

MG VR Experience

A VR empathy simulation that lets users feel the vision, balance, and fatigue challenges of Myasthenia Gravis firsthand.

EMPATHY / HEALTH VR SIMULATION • META QUEST 3 • UNITY

GENRE	Empathy / health-education VR simulation (non-competitive)
PLATFORM / ENGINE	Meta Quest 3, built in Unity, session data logged to Firebase
TARGET AUDIENCE	Empathy-training participants experiencing MG symptoms firsthand (specific deployment audience not specified in the document)

Prepared by Tensei Games • Senior Design & Production Review

1 Executive Summary

MG VR Experience puts the user inside three everyday scenarios — an Apartment morning routine, an Outdoor Park walk, and a Driving trip — while simulating three real symptom categories of Myasthenia Gravis: Vision impairment (Blurry, Droopy/ptosis, Diplopia), Balance instability (joystick-corrected sway), and progressive Muscle Fatigue (a pacing-and-recovery system built around a Stop & Breathe ritual). The design brief is explicit and unusually disciplined for this genre: 'not a game for competition, but an empathetic simulation,' with stated principles of immersion over abstraction, clarity in discomfort, one-touch onboarding, and positive closure.

THE PITCH IN ONE LINE *Three real MG symptoms, three everyday scenarios, one goal: let the user feel — briefly and safely — what a routine morning costs someone living with Myasthenia Gravis.*

The document is dense with production-ready detail: exact difficulty math, a full scoring model, three worked example playthroughs per scenario (best-case, normal, and failure runs), fifteen numbered use cases, and Firebase-backed session logging. That level of specificity is a real strength — it means very little in this study needed to be inferred. The central design question this study examines is whether the mechanics stay true to the stated empathy-first philosophy throughout, or whether familiar game conventions (a points score, a health-bar-style fatigue glow) quietly reintroduce the competitive framing the brief explicitly rejects.

2 Core Loop & Mechanics Analysis

The three systems are well-differentiated by the kind of player behaviour they demand, which is the right foundation for a layered simulation: Vision is passive (the player adapts perception, no input), Balance is continuous corrective input (joystick against sway), and Fatigue is pacing-and-recovery (a walk-speed slider plus a periodic breathing ritual). Layering them — up to three vision states plus one balance level plus one fatigue level simultaneously — is exactly how real comorbid symptom burden compounds, and the design correctly restricts balance and fatigue to one selected severity each while allowing vision effects to stack, since vision impairments are additive perceptual noise rather than competing input demands.

What works

- **The Stop & Breathe recovery ritual is the standout mechanic.** A radial breathing ring that the player syncs a trigger squeeze-and-release to, triggered mandatorily at 85% fatigue, does double duty: mechanically it's a legible, calming minigame, and thematically it directly enacts the real self-management technique MG patients use to manage fatigue. Making it mandatory in Driving (the car auto-pulls over) but a choice-with-consequences in Apartment/Park correctly mirrors real stakes — ignoring fatigue at the wheel is categorically more dangerous than ignoring it on a walk.
- **Diegetic feedback throughout, almost no HUD.** Balance is shown as two vertical lines at the screen edges rather than a bar; vision impairments are camera post-processing, not a menu toggle; haptics and audio ('stumble' sound, tire screech, breathing audio) carry most of the penalty signal. This is a disciplined execution of the brief's own 'immersion over abstraction' principle and it's worth naming as a genuine craft strength.
- **Three scenarios form an implicit difficulty and stakes curve.** Apartment (task-based, low stakes) to Park (ambient social distractions, moderate stakes) to Driving (fatigue builds faster, a missed recovery ends in a

crash cutscene) reads as a natural escalation even though the document presents them as freely selectable rather than a sequence.

Where friction lies

- **The scoring model quietly contradicts the stated 'not a game for competition' philosophy.** Harder settings grant a higher starting score (Base 500 + up to 300 each for Vision/Balance/Fatigue, so a full Hard/High/all-vision run starts at 1400 versus 700–800 for an easy run). A player optimizing for a high final number is incentivised to pick the hardest impairments, which converts a symptom-severity choice into a score-maximising choice — the opposite of the document's own stated intent. This is the single biggest tension in the design as written.
- **The fatigue glow explicitly borrows 'low health' game language.** The document states the green→orange→red edge glow is 'inspired by low health visual feedback in traditional games.' That's a pragmatic legibility choice, but importing a health-bar/game-over visual vocabulary into a simulation of a real chronic illness is worth deliberately re-examining rather than defaulting to convention (see Section 4).
- **No explicit onboarding order across the three scenarios.** The Scenario Selection screen presents Apartment, Park, and Driving as an open choice with no stated recommended order, so a first-time user could select Driving — the highest-stakes, fastest-fatigue scenario — before ever experiencing the calmer Apartment introduction.

