


How does learning change the pleasure of the same musical event?

Amaç Erdem | amace@bu.edu | ORCID 0009-0000-1711-0036

Abstract. The same musical event may feel different after different listening histories. We will combine participant-provided listening records with Musical Intelligence (MI) features from verified recordings to estimate individual exposure profiles. Recognition and continuation tasks will test these estimates. A controlled experiment will then manipulate learned expectations while keeping test audio unchanged and measure pleasure. Separate comparisons will test whether history-informed MI adds predictive value beyond self-report, listening frequency and acoustic baselines. Exposure is treated as evidence about listening experience, not a direct measurement of memory.

Academic edition r3 · 12 October 2026 · Source: v3.0, 7 October 2026. Proposed design; not preregistered; no new participant results.



Assigned mappings · schematic illustration · no results shown

Fig. 1. The critical comparison holds the cue–target waveform constant across assigned histories.

1. The unresolved question

Research connects musical pleasure to surprise and uncertainty, but reports different directions of their interaction [1, 2]. Learning can also create preferences for unfamiliar musical structures [3]. We ask how measured listening experience and experimentally learned context relate to the response to an unchanged musical event.

We will manipulate the probability of two target continuations and the uncertainty of the cue that precedes them. Cue meanings will be counterbalanced across listeners. Thus, an identical cue–target recording can have different learned meanings without changing its waveform.

2. Two distinct tests

H1. Contextual uncertainty changes the difference in pleasure between less- and more-expected targets. The primary contrast is a difference of differences, tested in both directions.

H2 · Secondary. MI-derived exposure profiles improve held-out prediction of familiarity, expectation and pleasure beyond self-report and listening-history baselines. H1 tests experimentally assigned learning; H2 tests prediction from observed histories. Neither establishes a unique neural mechanism.

3. What is held constant?

Within each musical family, target A is the same audio file wherever it occurs; the same applies to B. Each target is heard equally often overall. Four cues assign A and B probabilities of either 1/2 or 1/8, under low or high uncertainty. Three other continuations preserve these distributions.

This separates **conditional predictability** from simple target familiarity. Cue acoustics and target identities are counterbalanced, with sequence and repetition effects included in analysis. Specific cue–target pair familiarity remains linked to probability and limits interpretation.

4. Scope and participants

The planning sample is **192 adult non-musicians**, plus a separate pilot of 24. Twelve new musical families will be tested; each listener learns one. English comprehension, headphone screening [6] and task practice are required.

Musical experience and reward sensitivity will be measured [4, 5]. History donation is optional; its absence does not exclude a listener from H1. H2 uses the consenting subset with usable records, with coverage and selection reported. Neither dopamine nor brain activity is measured.

Listening history and the exposure profile

Identify what was heard

Participants may provide track records from Spotify, Apple Music or another source. Keep title, artist and available recording identifiers. Counts, dates and listening durations are used only when supplied; a list of names cannot establish frequency or recency.

Locate candidate recordings through catalogues or public pages, including YouTube. Confirm artist, version and duration rather than accepting the first title match. Record live/cover/remaster differences and unresolved matches. Track-level listening duration does not identify which moments were heard.

Build a musical exposure profile

Analyse the verified recording using a frozen MI configuration. R^3 describes acoustic structure; T^3 describes its temporal organisation. A separate comparison adds the six prespecified C^3 outputs already shown in the architecture. Their names remain computational labels.

Combine those features with the available history to estimate exposure to musical structures. Prespecify aggregation, recency weighting and missing-data rules in development data. Report matched-track and listening-time coverage where possible. A profile is an estimate to test, not a diagnosis or a reconstruction of autobiographical memory.

5. Learning, listening and reporting

Before testing, consenting listeners provide Spotify, Apple Music or equivalent track records. Names and artists locate candidate recordings; versions are verified before MI analysis. Frequency and recency are used only when present. A title list supports a repertoire profile, not a timed exposure history.

An added recognition and continuation session validates the frozen profile, using separate calibration and held-out items. The controlled experiment retains its planned 50–60 minutes; the pilot will determine the added history-task burden. Pleasure and later secondary responses remain in separate blocks.

After consent, sound checks and practice, the session contains 128 listening trials followed by 128 pleasure trials. Resolution and replay follow in separate blocks.

Continuous pleasure · primary

“How much pleasure do you feel right now?”

