CURIO

A MORE HUMAN FUTURE FOR MUSIC DISCOVERY

A Research through Design Exploration of Personal Objects and Interpretive AI in At-Home Music Streaming

Joe Oliver Grimshaw // Graduation Report

AI Statement

This project treats artificial intelligence as a design tool rather than a substitute for human judgment. I used AI critically and transparently throughout. Because AI was both a tool used in the project and an object of study, the following statements detail where and how it was involved.

1. Literature search. Undermind helped me gain an initial overview of unfamiliar fields, while Claude helped me trace references and follow citation trails. Neither tool determined how I interpreted a source. I read every source I cited in full, and no claim or citation was included solely based on an AI-generated summary.
2. Building. I used Claude Code (Model Opus 5) to support the development of exploration interfaces and Design provocations through prompt-based coding. I relied extensively on Claude when designing and selecting the electronic systems. This allowed me to address gaps in my technical knowledge that would otherwise have limited the scope and complexity of what I could prototype. I remained responsible for defining the intended behaviour, testing the outputs and deciding what was incorporated into the final prototypes. The Lab [§Chapter 4.4] exemplifies this approach by making the system's behaviour transparent and inspectable, allowing me to retain control over how the system acts.
3. System guidance and context. I created a set of markdown files that instructed the system how to respond during different tasks. These included guidance for creating diagrams, safeguards requiring it to flag rather than introduce uncited sources, and detailed context documents explaining the project's aims, methods and terminology. Giving the system access to this material helped ensure that its responses remained relevant to the project rather than being based on generic assumptions. Because agency was central to the project's exploration, particular care was taken to ensure that the system supported my goals and decision-making without directing the exploration itself.
4. Writing structure. I used Claude at different stages of the writing process, mainly to help with word choice, structure, spelling, and making argumentation more precise. To keep the thesis consistent with my own style and argument, I created guidelines that set out my preferred terminology, reasoning, and tone. I reviewed and edited every suggestion before deciding whether to include it in the final text.
5. Critical reading. I provided Claude with structured context documents and used it to pressure-test my framing, identify gaps and review sections of writing. I treated its responses as prompts for further reflection rather than as authoritative answers. Suggestions were evaluated against the research material, the design process and my own judgment.
6. Low-consequence work. I also used AI for preliminary and supportive tasks, including first passes, early diagrams, restructuring notes, document organisation, light image editing in Google Photos and visual inspiration through Vizcom. Using AI for these tasks reduced time spent on low-impact tasks and left more time for design exploration and reflection.

The ideas, interpretations, design decisions, analysis and conclusions presented in this thesis are my own. AI contributed to how parts of the work were explored, developed and expressed, but it did not determine the meaning of the research or the claims drawn from it. This is especially important in research through design, where the conclusions emerge from the relationship between making, reflection and interpretation and cannot be delegated to a model.

This approach follows TU Delft's guidance on the responsible use of AI. I take full responsibility for the content, accuracy and final form of this thesis.

CURIO: A More Human Future for Music Discovery

A Research through Design Exploration of Personal Objects and Interpretive AI in At-Home Music Streaming

PROGRAM

MSc Design for Interaction

Industrial Design Engineering

Delft University of Technology

GRADUATION

August 2026

MASTER THESIS

Joe Oliver Grimshaw

SUPERVISORS

Dr.D.S. Murray-Rust

Prof.dr. G.W. Kortuem

TU Delft Delft University of Technology

Preface

Left page:

Figure 0.1

An image from Antoine Geiger's SUR-FAKE series illustrating how digital technology can absorb human attention, distort social interaction and create a sense of dependence on mobile devices.

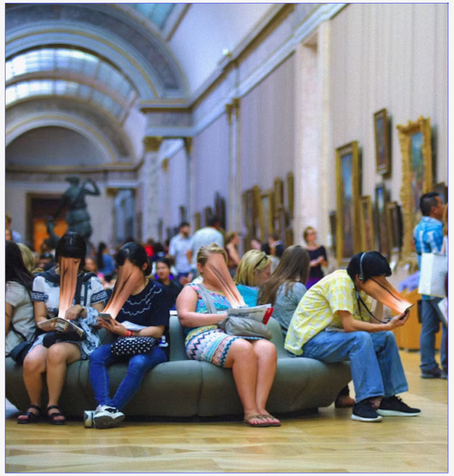
(Geiger, 2015)

Graduation projects begin in many ways: with a lecture, an industry brief, or a last-minute consultation with ChatGPT. Mine began with a feeling I could not shake. I had grown increasingly uneasy about the way digital products were changing my relationship with the world, though I struggled to explain exactly what felt wrong. I only knew that the rich experiences, with all their interaction, effort, and reflection, were being streamlined into efficient conveniences. This project became an attempt to understand that unease and to consider how I might respond to it as a designer.

Nostalgia shaped the project's starting point, but this is not a Luddite's guide to design. Rather than arguing for a return to the past, the project looks back at earlier experiences to ask what may have been lost, which human needs are no longer being met, and how future technologies might respond. Growing up, with access to my Dad's music collection, listening to music was tactile. I flicked through records, carefully lowered a needle onto vinyl, and rearranged old CDs simply for the pleasure of handling them. As streaming became my primary way of listening, these rituals gradually disappeared. Music became more accessible than ever, but it also began to feel weightless and intangible.

Throughout the project, I repeatedly returned to the experience of visiting a record store. I did not simply miss the physical format; I missed the searching, the suspense, the interaction, and the possibility of encountering something unexpected. You work through a crate without necessarily knowing what you are looking for. Nothing is selected in advance. You discover something for yourself, and because you found it, you develop a sense of ownership over it. It becomes part of your identity.

This led me to question design's obsession with frictionlessness [Figure 0.1]. Friction is usually treated as a problem to eliminate, yet effort can produce attention, uncertainty can generate curiosity, and choice can create meaning. These ideas became central to my exploration of how design might preserve participation rather than reducing every interaction to convenience. This concern extends far beyond music: maps choose our routes, TikTok selects what we watch next, smartwatches tell us when to sleep, algorithms are increasingly shaping our decisions. Convenience is valuable, but it can quietly turn participation into passivity. Artificial intelligence could intensify this shift, or it could offer an alternative. This thesis explores that possibility utilising a Research through Design investigation into unfamiliar territory. We live in a world designed by people. Why, then, does so much of it fail to feel human?



Acknowledgements



Although this project was mine to undertake, I could not have completed it without the people who generously shared their time, support, and advice, helping me find my way whenever I felt lost.

To my supervisory team, Dave and Gerd, thank you, first, for believing in the idea enough to commit your time to me over the past 20 weeks. I was ecstatic when I heard that you both wanted to get involved. I can't overstate how fantastic it has been to receive guidance from people whose papers I have read and whose work and approach I greatly admire. Your encouragement to experiment and the freedom you provided allowed me to create a project that expresses something about myself. For that, I feel a great sense of pride in the work I have produced.

It has been a privilege to be supervised by two people who care more about compelling ideas than rigid rubrics. Your critical and insightful advice gave me the confidence to embark on a journey without a clear view of where it would end. This is a lesson I hope to carry forward: to trust my instincts and have the courage to pursue them.

To the participants who gave up their time to help me fully bake the half-baked ideas I served to you, thank you for your patience and willingness to engage with my mad creations.

To my beautiful partner, I know I can be a pain to be around when I get stressed, so thank you for continuing to be there. I'm sorry you have had to listen to me talk about little else for a while now. Your patience and perspective always calmed the storm, and I couldn't have got here without you.

And to my parents, I have all of this to thank you two for. Your continued support with my move to the Netherlands over the past two years has allowed me to have such an enriching and unforgettable experience. To Dad, even though I mock your taste, you are the reason I fell in love with music [Figure 0.2]. As you can see from the following project, it is a passion that I will take with me wherever I go.

Right Page:

Figure 0.2

Photograph taken by the author at his first concert.

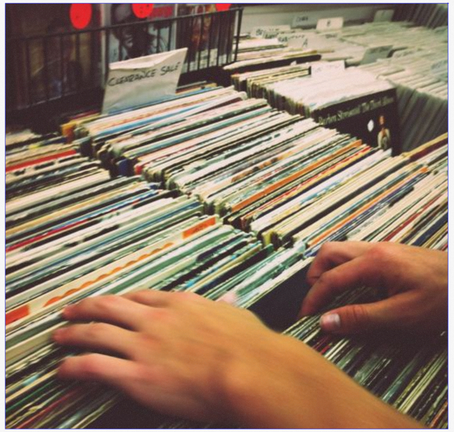
Abstract

Mainstream music recommendation systems rely heavily on prediction algorithms that translate listeners' unformed desires into explicit, nameable categories, leaving limited space for open and exploratory listening. This study addresses this limitation by exploring how to support Stimmung, the felt and preverbal attunement through which listeners sense and orient themselves toward a musical desire before it can be consciously articulated or named. To support this form of emergent musical desire, the study introduces nourishment as an interaction design condition that sustains the listener's active interpretive work rather than replacing it with passive, automated delivery.

Using a Research through Design methodology, the project identified dimensions of experiential loss associated with the shift from physical record stores to frictionless streaming. These insights informed the iterative, autobiographical development of Curio, a domestic, object-based AI system. Curio invites users to place personally meaningful household objects on a surface. A multimodal AI then interprets the objects' material qualities without identifying them by name, using this interpretation to retrieve unfamiliar music. This exploration resulted in an annotated portfolio of provotypes examining how agency is distributed among the user, the object, and the machine.

The research demonstrates that making AI reasoning fully transparent does not necessarily empower users. Instead, explicit justifications can function as persuasive devices that encourage passive acceptance of algorithmic recommendations. By contrast, withholding definitive explanations and introducing "productive ambiguity" can invite listeners to engage in meaningful co-interpretation. The study also shows that interaction design must correspond to the level of commitment associated with different media formats. Albums, as high-commitment formats, warrant slow and deliberate gestures, whereas songs, as lower-commitment formats, require more fluid and continuous interaction landscapes.

