether a single 48-hour callback rule is safe for every flagged symptom, since some post-cataract symptoms are emergencies, not routine reviews.

COHERES The canvas tells one consistent story: a tightly bounded post-op follow-up, autonomy held to calling and summarising with the discharge decision left to a clinician, and matching tools, rules and oversight. It comes apart only on time-sensitivity — the uniform 48-hour callback does not yet distinguish a routine review from an emergency.

DO THESE NEXT

Rules & Boundaries — Split the callback rule in your Rules & Boundaries cell (Cell 07): name the symptoms that demand immediate or same-day action and keep 48 hours only for genuinely routine flags.

Oversight & Accountability — Add a real-time escalation trigger to your Oversight cell (Cell 09): state what makes Dora hand a live call to a clinician straight away when a patient reports an acute symptom.

Context & Knowledge — State in your Context & Knowledge cell (Cell 05) the trusted source that confirms a case is "uncomplicated" before the agent calls, since the whole boundary rests on it.

TOP RISKS

- Your Rules & Boundaries cell (Cell 07) routes every flagged symptom to a callback within 48 hours, but some post-cataract symptoms (such as sudden vision loss) are emergencies that 48 hours could be too slow for.
- Nothing in your Oversight cell (Cell 09) describes what stops Dora or escalates urgently mid-call if a patient reports something acute during the conversation.
- Your own context asks about extending to pre-operative assessment; that is a different workflow with different symptoms and risks, and it reopens the boundary, the callback rule and sign-off rather than inheriting them.

CROSS-CELL FLAGS

CELLS 07, 09

The single 48-hour callback rule and the review-after-the-call oversight leave no path for an urgent symptom raised during the call.

On a clinical pathway, a fixed delay applied to every flag turns a foreseeable emergency into a missed one — the exact harm the agent exists to prevent.

CELLS 03, 09

Autonomy is set at execute for the live call, but there is no real-time stop or escalation trigger, only after-the-fact summary review.

An agent acting end to end with a patient needs a way to break out mid-task when it is out of its depth, not just a reviewer afterwards.

THE CANVAS, CELL BY CELL

Cell 01 · Target Workflow & Success Metrics

LOOKS CLEAR

WHAT YOU WROTE

Workflow: follow-up about three weeks after uncomplicated cataract surgery. Metrics: sensitivity in catching patients needing review, calls completed autonomously, unplanned care after a clear call.

COACH

Does "unplanned care after a clear call" capture the worst case — a patient who needed urgent attention but was cleared — clearly enough to act on?

Cell 02 · Users & Stakeholders

LOOKS CLEAR

WHAT YOU WROTE

Direct users: post-operative patients on the telephone, and the ophthalmology team receiving flags and summaries. Affected: waiting-list patients who gain the freed capacity, and the accountable trusts.

COACH

Have you separated the patient who could be harmed by a missed symptom from the patient who simply benefits from freed capacity, so each is designed for?

Cell 03 · AI Role & Autonomy

LOOKS CLEAR

WHAT YOU WROTE

Role: autonomous clinical caller. Autonomy: execute, inside a tight loop. It conducts the call end to end and recommends discharge or callback; a clinician reviews the summary before discharge is confirmed.

COACH

You call this "execute" but the discharge is confirmed by a clinician — is there anything the agent finalises during the call itself that no one reviews?

Cell 04 · Performance Needs

LOOKS CLEAR

WHAT YOU WROTE

Regulator-grade bars, published before routine use: 93.75% sensitivity for patients needing review, strong symptom-level agreement with ophthalmologists, and 96.5% of calls completed unaided.

COACH

"Strong symptom-level agreement" is the one bar without a number.

What measurable level of agreement with ophthalmologists counts as passing, and who signs that off before routine use?

Cell 05 · Context & Knowledge

NEEDS DETAIL

WHAT YOU WROTE

The cataract pathway and the five symptom domains that matter after surgery: redness, pain, vision, new floaters, and flashing lights. Restricted to uncomplicated surgeries by design.

COACH

You restrict the agent to uncomplicated surgeries, but you have not said how it knows a given case is genuinely uncomplicated.

Where does the agent get the confirmed "uncomplicated" status for each patient before it calls, and how trusted is that source?

Cell 06 · Tools & Action Channels

LOOKS CLEAR

WHAT YOU WROTE

Telephony: it places real calls and converses in natural language, then writes a structured summary into the clinical record. It cannot prescribe, book theatre time, or act beyond routing a callback.

COACH

Since it converses freely with patients by phone, could anything a patient says steer it into advice or action it should not give, and what blocks that?

Cell 07 · Rules & Boundaries

AT RISK

WHAT YOU WROTE

Only uncomplicated cataract cases. Any flagged symptom routes to a clinician callback within 48 hours. Emotionally complex situations and non-verbal cues are excluded; those belong to humans.

COACH

Every flagged symptom gets the same 48-hour callback, but some post-cataract symptoms are emergencies where 48 hours could be too slow.

Which symptoms should trigger immediate or same-day action rather than a 48-hour callback, and how does the agent tell them apart on the call?

Cell 08 · Memory & Learning

LOOKS CLEAR

WHAT YOU WROTE

Each call becomes a structured record in the patient's pathway. Improvement is audited in public: a supervised validation study first, then a published evaluation of routine deployment.

COACH

Between the validation study and the published evaluation, how does a wrong call get fed back so the agent actually improves rather than just being measured?

Cell 09 · Oversight & Accountability

NEEDS DETAIL

WHAT YOU WROTE

Clinicians review every summary before discharge and call back every flagged patient; the trusts own the pathway; named investigators published the results that license the autonomy.

COACH

Review of summaries is well covered, but nothing describes stopping or escalating urgently during a live call.

If a patient reports something acute mid-call, what makes the agent hand off to a person straight away rather than waiting for the summary review?