Only this response is requested during the primary test. A stationary, initialised slider is a valid response.

The response scale runs from 0 (none) to 100 (very strong).




Fig. 2. Probability and uncertainty are separated by design. A and B are the focal targets; C–E are fillers. Across four cues, A and B each occur 40 times in both exposure and test. Each test yields 80 target trials: 32 more-expected and 8 less-expected trials at each uncertainty level. Test frequencies preserve training probabilities; cells are not artificially balanced by overplaying rare targets. Cue-to-distribution assignments vary across listeners.




Fig. 3. An event-aligned outcome. Pleasure is sampled against the audio clock at 20 Hz. The fixed response window allows for reporting delay; it does not identify the exact onset of an internal feeling. Interrupted trials and uninitialised responses are flagged.

Predicted continuation · learning check

“Out of 100 occasions, how often would each ending follow this cue?”

Four cue-specific allocations before exposure, after exposure and after the primary test. No outcome feedback.

Allocations across A–E sum to 100.

Perceived resolution · secondary

“Press when the musical idea feels settled.”

Sixteen later trials. Closure is kept distinct from pleasure; non-response is not automatically failed attention.

The R key marks a moment.

Replay preference · secondary

“Which sequence would you like to hear again?”

Eight pairs; one choice is implemented at the end. Curiosity and pleasure may make different predictions.

6. What Musical Intelligence can explain

The original audio architecture is retained. R^3/T^3 descriptors and six existing C^3 outputs form candidate representations. A separate, proposed listener-profile layer combines these measurements with observed listening history; it is not an implemented memory mechanism.

Audio representation + listener history. Frozen MI returns the same scores for the same recording. The proposed profile layer adds track identity, repeat exposure, recency and feature similarity where available. Its aggregation rules are fitted in development data and frozen before held-out testing. Profile scores are hypotheses about familiarity, not memories or autobiographical diagnoses. The current MI comparison remains offline.

7. Analysis, decisions and next steps

7.1 One primary behavioural test

For each target trial, we calculate the change in pleasure specified in Figure 3. A mixed-effects model estimates probability, uncertainty and their interaction, accounting for listeners, musical families, recordings and trial order. H1 tests the interaction at two-sided $\alpha = .05$; effect sizes and 95% intervals are primary reporting units.

Probability allocations test whether learning occurred. Main-study listeners are not excluded for weak learning or low pleasure. Failed learning limits interpretation of a null H1. Resolution and replay are secondary; neither establishes causal mediation.


7.2 Controlled-learning comparison

Four models predict the same pleasure outcome: **A**, R^3/T^3 acoustics and timing; **B**, A plus six MI scores; **C**, A plus a sequential cue–continuation learner; **D**, C plus the MI scores. B–A tests MI’s audio representation; D–C tests its contribution beyond an explicit learning account (Fig. 5 in the original external figure atlas).

The learning model receives the actual ordered exposure history. All models share matched acoustic access and fitting budgets. Entire musical families and their listeners are withheld together; tuning uses development folds only. The primary predictive metric is held-out mean absolute error. Predictive gains do not validate a named brain mechanism.

7.3 Pilot and registration gate

The pilot uses separate musical families to check learning, discrimination, timing and burden. $N = 192$ is a planning figure, not a power result. Simulation must support detection of a four-point interaction on the 0–100 pleasure scale before the final sample is registered. Stimuli, randomisation, exclusions, windows, code and model versions are frozen before recruitment.



SPH Spatiotemporal Prediction Hierarchy
 WMED Working Memory-Entrainment Dissociation
 SRP Striatal Reward Prediction
 PWUP Precision-Weighted Uncertainty Processing
 UDP Uncertainty-Driven Pleasure
 RPEM Reward Prediction Error in Music

Fig. 4. The implemented routes used in the model comparison. Coloured routes use the manuscript's F2/F5/F6 palette; unused routes remain gray. Dashed lines read R^3 directly; solid lines pass through T^3 or upstream mechanisms. The selected paths use 17 R^3 features and 49 temporal keys. PWUP and WMED supply internal values. The 28 September source audit is historical, static evidence. Counts are declarations; archived artwork values are not study results. Regional and chemical outputs are outside the experiment.

Table 1. Prespecified MI candidates. Incremental prediction is tested jointly; named-mechanism interpretations remain exploratory.