Ultimately, this study transforms AI-mediated music discovery from a transactional exchange into an embodied and reflective ritual. It contributes transferable design knowledge through a framework comprising the stages Locate, Name, Decide, Translate, and Provoke. This framework enables designers working in other automated domains to reintroduce nourishment into digital contexts characterized by abundance.



Left page

Figure 0.3

A person crate-digging, carefully feeling their way through a collection of records in search of a meaningful discovery. (Harper's Bazaar France, n.d.)

Right Page

Figure 0.4

Mock-up illustrating people interacting Provocation 4 one of my final research materialisations.

Table of Contents

01 Introduction 16
1.1 Project Background 18
1.2 Design Brief 19
1.3 Research Questions 20
1.4 Research Process 22

02 Literature Review 24
2.1 Limits of Streaming Discovery 26
2.2 AI in Music Recommendation 28
2.3 Existing Landscape of Query Removal 31
2.4 Stimmung and Uncertain Intent 36
2.5 Discovery in the Record Store 38
2.6 Physical Objects as an Interface 42
2.7 Objects as Carriers of Personal Meaning 43
2.8 Prediction to Interpretation 44
2.9 The Research Gap 46

03 Methodology 48
3.1 What is RtD 50
3.2 How was RtD Applied 52

04 Explorations 54
4.1 Research Process - User Journey Map 56
4.2 Loss Dimensions 60
4.3 Research Tool 62
4.4 Exploration 1 - Intro 64
4.5 Exploration 2 - Intro 74
4.6 Exploration 3 - Intro 86
4.7 Exploration 4 - Intro 94
4.8 Exploration 5 - Intro 102
4.9 Exploration 6 - Intro 110
4.10 Tensions form the Explorations 116
4.11 Reflection and Synthesis 118

05 Research Provocations 120
5.1 Provotype 1 122
5.2 Provotype 2 132

06 Design Provocations 142
6.1 Two Economies of Commitment 144
6.2 Provotype 146
6.3 Provotype 4 150

07 Discussion 154
7.1 The Case for Using the Past as Design Material 156
7.2 The Case for Personal Objects as Inputs 158
7.3 The Case for Effective Ambiguity 160
7.4 The Transferable Method 162

08 Conclusion 164
8.1 Answering the Research Questions 166
8.2 Reflection & Limitations 173
8.3 Personal Reflection 175
8.4 Refrences 176
8.5 Image Refrences 181

List of Abbreviations

3D	Three-dimensional
AI	Artificial intelligence
API	Application programming interface
DJ	Disc jockey
GPS	Global Positioning System
MIDI	Musical Instrument Digital Interface
ML	Machine learning
MVP	Minimum viable product
NFC	Near-field communication
PETG	Polyethylene terephthalate glycol-modified
PrEmo	Product Emotion Measurement Instrument
Q&A	Questions and answers
RtD	Research through Design
UMG	Universal Music Group

Glossary

A

AI as a collaborator. A framing in which AI contributes a provisional, associative reading that the listener can question and develop, rather than making a final decision on the listener's behalf.

AI reading. The system's associative and qualitative account of an object. It is deliberately incomplete and revisable: it gives the listener something to work with without claiming to identify the object or explain what it definitively means.

Ambiguity. A design quality that leaves room for the listener to interpret and revise a relationship between object and music. In this project, it is not the same as hiding how a system works.

Annotated portfolio. A set of related design artefacts whose contrasts generate knowledge. Curio's four provotypes are read together, not as competing versions of one final product.

Association. The proposed connection between a personal object and music that AI produces. It is not presented as a correct or definitive interpretation; the listener decides whether and how it resonates.

Associative AI. AI used to propose a possible connection between the perceptible qualities of a personal object and existing music, without treating that connection as a fixed category or correct answer.

Autobiographical design. A Research through Design approach in which the designer's own sustained use, reflection and situated experience are part of the research process.

C

Cognitive commitment cost. The mental effort and attention a listener takes on when choosing a musical unit. The project argues that an interface should not demand more effort than that unit warrants.

Curio. The project's domestic, object-based music-discovery system. A listener presents a personally meaningful object; AI responds with music while leaving the relationship open to interpretation.

D

Design implication. A provisional design direction drawn from literature, fieldwork or an exploration. It guides making without prescribing a final solution.

Design provocation. A deliberately constructed artefact used to make a design position visible and discussable, rather than to claim a market-ready solution.

Design space. The range of possible arrangements of listener, object and AI explored by the project, especially different ways of distributing interpretive work.

Domestic music discovery. Exploratory music listening situated in the home, where personally meaningful objects and everyday surroundings can take part in the encounter.

E

Effective ambiguity. Ambiguity that is anchored in the listener's unresolved experience, recognisably connected to something real, and navigable through embodied action. It avoids both over-explanation and randomness.

Embodied action. Bodily activity that is meaningful to the interaction, such as searching for, selecting, placing or moving an object, rather than merely operating a control.

Encounter before a clear request. A situation in which a listener can be drawn towards music before stating a precise preference, query, genre or mood.

Exploratory listening. Listening undertaken when a person cannot yet clearly say what they want. The project treats this not-knowing as a valid starting condition for discovery.

Glossary

F

Finding Resonance. Provotype 2. An interaction in which the listener can move between explanatory and atmospheric readings, testing whether adjustable explanation can support rather than settle interpretation.

Five-Step Transferable Method. The project's method for reintroducing nourishment beyond music: Locate a superseded practice, Name its losses and gains, Decide what to recover/retain/hold open, Translate through a physical-digital exchange, and Provoke a position through making.

I

Interpretive authority. The authority to decide what an object-music relationship means. Curio aims to keep this with the listener even when the system supplies an association.

Interpretive distribution. The arrangement of interpretive work among listener, object and AI. The provotypes compare different distributions rather than seeking one universally correct balance.

Interpretive work. The ongoing activity of noticing, questioning, connecting and revising meaning. The central design question is where this work sits among listener, object and AI.

J

Journey map. A visual account of an experience across time. The record-store journey map was used to identify experiential qualities that streaming can make less available.

L

Listener agency. Control over where discovery goes next, rather than a general feeling of being in charge. In Curio, this includes choosing the object, interpreting the response and steering subsequent action.

Listening persona. A region of taste, memory and association that becomes carried by a listener's objects over time. In Long Play, a persona can be brought to the system or held back by choosing which objects to carry over.

Long Play. Provotype 3. A speculative album-oriented interaction where an object's arrangement gradually orients a personal library and lifting it commits the listener to an album, played whole and in order.

Loss dimensions. Nine experiential qualities identified from the record-store journey and literature. They serve as a framework for restoring or reinterpreting qualities of physical music discovery, not as claims about vinyl in general.

M

Material cue. A physical or visual quality that can invite attention or interpretation, such as texture, wear, artwork, typography, placement or form.

Matching. A contrasting system behaviour in which an object is reduced to a set of values and music is retrieved on that basis. The project treats this as closing the object's interpretation before listening begins.

Meaningful association. The lower threshold for an object-music association: it should feel resonant or enjoyable to the listener, rather than amounting to an arbitrary pairing.

Music recommendation. The selection or ranking of music a system offers a listener. The project critiques recommendation when it resolves an uncertain desire before the listener can take part in interpreting it.

N

Nourishment. A project-specific analytic label for an interaction that sustains a listener's active interpretive work, allowing time, questioning and adjustment instead of settling uncertainty too quickly.

O

Object-based music system. A system in which a listener-selected physical object acts as the starting point for music discovery. The object carries personal history and provides a material anchor for AI interpretation.

Opacity. A lack of visibility into a system's operation. Unlike ambiguity, opacity alone does not create interpretive value; it may simply create confusion.

Output medium. The musical unit an interaction ultimately offers, such as an album or a single song. Its temporal size shapes the amount of physical and cognitive commitment an interface can reasonably ask for.

P

Personal object. An object selected by the listener because it has personal meaning, history or resonance. It is not merely data for the system to decode; it helps establish the situation from which discovery develops.

Placation. A project-specific analytic label for an interaction that resolves uncertainty on the listener's behalf, leaving them mainly to accept or reject a finished result.

Prediction paradigm. The dominant recommender-system logic of using past behaviour or captured context to infer what a listener will want next, often turning an uncertain desire into a usable signal before discovery begins.

Provenance. The traceable origin of a recommendation: the sources, people, communities or reasoning through which it was formed. The Provenance Read-out exploration tested how visible attribution might shape discovery.

Provotype. A provocative prototype whose purpose is to expose and test a design position. Curio's provotypes make different distributions of interpretation and agency tangible.

Q

Query removal. An interaction strategy that reduces or removes the need for a typed, explicit request. Curio replaces the conventional query with a listener's physical choice of object.

R

Record-store journey. The documented process of browsing, selecting and listening in a record store, used as a research activity to identify what physical music discovery can afford.

Research through Design (RtD). A research approach in which designing and making are ways of generating knowledge. Curio uses grounding, exploration and synthesis, with reflective work continuing across the process.

Resolution. The point at which a system closes an uncertain relation by presenting a fixed meaning or recommendation. In Curio, resolution is contrasted with exploration that leaves the relation revisable.

S

Sensitising concept. A provisional idea that directs attention during research without fixing conclusions in advance. Observations from the record-store journey played this role in later explorations.

Short Play. Provotype 4. A speculative single-track interaction in which the listener moves an object through an unlabelled, responsive field of qualitative regions, receiving direction without a prescribed endpoint.

Stimmung. A German term often translated as mood or atmosphere. In Curio, it names a listener's emerging, relational sense of what they may want from music before they can put it clearly into words.

Glossary

T

Temporal granularity. The duration and scale of the unit being chosen. The project distinguishes the high-commitment album from the low-commitment song and designs interaction accordingly.

