

Feature Matrix & Product Roadmap

Museum AR Smart Glass Guide

HumanityAI'd — July 2026

CONFIDENTIAL

1. Feature matrix (v1.0 — shipping today)

Capability	Status	Notes
Hands-free exhibit recognition (CLIP + FAISS)	▪ Live	Verified 0.995 similarity on enrolled exhibits
Multi-angle enrolment (1–3 photos/exhibit)	▪ Live	Live index update, no retraining
8 languages (en, ar, fr, es, zh, ru, hi, de)	▪ Live	Engine supports 23
Arabic + RTL content	▪ Live	name_ar fields, RTL data model
Neural TTS narration (Chatterbox multilingual)	▪ Live	Museum-tuned loudness, Opus
Human review/approval of narration	▪ Live	Arabic requires explicit approval
Dwell-based auto-play ("look & listen")	▪ Live	Tunable trigger/cooldown
WebRTC live streaming from glasses	▪ Live	H.264, DataChannel results
Snapshot recognition (fallback)	▪ Live	POST /recognize
Guided tours (ordered stops)	▪ Live	Multilingual tour metadata
Sessions & engagement analytics	▪ Live	Dwell, play-through, popular exhibits
Admin content panel	Live	CRUD museums/galleries/exhibits/translations
On-prem Docker deployment	Live	Single GPU host, no cloud
Prometheus/Grafana ops	Live	Metrics + dashboards
Glasses settings (web panel)	Live	Display/config for the fleet

2. Known limitations (honest)

- Single-node / single-worker backend — one GPU box $\approx$ one museum, capped concurrent streams.
- Narration generation is an admin/offline activity (GPU-intensive); runtime serving uses pre-generated audio.
- Authentication and transport security are POC-grade — hardening scheduled (see Due-Diligence report).
- WebRTC streaming path pending on-hardware capacity verification; snapshot path fully proven.

3. Roadmap

v1.1 — Hardening (0-3 months)

- JWT validation + role-based access; per-device provisioned tokens; TLS + pinning on glasses
- Background TTS job queue; thread-pooled inference (unblock event loop)
- Automated tests + CI on both repos; dependency CVE upgrades
- FAISS rebuild-from-DB durability; `museum_id` multi-tenant filtering

v1.5 — Reach & scale (3-6 months)

- **BYOD mode**: visitor's own phone as an alternative to glasses (widen adoption)
- Horizontal scale: stateless API + separate inference service (multi-museum, 100+ concurrent)
- Thermal-adaptive capture ladder on glasses (battery/heat)
- Offline narration bundle on glasses for zero-latency common exhibits

v2.0 — Experience (6-12 months)

- Indoor positioning / wayfinding overlays ("navigate to the next tour stop")
- Personalization: adaptive tour depth (kids / expert / accessibility modes)
- Live analytics dashboard for curators (heat-maps of attention)
- Voice Q&A: "tell me more about this artist" (on-prem LLM)

v2.x — Platform

- Multi-site management console (museum groups / ministries of culture)
- Content marketplace: share/curate exhibit packages across institutions
- Accessibility: sign-language avatar track, audio-description mode for low-vision visitors

4. Differentiators to protect

1. **On-prem data sovereignty** — no cloud egress of visitor imagery (procurement-critical in the GCC).
2. **Dwell "look & listen" UX** — no buttons, no numbers, no app.
3. **Arabic-first multilingual pipeline** with human approval — built for the region, not translated into it.
4. **Minutes-to-enrol content tooling** — museums stay self-sufficient.