DESIGN VERDICT *The three symptom systems are well-modelled and the breathing-recovery mechanic is a genuine design highlight; the scoring model is the one piece of the mechanics that works against the simulation's own empathy-first premise.*

3 User Motivation & Impact

This is not a retention product, and the standard day-1/7/30 lens doesn't apply — the honest question is whether a single session (or a handful, across difficulty settings and scenarios) leaves the user with a genuinely shifted mental model of what living with MG costs day-to-day, which is the document's own stated goal ('understand the daily challenges of MG patients without overwhelming them').

What drives the intended impact

The reflection text is the platform's actual delivery mechanism for empathy, and it's well-built: every scenario ends in a static, motionless reflection room that translates the session's mechanical failures into a first-person real-world parallel — 'You missed 4 balance corrections. For someone with MG, that could mean multiple dangerous stumbles in a single morning routine.' This is the correct design move: the score and mistakes are not the point, they're raw material for the sentence that actually teaches something. The 'Replay / Try Different Impairments' loop is well-judged too — encouraging a user to redo the same scenario at a different severity is the single most effective way this format can build a fuller picture, since MG symptom severity varies day to day for real patients.

Where the intended impact is at risk

Correctly, there are no leaderboards, no timers, and no social/competitive framing anywhere in the document — that restraint is exactly right for this genre and should be protected. The one place the design still courts a

competitive read is the score itself: if a user leaves focused on 'I got 1030 points' rather than the reflection sentence that followed it, the mechanic has been gamified in a way the design philosophy explicitly warns against. Recommend user-testing whether the reflection text alone (with the score de-emphasised or shown only as a percentage/summary rather than a number to beat) produces a stronger self-reported empathy shift than the current points-forward presentation.

IMPACT VERDICT *The reflection-and-replay loop is a well-designed empathy delivery mechanism; protecting it from the competitive pull of its own scoring system is the main open question.*

4 Narrative & Aesthetic Alignment

The document specifies very little conventional 'art direction,' and that's appropriate — the aesthetic decisions that matter here are perceptual and diegetic rather than stylistic, and most are well-aligned with the empathy goal.

What's aligned

- **Reflection copy is the platform's strongest narrative device.** Tying every mechanical failure to a specific, plain-language real-world consequence for an MG patient is what turns a UX document into an actual teaching tool. It should be treated as the single most important piece of writing in the whole project and reviewed with the same rigour as the mechanics.
- **Audio and haptics as a genuine second sensory language.** Directional audio hints replacing UI markers, breathing/heartbeat/tremor audio layered under the fatigue system, and haptic-first penalty feedback (rather than pop-up text) all serve the 'immersion over abstraction' principle and reduce the risk of the experience feeling like a UI-heavy tutorial rather than a lived scenario.

Where it needs sharpening

- **The health-bar-style fatigue glow (green → orange → red) imports familiar game-over connotations.** This is the one visual choice in the document that reads as borrowed from conventional game feedback rather than derived from the lived experience being simulated. It is legible and functional, but it's worth testing with actual MG patients or advocacy groups whether a 'failing health bar' visual metaphor for fatigue feels accurate and respectful, or whether it inadvertently frames a chronic symptom as a game-over state to avoid.
- **Failure-case reflection copy is emotionally strong language that deserves a sensitivity pass.** Lines like 'this could mean total physical collapse before even leaving their apartment' and 'a dangerous day where they cannot drive safely at all' are vivid and memorable, but strong enough that they should be reviewed with MG patients or clinicians to confirm they land as informative rather than alarmist about a real, currently-lived condition.

AESTHETIC VERDICT *The reflection copy and diegetic feedback are genuine strengths; the fatigue health-bar visual language and the most severe reflection lines are the two spots that most need a sensitivity/accuracy pass with the affected community before wider release.*

5 Viability & Deployment Model

The document does not specify a business model, distribution channel, or target deployment context — appropriately, since this reads as an institutional/educational tool rather than a commercial release, and the document should not be read as implying monetisation intent where none is stated. What it does specify is a Meta Quest 3 target and Firebase-backed session logging, which points toward a managed or sideloaded institutional deployment (clinics, patient advocacy organisations, employer disability-awareness training, or medical/caregiver education programs) rather than a public app-store release, though the document doesn't name a specific deployment partner or audience.

What's sound

- **Data logging is deliberately scoped to session mechanics, not health data.** Use Case 14 explicitly states 'only session-level data is logged, no personal health/biometric data' — a genuinely good privacy-conscious design decision for a tool that simulates a medical condition, and worth highlighting to any institutional buyer evaluating the product for compliance concerns.
- **Login is deliberately lightweight.** Username/email only, no password, is a sensible friction-reducer for a short-session empathy tool likely used in group/training settings (e.g., a workshop or classroom) rather than a personal account-based product.