The Lab. A companion research interface that records images, prompts, AI outputs, recommendations, timings and costs, so the system's behaviour can be reviewed alongside the designer's response.

The Space Between. Provotype 1. A baseline that materialises placation: the system reads an object clearly, turns that reading into a literal music match and leaves little room for the listener to steer or interpret.

Transparency. The degree to which a system exposes its reasoning. The project finds that transparency is not automatically agency: a complete explanation can persuade the listener to accept a result rather than help them steer it.

Traceability. The upper threshold for an association: the listener should be able to understand something of how it emerged, without receiving a fully explicit or definitive machine explanation.

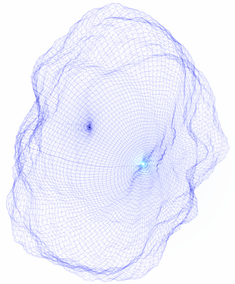
Tuning. Gradual, reversible adjustment of a discovery direction through interaction. It describes navigation and calibration rather than declaring a final mood or accepting a one-off recommendation.

U

Uncertain intent. A listening orientation that is sensed but not yet nameable. It is the condition from which Curio begins, not missing information that AI must immediately resolve.

V

Vibe Map. Exploration 6 and the basis for Short Play: a bounded, unlabelled field generated from qualities of a scanned object. Moving the chosen object through it lets the listener explore the AI's associative logic without a fixed destination.



Introduction

This chapter establishes the project's point of departure, providing a concise overview of the context, motivations, and early process that set the research in motion.

Project Background
Design Brief
Research Questions
Research Process

1.1 Project Background

Left page:

Figure 1.1
Spotify's AI-powered DJ interface and TikTok-style endless scroll, illustrating how artificial intelligence is being used to personalise music recommendations and shape the listener's experience. (Pierce, 2023).

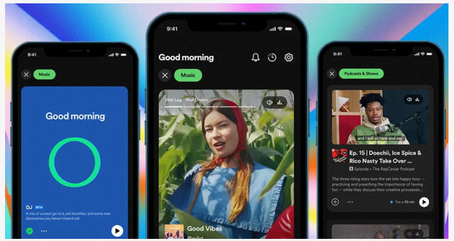
Recorded music is now predominantly accessed through streaming platforms. At the end of 2025, Spotify alone reported 289 million subscribers and a catalogue of more than one hundred million tracks (Spotify, 2026)[Figure 1.2]. Listeners can now reach an extensive range of music within seconds, while artists can connect with audiences without relying exclusively on the institutions that have traditionally controlled distribution. This project takes the expansion of access as its starting condition. Its concern, however, lies less with the size of the catalogue than with the conditions under which its contents become navigable.

To recommend music, a streaming platform must establish some account of what a listener might want. It may ask the listener to provide that account through an explicit input, or infer it from previous listening behaviour [Figure 1.1]. The latter approach can make established behaviour the starting point for future discovery, potentially reinforcing familiar patterns. This reliance on past behaviour is sometimes discussed through the concept of the filter bubble (Pariser, 2011). This project focuses less on the general filter-bubble effect and more on how recommendation systems establish a starting point for discovery.

The present project addresses a more specific problem within this starting-point model: exploratory listening, in which a listener may not yet be able to provide a clear account of what they want. More recent developments have used AI either to broaden the form that input can take or to remove it altogether [Figure 1.1]. Spotify's Prompted Playlist allows listeners to describe what they want in natural language (Spotify, 2026a), while its AI DJ requires no input at all and assembles a session from listening history and inferred context (Spotify, 2023). Research on context-aware and affective recommendation extends the same logic further still, capturing mood, activity, location, or physiological state as the signal that shapes what is offered (Kaminskas & Ricci, 2012; Schedl et al., 2018). Although these implementations differ in what they ask of the listener, they share the underlying assumption that an uncertain desire must first be converted into a usable signal before discovery can proceed.

For many listening situations, this assumption is unproblematic. A listener who wants a particular album, something to work to, or more of what they played yesterday can state that intention, and the system is designed to respond. The assumption becomes more constraining in exploratory listening, where a listener may be reaching towards an atmosphere that is felt before it can be articulated. Existing systems generally treat this ambiguity as missing information to be resolved before discovery can proceed. In converting an indistinct feeling into a predefined signal, they narrow the listener's intention to the categories available to the system and perform the initial interpretation on the listener's behalf. The listener is left to accept or reject the result rather than shape the interpretation that produced it.

The resulting gap is therefore not reducible to recommendation diversity. It concerns the limited capacity of current music-discovery systems to engage with affective intentions that remain ambiguous or preverbal. This thesis investigates that gap through the concept of Stimmung: the felt atmosphere or attunement through which a person encounters the world before that experience is fully expressed in language (Gumbrecht, 2012). The project is concerned with whether an encounter can nourish the listener by developing what they were reaching for, rather than placating them by settling it on their behalf.



1.2 Design Brief

The space. The project explores domestic music discovery for a listener who cannot yet say what they want. This not-knowing is treated as the condition to design within, rather than a gap to close before music can play. Instead of entering a typed request, the listener offers a personal object. AI acts as an interpretive partner, reading the object associatively rather than assigning it to a category or making a recommendation.

Context. The project is situated in the home because this is where exploratory listening is most likely to occur. Listening varies across settings, but domestic listening is more often personally invested and emotionally significant (North & Hargreaves, 2008). It is also where people already keep objects with personal meaning, so the material the system interprets is already present.

User. Anyone who wants to explore music without exact knowledge of what they are looking for. It describes an experience, not a type of person: once the listener knows what they want, search already serves them well.

The constant. One interaction holds across the set. A person places a personal object on a music-playing surface. The system reads the object associatively rather than matching it to a fixed category, then responds with music. Discovery begins with a physical act rather than with an intention typed into a field.

What varies.

1) How the object is read. What the system reads and how it interprets it determine the object's role, shaping the resulting experience.

2)How much of the system's interpretation is revealed. Too little and the object feels decorative. Too much and the encounter turns back into the explicit request the space was defined to avoid.

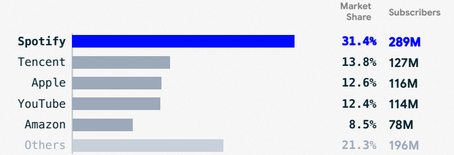
3)The role of the object after it has been read. Its work can end at supplying the input, or it can stay in the listener's hand as the means of moving through what the reading opened up.

4) The music output medium. An album asks for more time and more commitment than a song, so the two invite different kinds of decision, and offering one is not the same design problem as offering the other.

What is retained. Access to the catalogue remains unbounded,; none of the designs creates value by restricting it. Instead, each narrows only what the listener encounters at a given moment: one object, one field or one side of a record. The view is bounded, but access to the catalogue is not.

Out of scope. The project selects from existing music rather than generating new tracks. It operates through a listener's Spotify account in the home, so both the available catalogue and what can be known about a track are determined by Spotify. Other boundaries were not fixed in advance. As a Research through Design project, what Curio was and was not, became clearer through making.

Why Spotify. This report critiques streaming as a way of encountering music, using Spotify as its principal case. As [Figure 1.2] shows, Spotify held 31.4% of the global market's 921.6 million paying subscribers at the end of 2025 (289 million users) and more than double the share of its nearest competitor (MIDiA Research, 2026). Its scale and influence help set patterns adopted across major platforms, so examining Spotify is not simply a study of one company's choices. However, the figure also shows this claim's limits: around two-thirds of subscribers use other services. The critique therefore applies to major platforms to varying degrees, and not to smaller services built on substantially different principles such as Bandcamp or Pandora .



Right page:

Figure 1.2
Global music subscriber market share, Q4 2025. Spotify leads a market of 921.6 million (MIDiA Research, 2026).

1.3 Research Questions

// Overview

§Chapter 2 identifies a common limitation across contextual music recommendation systems. Although these systems use different signals, they follow the same underlying process: captured information is interpreted and translated into a predefined, nameable category before music is recommended. Interpretation is therefore completed by the system rather than developed with the listener. Listeners may accept or reject the recommendation, but they can rarely actively question, further refine, or directly contribute to the interpretation behind it. This project departs from that approach. Rather than classifying feelings and associations in advance, it investigates how music discovery can remain open to this "vibe" that listeners sense but cannot yet clearly identify or express.

Primary.

How might a domestic, object-based music system use associative AI to support music discovery without requiring the listener to declare or resolve their listening state before the encounter?

SQ1, What was lost? // The Record Store

What qualities does record-based discovery, ownership, and use require from listeners and provide to them that streaming has removed? How might these qualities be translated into streaming without reproducing the limitations of physical media?

SQ2, How can this be brought back? // Physical Objects

How can AI-generated associations between a personal object and music remain meaningful and traceable without reducing the object to a single authoritative interpretation?

SQ3, Where should interpretation sit? // AI as a Collaborator

How can interpretive work be distributed among the listener, their object and the AI? And what does each arrangement mean for the listener's agency over where the discovery goes?

// Explanation

The primary question has two connected concerns: the system to be built and the condition of encounter it is designed to preserve. The first rests on an argument made in Chapter 2 (§2.4), where a felt state that is sensed clearly but resists naming is termed a Stimmung (Gumbrecht, 2012). The term is used not as a better model of that state, but as a reason not to model it. A listener's inability to specify what they want therefore becomes the condition within which the design operates, rather than a deficiency to be corrected.

The design commitment that follows is that interpretive work should remain with the listener. This is a stipulation rather than a finding, and it carries no judgement of listeners who prefer an answer supplied for them. What the project contests is the withdrawal of the alternative. Nor does it prescribe a lesser role for the machine. Associative AI is asked to do more interpretive work than a conventional recommender, while stopping short of concluding what the listener feels or why a given track answers it.