SPH.P1.prediction_error Audio-derived prediction-error candidate	UDP.P0.context_assessment Audio-derived context candidate	UDP.P2.reward_computation Pleasure-related candidate
SRP.P0.wanting Approach-related candidate	SRP.P1.liking Pleasure-related candidate	RPEM.current_rpe Reward-error candidate, not measured RPE

7.4 Listening-history validation

Compare self-report alone, self-report plus available listening counts/recency and artist/genre baselines, then the same model plus an MI-derived exposure profile. Profile ablations use R^3 , $R^3 + T^3$, and $R^3 + T^3$ + the six C^3 candidates, with matched fitting budgets. Freeze histories before testing; evaluation items never tune the profile.

Primary history validation is recognition discrimination on held-out familiar candidates and newly created, matched foils; continuation prediction and pleasure are secondary outcomes. "Absent from the export" does not mean unheard. Hold out listeners and recording families; report matching coverage and sensitivity to omitted tracks. Predictive improvement is not a causal effect of past listening.

7.5 Recording identity and retention

Match title, artist, recording version, duration and identifiers where available. Catalogues and public pages, including YouTube, may locate candidates; uncertain, live, cover or remastered substitutes are flagged. Extraction requires an authorised audio source and applicable-use review; academic purpose and deletion alone do not establish permission.

No audio is redistributed. Temporary analysis audio is deleted after extraction and quality checks, within a planned seven-day maximum. Retain only necessary pseudonymous features, recording identifiers, checksums, match decisions and extraction/deletion logs under the consented retention policy. Experimental playback has separate source and retention requirements.

7.6 Readiness

The history layer is proposed. Record import, matching, permitted audio access, extraction, behavioral calibration and its

separate precision analysis remain to be completed. The app illustrates the design and collects no history or responses. Ethics approval and registration precede recruitment; neither memory diagnosis nor autobiographical inference is claimed.

Additional history safeguards

Track-level listening duration does not identify which moments were heard. Prespecify aggregation, recency weighting and missing-data rules in development data, and report matched-track and listening-time coverage where possible. Recognition foils permit false-alarm estimation. Record exports have incomplete coverage across devices, services and everyday listening; structural familiarity in new music is a prediction requiring independent behavioral validation.

Temporary analysis audio is deleted within seven days of acquisition, including unresolved or failed items. Platform access requirements remain applicable. Parsers, recording matching, the profile model and research data collection remain to be implemented and validated.

References

- [1] Cheung et al. (2019). Uncertainty and surprise jointly predict musical pleasure. *Current Biology*. 10.1016/j.cub.2019.09.067.
- [2] Mas-Herrero & Marco-Pallarés (2025). Predictive processes shape individual musical preferences. *PNAS*. 10.1073/pnas.2500494122.
- [3] Kathios et al. (2024; online 2023). Generating new musical preferences from multilevel mapping of predictions to reward. *Psychological Science*. 10.1177/09567976231214185.
- [4] Müllensiefen et al. (2014). The musicality of non-musicians. *PLOS ONE*. 10.1371/journal.pone.0089642.
- [5] Mas-Herrero et al. (2013). Individual differences in music reward experiences. 10.1525/mp.2013.31.2.118.
- [6] Woods et al. (2017). Headphone screening to facilitate web-based auditory experiments. 10.3758/s13414-017-1361-2.
- [7] Albury et al. (2026). Stimulus predictability and liking enhance auditory-motor encoding and memory for melodies. 10.1111/nyas.70310.

Model provenance

MI Systems Neuroscience manuscript, 22 September 2026; static source audit, 28 September 2026. The original map and channels are unchanged. The v3 listening-history specification distinguishes proposed additions from implemented audio routes. No engine or listener-profile validation is claimed.

Protocol and materials

Source proposal: v3.0, 7 October 2026. mymusicalmind.io/learning/. This edition includes the substantive listening-history companion, preserves the four numbered manuscript figures as original vector PDF objects and leaves the existing figure atlas external. The original six-channel and 20-Hz specifications are retained; this formatting edition does not amend the study protocol.

History-data documentation

The source companion records a review on 7 October 2026 of Spotify export fields, Spotify Developer Policy, Apple Music API scope and YouTube terms. These links document source-access requirements; no account connection or audio acquisition is part of this edition.