

AMRS *A Sequential Music Decision System for Adaptive Delivery*

AMRS (Affective Music Recommendation System) is LUCID's **music sequencing and decision layer**. It determines which music to play, and in what order, based on a listener's measured affective state and the intended goal of a session. AMRS operates on **state signals produced by BioMIR**—valence, arousal, absorption, and engagement—and makes sequencing decisions within **human-defined constraints**.

SYSTEM OVERVIEW

AMRS models music delivery as a **sequential control problem**, where decisions influence future listener state and are evaluated over time.



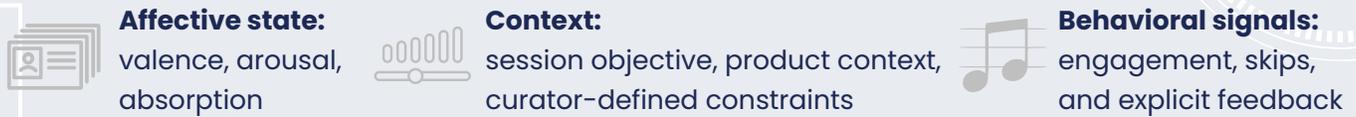
Rather than selecting the next item independently, AMRS evaluates music trajectories and adapts delivery based on **measured response**.

HOW LUCID IS DIFFERENT

Dimension	Traditional Music Systems	LUCID'S AMRS
Control Loop	Open-loop; static playlists or one-step recommendations	Closed-loop , adaptive system driven by measured human state
Decision Unit	Single-track, stepwise decisions	Sequential decision-making across full music trajectories
Safety & Compliance	Limited control, black-box behavior	Safe, gated deployment suitable for regulated settings
Signal Inputs	Behavioral data (clicks, skips, popularity)	Physiological + affective signals via BioMIR
Personalization Basis	Correlation-based (similar users, similar tracks)	Biologically informed, state-aware personalization
Learning & Testing	Live experimentation on users	Offline, world-model-based RL with counterfactual evaluation
Human Governance	Minimal human oversight after launch	Curator-defined rules, constraints, and boundaries
Optimization Horizon	Immediate engagement or similarity	Long-term, session-level outcomes

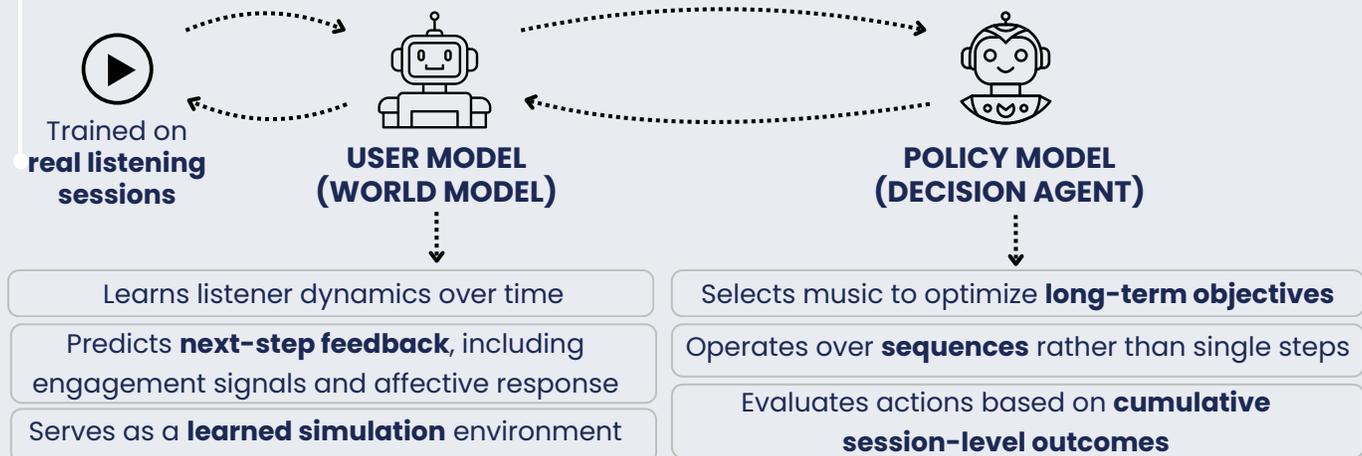
STATE REPRESENTATION & INPUTS

AMRS consumes a structured representation of **listener state and session context**, primarily derived from BioMIR:



SEQUENTIAL DECISIONING ARCHITECTURE

AMRS uses a **two-phase, world-model-based reinforcement learning architecture**, developed in collaboration with MILA.



LEARNING & OPTIMIZATION

AMRS follows an **offline-first training and evaluation approach**:

- ✓ Policies are trained and evaluated within the learned user model
- ✓ Counterfactual policies can be explored **without live user exposure**
- ✓ Candidate policies are **gated before deployment**

Optimization is driven by **multi-objective reward shaping**, combining:



HUMAN GOVERNANCE & PLATFORM ROLE

AMRS operates within a **human-governed framework**:

- Music pools, transitions, and boundaries are **curator-defined**
- The system does not autonomously expand beyond these constraints

Within the broader platform:

- **BioMIR** measures and represents human state
- **AMRS** performs sequential decisioning
- **ABS** applies modulation during playback



AMRS applies offline-trained sequential decisioning to guide music delivery using measured affective state.