SQ1 distinguishes between what record-based practice required from listeners and what it provided to them. Some qualities depended on the listener's effort, attention, or participation. Others were produced by the material and spatial conditions of physical media. A quality falls within the scope of SQ1 if it operated in either of these ways. This distinction is important because an account limited to what listeners previously did would overlook the qualities that were provided by the conditions surrounding the practice. Looking back at earlier forms of music engagement does not imply that they were preferable. It serves a methodological purpose. Qualities that once operated implicitly often become identifiable only after they have been withdrawn. The object of study is therefore not record-based practice itself, but the qualities that became absent in the transition to streaming. The second part of SQ1 establishes the constraint that any proposed response must meet. Physical media supported these qualities alongside the limitations that contributed to their displacement. They therefore cannot be recovered simply by reinstating the conditions through which they originally emerged. For this reason, "translated" is the selected term. The question asks how these qualities might be reformulated so that they can persist under conditions of abundance and immediate access, while recognising that not all of them may be transferable.

SQ2 turns on three terms that require explicit definition. Associations describes what the AI is asked to produce: a proposed connection between a personal object and music that does not present itself as the correct or definitive interpretation. This differs from matching, which reduces the object to a set of values and retrieves music on that basis, effectively closing its interpretation before listening begins. Meaningful establishes the lower threshold and refers to the listener's experience of the association. The connection between the object and the music must feel resonant or enjoyable. This does not require the listener to agree with a fixed explanation of the object, but it does require the association to offer more than an arbitrary pairing. Traceable establishes the upper threshold. The listener should be able to understand something of how the association emerged, without the system making its reasoning fully explicit or presenting it as definitive. A single authoritative interpretation would exceed this threshold because the meaning of a personal object belongs partly to the listener and cannot be inferred from its visible features alone. Any machine interpretation must therefore remain incomplete. This incompleteness creates space for the listener's own reading, reflecting Gaver et al.'s (2003) account of ambiguity as a design resource. The question can consequently fail in either direction. Too much explanation fixes the object's meaning in advance and turns it into a conventional query; too little leaves the association opaque and risks making the object merely decorative.

SQ3 asks how interpretive work might be distributed across the listener, the personal object, and the AI. It does not assume that there is one correct distribution. The qualities identified through SQ1 may conflict both with the strengths of streaming and with one another. A single design would have to resolve these tensions in one particular way, thereby presenting a degree of certainty that the research does not support. Developing a set of designs allows each arrangement to assign interpretive work differently and keeps these tensions open for examination. This approach corresponds to Gaver and Bowers' (2012) concept of the annotated portfolio, in which design knowledge emerges through the relationships and contrasts between a group of artefacts rather than from any single artefact in isolation. Comparing the arrangements makes it possible to examine, for example, what changes when greater interpretive responsibility is given to the listener, the object, or the AI. The second part of SQ3 considers what each distribution means for the listener's agency. Here, agency refers specifically to the listener's control over where the process of music discovery goes next, rather than to a general feeling of being in control. Each arrangement enables some possible directions while limiting others. These differences become more visible through comparison, because what one arrangement withholds can be recognised more clearly alongside an arrangement that makes it available.

Closing Section

These questions do not ask what listeners in general prefer, and nothing in this report claims to answer that. The listener is the one specified in §1.2. The evidence is sustained first-person use and close work with a one other participant, so the claims concern what the approach makes possible rather than what a population would choose (Gaver, 2012; Neustaedter & Sengers, 2012).

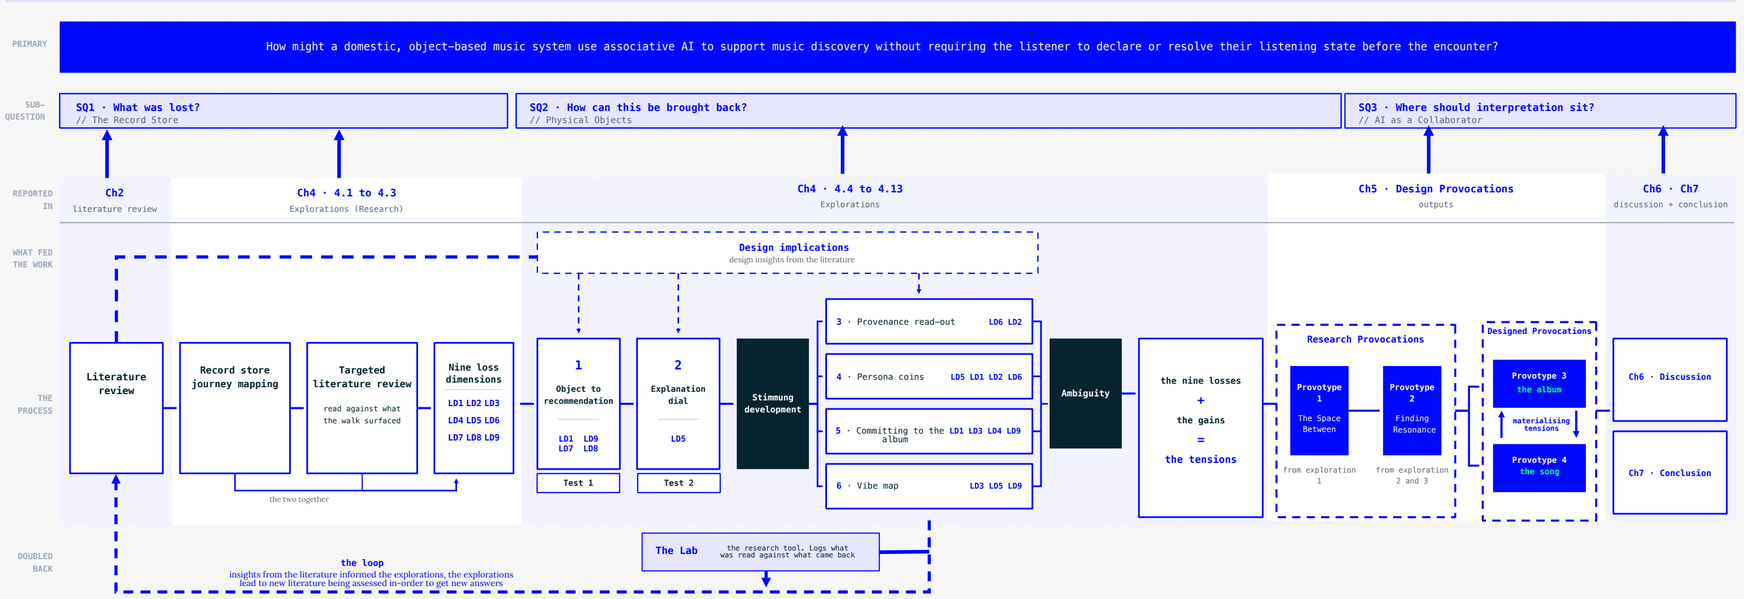
1.4 Research Process

Left page:
Figure 1.3
Research Process Diagram

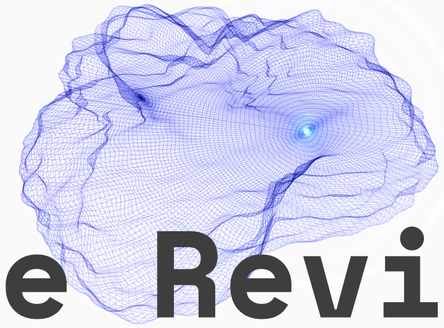
The project followed a Research through Design process, discussed in greater detail in §Chapter 3. As shown in [Figure 1.3], it began with a literature review examining what streaming had displaced, why those losses mattered and how design might respond. The design implications developed at the end of each literature-review section provided initial directions for design explorations, considered alongside the dimensions of loss identified through the documented record store journey mapping.

Research and making then proceeded iteratively. Design provocations generated questions that led back to the literature, while formal testing was undertaken only when existing research could not provide an adequate answer. This approach made efficient use of the available time and reserved testing for producing knowledge that could not be gained through reading alone. Each exploration therefore served both as a potential design response and as a means of revealing further questions.

Reflection on later provocations also reshaped the interpretation and development of earlier explorations. The literature, design implications, loss dimensions, provocations, and reflections therefore formed a cyclical, mutually informing system rather than a linear sequence. In [Figure 1.3], dashed arrows indicate these returns, while the main line shows the approximate order in which material entered the project. The resulting provotypes represent different stages and perspectives within the exploration of the design space, rather than steps towards a single final solution. The first two provocations materialised research concepts explored in my explorations, and the final two act as materialisations of design tensions that revealed themselves from the insights gained through the exploration phase.



Chapter 2



Literature Review

This chapter traces the evolution of the research that shaped Curio, showing how the literature progressively refined the project's scope and informed its first concept direction

Each section closes with a set of design implications: short statements of what the literature in that section demands of future design work, these are written as constraints rather than findings.

Limits of Streaming Discovery
AI in Music Recommendation
Existing Landscape of Query Removal
Stimmung and Uncertain Intent
Discovery in the Record Store
Physical Objects as an Interfaces
Objects as Carriers of Personal Meaning
Prediction to Interpretation
The Research Gap

2.1 Limits of Streaming Discovery

Music streaming has transformed access to recorded music. When listeners want to discover unfamiliar music, major streaming platforms generally direct them along one of two routes. Listeners can name an artist, song or genre and allow the system to work outward from that starting point, or they can accept a selection assembled from their previous listening behaviour. The first route assumes that listeners can articulate a useful preference before exploration begins. The second requires no explicit input, but asks the platform to infer a preference from past behaviour. In both cases, discovery begins from a relatively fixed account of what the listener is expected to want: either one stated by the listener in advance or one predicted by the platform. Preference therefore becomes a starting condition of discovery, rather than something that might develop through the encounter itself.