What needs scrutiny

- **There is no described mechanism for measuring the actual outcome the tool exists to produce.** The system logs gameplay telemetry (fatigue events, balance stumbles, collisions) but nothing that captures whether the user's empathy or understanding actually changed — no pre/post attitude check, no post-session survey prompt. For an institutional buyer (a hospital, an employer, an advocacy organisation) to justify deploying this at scale, they will likely need evidence of impact beyond 'sessions completed.' Recommend adding a short, optional pre/post reflection prompt tied into the existing Firebase logging, since the infrastructure to capture it already exists.
- **Session data tied to a real username/email, even without health data, is still sensitive by association.** A record showing a real name performed poorly in an MG-symptom simulation is a softer form of sensitive data than clinical health data, but still worth a clear consent/privacy notice at login — the document doesn't currently describe one beyond format validation on the input field.

DEPLOYMENT VERDICT *The privacy-conscious data scope is a genuine strength for institutional sale; the missing piece is a way to actually prove the empathy outcome the tool is built to deliver.*

6 Risk Assessment & Strategic Recommendations

The document is already unusually thorough for this stage — three fully worked example playthroughs per scenario, fifteen numbered use cases, and explicit safety handling for headset removal and controller disconnects are all signs of a mature production process. The risks below are the ones that stage of detail doesn't yet resolve.

Risk 1 — Scoring incentives run against the stated design philosophy (Design)

- Because harder impairment settings grant a higher starting score, a player motivated by the number rather than the experience is rewarded for choosing maximum difficulty — directly at odds with the brief's own 'not a game for competition' principle. Mitigation: de-emphasise the raw point total in the reflection screen (lead with the qualitative breakdown — balance stability %, correct selections, fatigue red-zone hits — and show score as a minor secondary detail, or drop the point total from the user-facing screen entirely while keeping it internally for analytics).

Risk 2 — VR comfort risk from stacked visual impairments (Technical / Safety)

- Camera blur, partial occlusion, double-vision trailing on head movement, and balance-driven camera sway are, individually and especially combined at Hard/High settings, a known recipe for VR motion sickness. The document's accessibility section covers adjustable impairment intensity and a Pause & Breathe button, but doesn't describe a comfort mode (e.g., vignetting, reduced camera movement, seated-mode) distinct from impairment intensity, nor guidance for a user who becomes physically uncomfortable mid-session rather than choosing to lower intensity in advance. Mitigation: add an explicit comfort-settings layer independent of impairment severity, and specifically playtest the combined Hard/High/triple-vision configuration for motion sickness before any group deployment.

Risk 3 — Symptom fidelity and tone are unvalidated against the real patient community (Content / Ethical)

- Specific figures in the document (4°/second balance sway, exact fatigue point rates, an 85% Stop & Breathe threshold) are precise enough to read as clinically grounded, but the document doesn't cite a source or clinical/patient advisory input. For a simulation whose entire value proposition rests on accurately representing a real, currently-lived condition, unvalidated specifics are a credibility and, more importantly, a respect risk. Mitigation: before wider release, review the symptom mechanics and the reflection copy's tone with MG clinicians and a patient advocacy organisation (e.g., a national myasthenia gravis foundation), and be prepared to adjust both numbers and language based on that feedback.

Bonus Risk — No explicit first-scenario recommendation (Onboarding)

- Presenting Apartment, Park, and Driving as an unordered choice risks a first-time user selecting Driving — the fastest-building fatigue and highest-stakes scenario — before the gentler Apartment scenario has taught them the three mechanics at a manageable pace. Mitigation: recommend (without hard-locking) Apartment as the suggested first scenario in the Selection screen, consistent with the onboarding tutorial's own low-stakes framing.

RISK SUMMARY *The technical and production groundwork is strong; the remaining risks are about protecting the empathy-first intent from its own scoring system and validating both fidelity and tone with the community the simulation represents.*

7 Final Verdict

MG VR Experience is a carefully constructed attempt at experiential empathy education, and the craft is real: three well-differentiated, symptom-grounded mechanical systems, a diegetic feedback language that avoids breaking presence, a mandatory-vs-optional recovery ritual that correctly mirrors real-world stakes across scenarios, and — the project's strongest single asset — reflection copy that converts abstract score deductions

into concrete, first-person understanding of what a routine task costs an MG patient. The privacy-conscious approach to session data is a genuine credit to the team's judgment on a sensitive product category.

The open risks are addressable rather than fundamental: reconcile the scoring model with the stated non-competitive philosophy, add a comfort layer distinct from impairment intensity, and — most importantly for a project representing a real medical condition — validate the specific symptom figures and the tone of the reflection copy with MG clinicians and the patient community before broader deployment. None of this requires rethinking the core design; it requires the same rigour already visible in the mechanics being extended to the human side of the simulation.

A note on subject matter: this project simulates a real, currently-lived medical condition. This study has evaluated it on its own stated terms — empathetic accuracy and emotional impact, not entertainment metrics — and that same standard should carry into production: treat clinician and patient-community review as a required step before wider release, not an optional polish pass.