Platforms such as Spotify shape how listeners are understood. Prey (2018) argues that personalisation does not reveal pre-existing listeners; it produces them through categories designed around advertisers' target markets. The platform's vocabulary for listening situations therefore comes before the listener's own. Popularity bias adds a further constraint, leading recommender systems to favour already well-known content over less familiar parts of the catalogue (Celma, 2010). Recent interface features, including mood-based browsing, 'vibe' search and natural-language prompts, make some preferences easier to express. Spotify's Prompted Playlist, for example, asks listeners to describe what they want in their own words (Spotify, 2026a), while Discover Weekly, which was previously assembled without direct input, now includes genre controls at the top of the playlist (Spotify, 2026a). These features give listeners more control over how they describe a request, but they do not remove the need to define a direction before encountering the available music. They may also assume that the listener's understanding of a term such as 'excited' corresponds closely enough to the way the platform has classified that mood. Agency is therefore distributed unevenly: the listener must either define a preference before encountering the available music or allow the platform to define one from past behaviour. Once the system begins selecting and sequencing tracks, alternative directions and opportunities to refine the choice may become less prominent, leaving the platform with greater control over how the exploration develops. The issue is therefore not only whether a platform can interpret an input accurately, but also who gets to shape the direction of discovery and when.

That starting-point model fits poorly with how music is often chosen in everyday life. An exploratory study of everyday listening reported that listening is frequently reactive and situated, meaning that music is selected in response to the circumstances of the moment rather than through a deliberate choice made before the listening situation unfolds (Sloboda et al., 2001). This does not mean that listeners make no choices. Instead, it suggests that their choices may arise from immediate factors such as activity, mood, location or social setting. A similar pattern was reported by Greb et al. (2019), who found that music selected in everyday life was influenced more strongly by situational characteristics than by stable individual differences. A related modelling study demonstrated the value of combining listening history with situational context rather than treating past behaviour as a single, undifferentiated preference profile (Pichl & Zangerle, 2021). These studies do not themselves claim that preferences develop through interaction with music; that is a further inference made in this review. Taken together, however, they suggest that a preference specified or inferred before listening may be an incomplete guide to what a person will value in a particular situation. Some preferences may become clearer as the listener encounters and compares music, rather than existing as fully formed choices waiting to be retrieved. The problem therefore concerns not only the accuracy of recommendation, but also the interaction through which listeners encounter, assess and refine possible choices.These findings therefore have implications for the way recommendations are delivered. When recommendations arrive through a continuous or 'lean-back' stream, the listener may take on a more passive role. This can be convenient, but it may also provide fewer opportunities to attend closely to differences between tracks, compare alternatives or reconsider what should come next. If listeners discover what they like by comparing tracks, a continuous stream may give them fewer opportunities to make those comparisons and adjust their choices.

Streaming platforms have been described as being organised around the collection of behavioural data (Eriksson et al., 2019). One possible consequence is that when recommendations are generated from established patterns, and the resulting listening is then recorded as further evidence of those patterns, personalisation may become self-reinforcing. Regular behaviour is easier to capture and model than temporary, uncertain or contradictory intentions. This is an interpretation of that platform logic rather than a direct finding of the work itself. There is some evidence that the way listeners access music is related to the range of music they hear, Anderson et al. (2020) found that user-directed listening, such as searching for music or using personally created playlists, was more diverse than listening through algorithmically programmed sources for most users. This finding does not show that algorithmic recommendations caused narrower listening, but it suggests that the mode of recommendation may influence which parts of the catalogue listeners encounter.

These dynamics matter beyond diversity metrics. Music has been described as a "technology of self": a resource through which people construct identity and regulate emotion in everyday life (DeNora, 2000). Choosing music can be part of that process. When a platform takes over more of the selection, it reduces the occasions on which listeners actively choose and reflect on what they hear. Passive listening is not inherently a problem. It can be useful, pleasurable and appropriate to the situation, and this review does not argue that it should be replaced. The concern is more specific: streaming interfaces make passive listening easy to sustain, while exploratory and meaning-making forms of listening can require more deliberate effort.

This matters because the relevance of discovery is not stable across situations. The same listener may seek unfamiliar music in one situation, familiar music in another. These shifts do not indicate incoherent taste. They show that people draw on different parts of a broad musical repertoire for different purposes and in different contexts. Predictive systems built mainly from past listening may struggle to respond to this variation. By treating repeated behaviour as evidence of a stable preference, they risk typecasting listeners and continuing to recommend from the same limited part of their taste. This may work well when a listener wants continuity, but it is less useful when they want to move beyond established patterns or explore another part of their musical interests.

Rather than dividing listeners into passive and exploratory types, this review focuses on changing listening situations. The same listener may want to explore unfamiliar music at one moment and listen to familiar music at another. The kind of support they need from a recommendation system may therefore also change with the situation. These moments cannot be assessed through predictive accuracy alone. Studying discovery on a major streaming platform, (Garcia-Gathright et al. 2018) found that what counts as a satisfying recommendation depends on the listener's goal, so identical interaction signals carry different meaning from one listening context to the next.. In music recommendation, novelty, diversity and serendipity can be increased while limiting the resulting loss of accuracy (Zhang et al., 2012). Serendipity has, however, been defined and measured inconsistently, making approaches difficult to compare (Kotkov et al., 2016). Discovery success should therefore be understood both as an outcome and as a process: the extent to which listeners can steer, question and therefore revise the direction of their exploration. It concerns whether AI-supported discovery can respond to changing situations and purposes, make less visible parts of the catalogue available for exploration, and allow listeners to refine their direction without requiring them to define either their preferences or their listening mode fully in advance.

Design implications.

D1 Allow the listener's sense of what they want to develop through listening itself, rather than requiring them to name an artist, genre, mood or activity before the first track plays.

D2 Respond to the listener's present circumstances and not only to their accumulated history, using situational context where it is available and can be made visible to the listener, since the same person may want lean-back listening in one situation and active exploration in another.

D3 Keep the listener's control available throughout the session rather than only at its start, so that they can pause, redirect, compare alternative directions or revise the exploration after playback has begun.

D4 Treat each skip, save or repeated play as a provisional signal rather than as evidence of durable taste, and give the listener a way to correct or reverse the personalisation that results.

D5 Support familiar listening while also making less visible parts of the catalogue reachable, expanding the range of what the listener encounters without resorting to random novelty.

2.2 AI in Music Recommendation

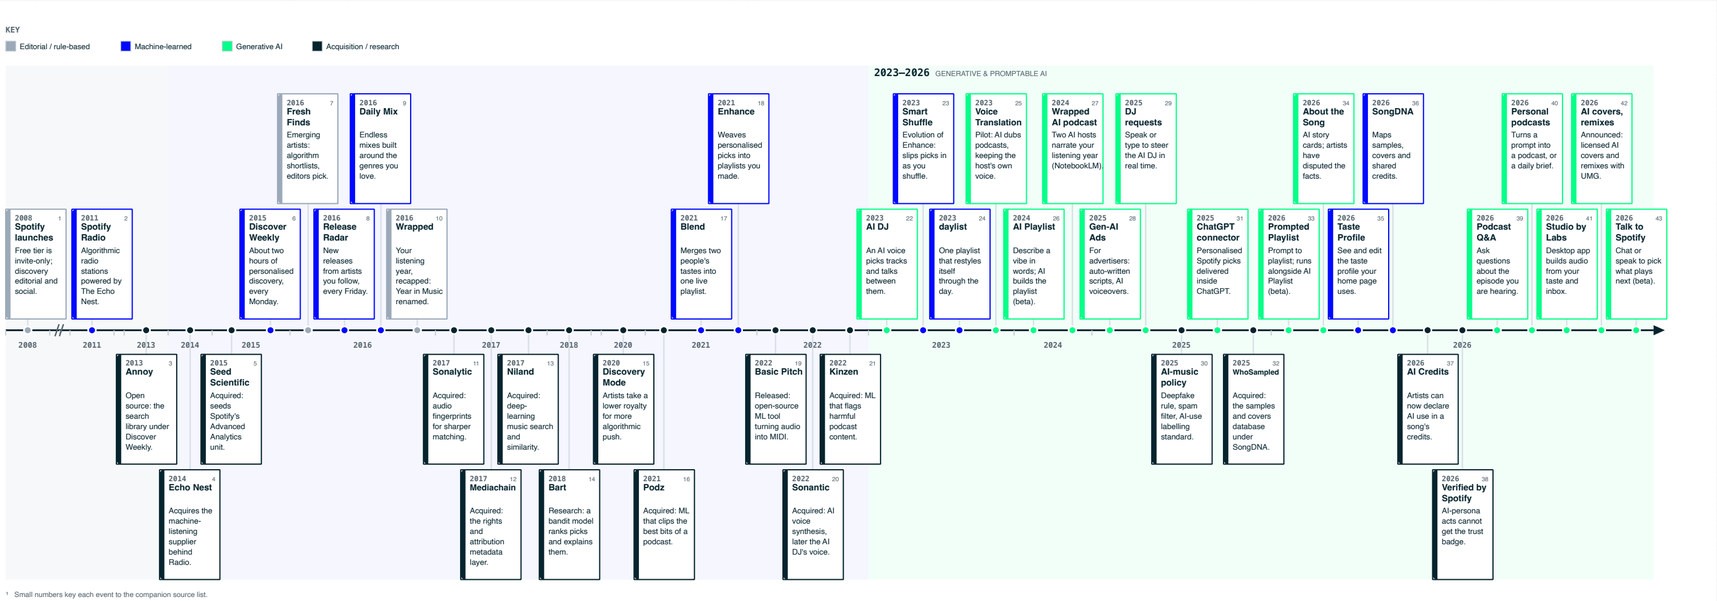
Right page:
Figure 2.1

Spotify's shift from editorial curation to machine-learned personalisation to generative AI, 2008 to 2026, with the acquisitions and research beneath it. The axis is chronological Numbers key to the source list that can be found in [Appendix A]

To understand what a different use of AI might involve, it is useful first to examine how AI is used in recommendation systems today. Collaborative filtering infers listeners' preferences from the behaviour of users with similar tastes; content-based methods match recommendations to audio and metadata features; hybrid systems combine these approaches; and context-aware systems incorporate situational factors such as mood, activity, location, and time. More recently, deep neural networks and increasingly generative and large language models have extended these approaches, using more powerful forms of modelling to personalise music recommendation (Epure et al., 2025). Different as they are, they optimise for one thing: predicting the next play accurately enough to hold attention. They all rely on prediction, and implicitly this means working from the past, from how a listener and others like them have already acted.

Every paradigm reads the past rather than the present moment. Under an engagement objective, the safest way to hold attention is to return more of what already works, so deployed prediction tends toward the familiar. It is easy to overstate this, so it is worth stating carefully: that tendency comes from the goal, not from prediction. The serendipity and beyond-accuracy work noted in §2.1 shows prediction can be tuned for novelty and surprise instead (McNeeet al., 2006; Kotkov, et al., 2016); the narrowing reflects commercial pressure rather than a technical limitation. Nevertheless, the prediction paradigm is not built to ask what if AI's capacity for interpretation were turned on something present, a physical object in front of the listener, rather than on their history? This is taken up in §2.8, once the loss it might answer to has been set out. For now, this approach uses past behaviour to guide the next play, and it is worth asking what is gained and lost by following it.

It offers convenience and accuracy, but the sections that follow examine what this costs. As recommendation grows more frictionless and more accurate, the physical and mental effort involved in finding and choosing music goes with it. What has already been lost turned this project from a critique of recommendation into a search for what a different use of AI might restore.



[1] Small numbers key each event to the companion source list.

2.2 AI in Music Recommendation

Design implications.

D6 Base the recommendation on the physical artefact present in front of the system rather than on the listener's accumulated listening history.

D7 Do not treat next-play accuracy or sustained attention as the sole objective, since holding a listener's attention and opening unfamiliar territory are distinct design targets that can pull against each other.

D8 Use listening history as one source of context that informs the system's response, rather than as the factor that determines what the listener is able to encounter.

D9 Preserve the listener's effort and agency in finding music, leaving room for choosing, comparing, interpreting and dwelling rather than reducing discovery to passive delivery.

2.3 Existing Landscape of Query Removal

The difficulty of asking a listener to name what they want is not unique to music systems. Information-retrieval research has long recognised that people may feel a gap in their knowledge without being able to formulate a complete request. Recommendation research similarly treats context as something that may need to be modelled rather than assumed (Belkin, 1980; Adomavicius & Tuzhilin, 2005).

In this review, an under-specified listening direction describes a situation in which a listener senses that they want music to take them somewhere, without being able to identify that direction as a specific mood, genre or typed request. Everyday listening research shows that music selection is strongly shaped by situational factors, although their influence varies between listeners (Greb et al., 2019). Context-aware music research therefore considers factors such as activity, location, time, weather, social setting and physiological condition (Kaminskas & Ricci, 2012; Murciego et al., 2021).

The following categories form an analytical framework rather than an exhaustive taxonomy of contextual music systems. They identify what supplies the listening direction, who interprets it and when that interpretation becomes the basis for song retrieval. The categories are not mutually exclusive; they might combine more than one context; however, the framework identifies the dominant role that each input plays in making a listening direction actionable.

Declared context

Some systems ask the listener to describe their current situation through a set of predefined controls. Contextual values such as activity or location are then converted into a query, ranking condition or playlist constraint.

ContextPlay allows listeners to control both musical preferences and contextual characteristics, including mood, weather, time, location and social setting. In a study with 114 participants, additional control over context improved perceived recommendation quality without increasing cognitive load (Jin et al., 2019).

The animated mood-picture prototype developed by Lehtiniemi and Holm (2012) allows listeners to select from animated pictures associated with particular moods or activities. The system then provides music recommendations from the genres associated with the selected picture. This preserves listener control but limits it to the moods and situations represented by the interface. The listener must therefore translate an underspecified listening direction into a predefined mood, activity, or scene before the system can respond.

These systems preserve a degree of listener control, but that control is limited to the categories made available by the interface. From the perspective used in this review, they place the initial work of resolving uncertainty with the listener. Before the system can respond, the listener must translate an under-specified listening direction into a mood, activity or other contextual value.

2.3 Existing Landscape of Query Removal

Sensed and inferred listener context

Other systems reduce the need for explicit declaration by estimating something about the listener from physiological, behavioural or other sensed signals. The important distinction is not the particular sensor, but what the system does with the signal: it uses the measurement to make an estimate about the listener's situation before selecting music.

The Affective Music Player learns an individual listener's physiological responses to music and selects existing songs intended to move them towards a specified goal state (Janssen et al., 2012). It is therefore an example of physiological goal-targeting rather than simply detecting a listener's current mood.

Wearable-sensor systems similarly use physical-world information to predict music preferences. In Towards Ubiquitous Personalised Music Recommendation with Smart Bracelets, data collected from smart bracelets was combined with participants' reports of mood and musical preference during a one-week field study with 30 participants. The authors used these data in their Multi-task Ubiquitous Music Recommendation model to improve personalised recommendations (Li et al., 2022).

These examples reduce reliance on explicit self-report in different ways. The Affective Music Player learns how music affects an individual and selects tracks towards a physiological goal, while the smart-bracelet system uses physical-world signals, mood reports and environmental information to predict context-sensitive music preference (Janssen et al., 2012; Li et al., 2022). They demonstrate context-sensitive selection and mood regulation more directly than open-ended discovery of unfamiliar music.

These systems reduce the burden of self-report, but shift interpretive responsibility to the system. The listener contributes bodily or behavioural signals, while the model turns them into a usable estimate, for example, that the listener is active, tired or in a particular mood. In the reviewed examples, this estimate is then used to produce a ranked playlist or targeted affective intervention.

External environment and image prompts

Context may also be related to the listener's surroundings rather than to a direct account of their internal state. Location-aware systems use a person's geographic position, while venue-aware systems focus on the type or character of the place they are in, such as a café or park.

Location-aware recommendation has been shown to improve satisfaction compared with context-free personalisation (Kaminskas et al., 2013). Venue-aware systems model relationships between venue types and musical content through a shared semantic space (Cheng & Shen, 2016).

Location and venue systems can also involve sensing or inference. For example, a phone may use GPS to identify a person's location, after which the system may interpret that location as a venue type. The distinction used here is therefore analytical: sensed or inferred context describes how the system obtains and interprets information, while environmental context describes what that information refers to, the surrounding place or situation.

Image-based systems provide a related but distinct form of external prompt. These systems use visual material to represent a scene, atmosphere or situation without requiring the listener to enter a textual query. Image-based retrieval systems use different forms of representation. Some systems extract mood or thematic tags from images (Park et al., 2023), whereas others represent images and music within a common space of meaning or emotion (Hsia et al., 2018; Zhao et al., 2020).

Hsia et al. (2018) provide a technical precedent for image-to-song retrieval and a comparison of text, image and hybrid-systems for music search found that participants reported better usability and user experience with image-based search than with text alone (Park et al., 2024). This provides evidence about usability, but not about long-term discovery or the development of personal meaning.

Location, venue and image systems are grouped here because they all use something outside the listener's verbal description as the starting point for recommendation: where the listener is or what they are looking at.

In each case, the system converts the external prompt into information it can use to find or rank songs. A place may be represented through tags, venue categories or semantic associations. An image may be represented through visible features, mood labels or a position in a learned image–music space. The system then uses this representation to determine which existing tracks are retrieved.

Among the reviewed approaches, image-based prompting provides the closest technical precedent for using external material to prompt the retrieval of existing music. For the purposes of this review, the limitation is that these systems model the image through features available to the system rather than through the owner's personal relationship with it. They use visible cues as indirect clues about what music might fit. The external prompt therefore replaces a textual query, while the system determines the representation used for the initial retrieval.

The relevant distinction is not simply between an image and a physical object. It is between an external input treated primarily as data and an object chosen as a personally meaningful act.

Discovery without a fixed query

A different response is to change the interaction rather than to sense more context. Exploratory music interfaces reduce the need for a fully formed request by allowing listeners to browse, compare and encounter music through interaction.

Musicream presents a flow of existing musical pieces that listeners can select, arrange and revisit. It supports unexpected encounters without requiring a specific artist or genre query, but its browsing remains structured by audio similarity and query-by-example interaction (Goto & Goto, 2009). EmoMTB similarly works with an existing collection of tracks. It provides a non-linear music landscape in which listeners can move between clusters, explore unfamiliar areas and leave their established musical comfort zone. Emotion-aware recommendations provide starting points for exploration rather than determining the entire listening path (Melchiorre et al., 2023).

These are not music-making systems. They do not primarily support composing or generating music; they change how existing tracks are browsed, organised and encountered. EmoMTB is therefore evidence of an exploratory interface, not evidence that open-ended interpretation itself succeeds.

These systems preserve more active participation than conventional ranked lists: listeners can steer discovery through movement, selection and comparison. However, the available paths remain bounded by similarity structures, genre groupings, mood labels or optimisation objectives. Exploration is therefore open-ended only within the relationships between tracks that the system makes computationally available (Goto & Goto, 2009; Melchiorre et al., 2023). As Wang et al. (2014) illustrate, recommender-system approaches to exploration may also frame discovery in relation to improving future prediction or maximising long-term reward.

2.3 Existing Landscape of Query Removal

Right page:
Figure 2.2

A comparison table for the content covered in §2.3

[Figure 2.2] compares these approaches according to how they initially treat an under-specified listening direction.

Approach	Listener contribution	System operation	How uncertainty is initially resolved
Declared context	Names a mood, activity or situation	Converts it into a query or ranking condition	The listener resolves it for the initial recommendation.
Sensed or inferred listener context	Provides bodily, behavioural or physiological signals	Classifies or estimates a current situation	The system resolves it for the initial recommendation.
External environment or image	Is situated within or supplies a place, scene or image	Maps it to semantic, affective or similarity-based features	The system operationalises a representation before the initial retrieval.
Encounter-based discovery	Browses, selects and compares	Sequences or exposes possibilities for exploration	Uncertainty remains open longer, but is bounded by computational relationships.

These approaches differ in who resolves an under-specified listening direction and when. Declared systems require the listener to resolve it before recommendation. Sensed systems resolve it through physiological, behavioural or environmental signals, while image- and environment-based systems translate it into semantic or affective similarity. These categories can overlap: for example, a system may sense a person's location through GPS and then use the resulting venue representation as an environmental prompt. Encounter-based interfaces leave uncertainty open for longer, but still organise discovery through similarity, genre, mood or prediction.

The central distinction is therefore not simply how context is captured, but whether the system treats its representation as sufficient for initial selection or keeps it provisional and revisable. In most of the recommendation systems reviewed here, a representation of the listener, situation or input is established before the first recommendation is ranked or retrieved. Exploratory interfaces are not necessarily query-driven, but they usually still rely on a computational structure that limits which relations can be encountered.

This points to a more specific gap. The literature contains separate examples of contextual recommendation (Jin et al., 2019; Kaminskas et al., 2013), image-to-music retrieval (Hsia et al., 2018; Park et al., 2023), tangible music interfaces (Frohlich et al., 2004; Kwon & Lee, 2018; Escobar et al., 2023; Jang et al., 2022), and exploratory browsing (Goto & Goto, 2009; Melchiorre et al., 2023). Within the literature reviewed here, what is less developed is their combination, by using a personally meaningful physical object to prompt associative interpretation and open-ended discovery of existing music.

The under-explored area is therefore not whether objects or images can be linked to music. It is whether associative interpretation can support discovery while leaving the relationship between the object and the music open to the listener. The resulting design question is not only how a system captures context, but how material interaction might preserve the listener's authority to interpret the relationship between a personally meaningful object and the music it returns.

Design implication.

D10 Treat whatever the system infers about the listener as a provisional working representation rather than a settled basis for selection, so that it remains open to revision once listening has begun and does not permanently bound what music the listener can be offered.

2.4 Stimmung and Uncertain Intent

Stimmung is a German term commonly translated as “mood” or “atmosphere,” although no single English term captures its full meaning. It can describe both a person’s disposition and the atmosphere of a room or situation. The concept therefore locates experience neither solely within the individual nor solely within the environment, but in the relation between them (Thonhauser, 2021). For this project, Stimmung names a listener’s developing sense of what they may want from music before they can describe it clearly.

The term’s musical etymology reinforces this relational sense. Stimmung derives from Stimme (“voice”) and stimmen (“to tune” or “to bring into agreement”) (Thonhauser, 2021), therefore implying resonance. This makes the concept relevant to music discovery because a listener may sense what they are seeking without being able to specify the specific mood or genre. Rather than treating this uncertainty as an incomplete query, Stimmung frames it as a meaningful condition through which experience unfolds.

Gumbrecht (2012) argues that literature and art do more than represent meaning: they can make an atmosphere felt. He describes Stimmungen as gradual variations that we can sense but may not be able to describe precisely or put into words. This matters for music interfaces because a listener may feel that a certain musical direction is right without being able to name the mood they want. Mood labels may still be useful, but requiring one at the beginning can turn an emerging experience into a fixed category.

The relational character of Stimmung also challenges the assumption that mood is a fixed feeling located entirely inside the listener. In this project, meaning develops through the interaction between listener, object, sound, and environment. The object is therefore not merely a means of eliciting or recording a pre-existing preference; it helps the listener develop and recognise what they want from the music.

Gumbrecht’s point is that an atmosphere can be felt without being clearly named. Heidegger’s (1927/1962) account of attunement (Befindlichkeit) takes this idea further. He treats mood as something that shapes how a person is already situated in the world, before they consciously identify or explain how they feel. From this perspective, not knowing exactly what one feels is not simply a gap in the listener’s knowledge that a system should correct. It is part of the experience itself. Designing with uncertainty is therefore defensible: the system can support musical discovery without first identifying the listener’s mood.

A corresponding design argument appears in (Gaver et al.’s 2003) account of ambiguity as a resource for design. Their account treats interpretation as part of the value of an encounter rather than as a problem to eliminate. Interaction can therefore support gradual adjustment rather than demand immediate specification. Moving an object and hearing the musical output change becomes an act of tuning: the listener retains control through navigation and calibration rather than explicit declaration.

Taken together, these accounts establish Stimmung as the project’s conceptual basis. The project is not concerned with detecting emotion or predicting the most appropriate track. Its scope is narrower; it investigates how interaction can support a listener who cannot yet specify what they want while preserving interpretive control. Uncertain listening intent is treated not as a query to be solved, but as an atmosphere into which the listener can gradually tune.

This focus also explains the move to the record store. The record store is relevant not simply as a nostalgic alternative to streaming, but as a setting in which discovery is distributed across objects, spatial arrangement, social interaction, and chance. These conditions make it possible to explore an emerging preference without requiring it to be fully articulated in advance.

These two lines of thought converge on the position that guides the remainder of this report. One treats interpretation as part of the value of an encounter, while the other treats atmosphere as perceptible but not fully nameable (Gaver et al., 2003; Gumbrecht, 2012). Together, they establish the contrast this project will examine: whether recommender systems tend to resolve an unnameable desire into a recommendation, a process I call i will define as “playcation”, or help the listener explore that desire through interaction, which I call “nourishment”.

These terms are my analytic labels, not empirical findings or claims that one mode is inherently superior. They identify where interpretive work is placed; whether it is settled by the system or sustained through interaction.

Design implications.

D11 Treat the listener’s not yet knowing what they want as the condition the interaction begins from, rather than as a gap the system must close before it can respond.

D12 Allow the object, the music and the surrounding environment to take part in forming the listener’s preference, rather than treating that preference as something already fixed inside the listener and waiting to be retrieved.

D13 Support gradual and reversible tuning of the result, so that the listener can adjust their direction over several moves rather than receiving a single recommendation to accept or reject.

D14 Give the listener feedback on how their input has been read without presenting that reading as the one correct interpretation, so that the ambiguity remains something they can work with.

D15 Keep interpretive authority with the listener and let discovery arise through material, spatial, social and chance encounters rather than through a single computed judgement.

2.5 Discovery in the Record Store

The previous section introduced Stimmung as a quality of attunement that a listener enters into rather than deliberately selects. This attunement is not produced by the listener alone; it is also shaped by the conditions through which music is encountered. Listening technologies organise attention by structuring the relationship between listener, sound and medium (Sterne, 2003). Different systems of music access therefore support different forms of attention.

This section uses the record store not as a model of unmediated discovery, but as a model of how physical stores shape discovery differently from streaming platforms. Its significance is that an object can solicit attention before a listener has translated a felt orientation into a clear request. The difference concerns when specification happens, not whether mediation exists: encounter does not always precede specification, but the record store makes that sequence possible. The store matters here because it combines material handling, spatial browsing and incomplete information, not because it is free from classification, commerce or cultural convention. A visitor can respond to a record before articulating a clear intention or preference. I refer to this condition as encounter before a clear request, a term that carries through the discussion that follows.

Research on record collecting explains why encounters of this kind become personally significant. (Straw, 1997) treats collecting as a practice of connoisseurship in which acquisition and identity formation are bound together, and specifically as one through which masculine expertise is performed. Moving through shelves, examining objects and making selections does not simply express an existing taste; it contributes to its development. Collectors form lasting relationships with music as an owned and handled object (Eisenberg, 1987), while music more broadly serves as a resource through which people regulate mood, identity and everyday experience (DeNora, 2000).

These sources do not establish that record-store browsing produces attunement. They instead identify a mechanism relevant to this project. When listeners explore without first stabilising their intentions, preferences and associations emerge through the encounter itself. Here, attunement refers not to a feeling that has already been clearly identified, but to this receptive and self-directed orientation as it develops through exploration.

The physical and visual features of the store support this process in two ways:

- Evidential: a worn sleeve indicates handling and circulation, while a thinned-out section suggests demand or popularity rather than qualities of the music itself.
- Representational: artwork and typography are designed to signify the record's contents, and a browser reads them as claims about how it will sound.

These signals are culturally coded rather than neutral or universally understood, and interpreting them depends on the knowledge brought by the interpreter. Responding to them nevertheless does not require a listener to formulate an explicit request. A browser is drawn towards an object without being able to state precisely why.

Record sleeves support discovery partly because they provide a visible representation of the music they contain. However, objects can also facilitate discovery without directly representing musical content. Their size, weight, texture, manipulability and resistance invite actions such as picking up, rotating, comparing, retaining or rejecting. These actions do not communicate what the music is; rather, they create the time and space in which listeners can notice the object and develop a response to it. Meaning is therefore not fixed within the object but emerges through the interaction between the object and listener. Record stores are not language-free environments; genre dividers, artist names, price labels and staff recommendations all structure the browsing experience. What differs is the position of language within this sequence. Visitors may first notice, handle or respond to a record before consulting its name, genre or broader context; the encounter precedes the description of what the user is interacting with.

Streaming platforms, however, reorganise this sequence, structuring discovery through searchable metadata, explicit categories and inferred preferences. A listener enters a song, artist, genre. These systems still produce unexpected encounters, and streaming has not eliminated explorative browsing. However, the catalogue reaches the listener already framed by classifications and predictions about the listener and the intended listening situation. The contrast is not between an open physical environment and a predetermined digital one. Record stores also shape attention through shelf organisation, commercial selection and cultural convention. The difference is when that structure becomes active. In the store, some encounters happen before the listener's intention has been defined or inferred. On a streaming platform, description or prediction frames the encounter first.

Vinyl's continued appeal indicates that convenience is not the only value by which music technologies are judged. In the United States, vinyl revenue reached US$1.4 billion in 2024, marking an eighteenth consecutive year of growth (Recording Industry Association of America, 2025). This figure does not reveal why people buy records, nor does it demonstrate a desire for record-store browsing. Vinyl is purchased for ownership, collectability, visual display, nostalgia, perceived sound quality and support for an artist (Fernandez & Beverland, 2019). Qualitative work locates part of this appeal in material interaction. (Bartmanski and Woodward 2015) place vinyl's value in the practices of handling, while (Magaudda, 2011) reads the return to physical music practices as a response to their removal from digital consumption. Handling, ritual and ownership remain meaningful even after they stop being necessary for access. The revival therefore supports a limited claim that efficiency is one criterion among several, It does not show that friction is valuable in itself.

Slow technology clarifies what kind of value is at stake. (Hallnäs and Redström, 2001) argue for forms of design that make things appear and be used more slowly in order to open time for reflection, treating presence rather than efficiency as the goal. (Redström, 2006) adds that artefacts can leave room for interpretation instead of defining their use in advance. Applied to browsing, this distinction separates productive slowness from mere obstruction. Moving through a crate, examining one object at a time and making provisional judgements sustains attention without prescribing a result. This process is slower than entering a search query, but its value does not lie in duration alone. It lies in the ambiguity, bodily involvement and interpretive space that the interaction preserves.

Encounter-based discovery is not the same as random choice. For an encounter to become meaningful, the listener's response must accumulate across actions and affect what becomes possible next. The design challenge is therefore to provide continuity without turning the object into a code that dictates the music. Ambiguity must leave room for interpretation while still giving the listener something to follow. The case therefore identifies four conditions for the design work that follows: material cues, self-paced exploration, incomplete information and encounter before a clear request.

None of these conditions requires vinyl, and materiality alone is not sufficient. A record sleeve does more than give the hand something to hold: its visual qualities are also designed to signify or evoke the music inside. Can encounter-based discovery be supported by material cues that do not represent the musical content, or does attunement depend on a correspondence between the cue and the music encountered? The analogy has a precise limit, and it lies at the point where the project is most original. Browsing in a record store is pre-verbal, but the object encountered still has a culturally established relation to music. Curio, by contrast, uses ordinary household objects with no designed relation to any particular music. Its objects therefore cannot tell the listener what the music is. Their role is instead to provide a material situation in which a response can be formed and carried forward.

2.5 Discovery in the Record Store

The record store therefore supports the claim that encounter-based, pre-verbal browsing works, but not the harder claim that interpreting a non-musical object can drive music discovery. That gap is the project's central proposition rather than a weakness in the analogy. Removing the built-in link between object and music keeps the listener's interpretation open without turning the object into a code that can be decoded into musical content. The listener cannot simply infer the music from the object, so the listener still decides what the connection means. Whether a reading without a built-in musical relation can still feel meaningful is one of the questions this project looks to explore.

This question sets the scope of the design work. Rather than recreating the record store, the project asks whether a material, self-directed interaction can support receptive discovery without requiring the listener to begin from an known input.

Design implication.

D16 Preserve the physical actions that slow the interaction enough for attention to form, allowing the listener to linger over a result, compare it with another, backtrack or change direction at their own pace.

D17 Let the listener encounter the object and hear the music before any artist, genre or mood description frames it, and make that information available afterwards so it extends the encounter rather than preceding it.

D18 Treat the setting as part of the discovery mechanism, since encounter-based discovery depends on a place that supplies objects within reach, time to dwell and existing personal associations, which in this project is the home.

D19 Let an object's handling, weight, texture and manipulation contribute to how the listener interprets what they hear, without encoding a fixed correspondence between particular objects and particular music.

D20 Let the listener's interpretations accumulate and shape the musical possibilities that follow, so that ambiguity opens a path through the catalogue rather than producing unrelated results.



Right page:
Figure 2.3

Image of the authors local record store.

2.6 Physical Objects as an Interface

The project begins with the idea that an ordinary object can serve as a computational interface without losing its physical or cultural meaning. Ishii and Ullmer's account of "Tangible Bits" (1997) established this possibility, arguing that digital information can be accessed through physical objects rather than separated from the material world. Because physical objects already have properties that users can perceive directly, interactions grounded in them can feel more immediate than those confined to screens and keyboards. Placing an object on a surface to initiate a computational process therefore extends an established principle of tangible computing rather than introducing an arbitrary input method.

Dourish's account of embodied interaction (2001) provides a broader explanation of why this matters. He argues that meaning in computing emerges through engagement with a material environment, not through the manipulation of abstract symbols alone. Drawing on Heidegger's idea of the ready-to-hand, Dourish describes interactions in which the system recedes from conscious attention and becomes part of an ongoing activity. Applied to this project, choosing and placing an object should feel like a meaningful gesture rather than a technical operation. The system should stay in the background while choosing and placing the object remain the central actions.

Affordance theory further supports this approach. Gibson (1979) argues that the perceived properties of objects suggest how they can be used, while Norman (1988) develops this into a design principle: form should make possible actions apparent. This supports the use of household objects rather than purpose-built tokens or tagged artefacts. Their differences are not something the design needs to eliminate; shape, weight, texture and material are part of what the AI interprets.

Tangible music systems provide an important precedent. The reacTable (Jordà et al., 2007), for example, demonstrates that moving physical objects across a surface can support expressive musical interaction. This project applies the same general principle to music discovery rather than music creation. Existing physical interfaces for streaming typically control playback after a choice has already been made, allowing users to adjust volume, skip tracks or manage playback. To the author's knowledge, they do not use an object with personal meaning to determine what is played. The theoretical basis for physical objects as expressive interfaces is therefore well established, but their use in music discovery remains largely unexplored. The project that follows aims to address this gap.

Design implication.

D21 Make placing the object the complete input action, requiring no typing, tagging or confirmation, and keeping any screen used elsewhere in the system out of the initial choice.

D22 Keep the computation in the background so that choosing and placing an object is experienced as a physical gesture rather than as the operation of a digital system.

D23 Work with the variety of ordinary household objects rather than a designed set of tokens, so that differences in shape, weight, texture and material remain part of what the system reads.

D24 Let the object's perceived qualities shape the music that is returned, so that it contributes to the recommendation rather than only triggering one.

2.7 Objects as Carriers of Personal Meaning

The case for using personal possessions as inputs is not simply that they can be handled. It is that they already belong to the user's life. (Csikszentmihalyi and Rochberg-Halton, 1981), in their study of object meaning in the home, find that people invest their most significant possessions with memories and associations that are specific, deeply felt and mostly invisible to anyone outside the relationship. The authors call them "signs of the self": material anchors for identity and memory. This matters to the project because the objects a person reaches for in their own home carry a depth of association that no supplied token or designed artefact could match. Work on material experience in design reinforces the point, treating the felt and associative qualities of materials as central to how objects mean (Karana et al. 2015).

Belk (1988) extends this in his account of the "extended self," showing that people experience personal possessions as helping to form identity rather than just an expression of it. For the interaction model, choosing a personal object as an input is therefore a self-expressive gesture carrying the accumulated meaning of a relationship, not the selection of an arbitrary visual stimulus. The instinct that guides the choice, reaching for this object rather than that one, is informed even when it cannot be put into words.

This also introduces an important limitation. Everyday objects do not automatically produce intuitive interactions. Their meanings vary between people and cultures, and an object's physical affordance does not necessarily indicate how it should influence a musical choice. The project therefore distinguishes between an object's observable properties and its personal significance. An AI system may identify a colour, shape or material, but it cannot infer why a particular object matters to its owner. Personal meaning must be treated as something the user contributes to the interaction, rather than something the system can simply read from the object.

This raises an open design question: if an object's meaning is personal, while the AI perceives only its visual and material qualities, how does this gap shape the interaction? Rather than needing to recover the owner's intended meaning, the AI can offer an associative interpretation based on what it perceives and on broader cultural patterns. The user may accept, reject or be surprised by this reading. The design question is therefore whether the distance between personal meaning and machine interpretation can become a productive space for reflection, prompting slower sense-making after the immediate act of choosing and placing the object.

Design implication.

D25 Use objects the listener already owns rather than props supplied by the designer, so that the personal history the machine cannot access still enters the interaction through their choice.

D26 Present the system's response as an association rather than a correct reading, treating the distance between what the object means to the listener and what the machine can perceive as something the interaction makes use of.

D27 Allow the listener to choose an object instinctively without having to justify it, and to reject a result, retry or choose another object without that mismatch being treated as a failure.

D28 Do not assume that a given object carries the same meaning for every listener or across cultures, and avoid building fixed object-to-music mappings on that assumption.

D29 Give each result time to unfold, allowing space for listening and reflection rather than presenting recommendations as a rapid feed.

2.8 Prediction to Interpretation

This project does not eliminate prediction; it reorients what the system is asked to predict. Rather than estimating a listener's preference from past behaviour, Curio begins with a present physical object and returns music in response to an AI-generated association based on the object. The distinction from the affective systems discussed in 2.3 is not simply that the input is present: those systems also respond to photographs, faces or bodily signals. The difference lies in how the input is treated. They translate it into a state to be identified and resolved, whereas Curio treats the object as an open prompt for interpretation.

Conventional recommendation generally begins with what is already known about the listener, such as their listening history, declared mood, or inferred affective state. Curio instead begins with what is present in front of the system, while recognising that the object's personal significance is not directly accessible to the AI. That significance enters indirectly through the listener's decision to choose the object and through their response to the resulting music.

The systems reviewed in §2.3 can be understood along two dimensions: whether recommendation begins with past listener data or a present input, and whether the system treats interpretation as resolved or open. Most occupy the side of resolved interpretation, whereas Curio combines a present input with unresolved interpretation. Its output is therefore not presented as a forecast of what the listener will like, but as a tentative musical interpretation that the listener must encounter and assess.

This interaction follows from the relation between Stimmung and machine perception introduced in §2.4. The listener approaches an object with an unspoken sense that it matters to them, before they can fully explain why (Heidegger, 1927/1962). The AI cannot access that significance directly. It can respond only to observable qualities of the object and to associations between those qualities and human experience learned from its training data. The interaction can therefore be understood as a loop: the listener's attunement shapes the choice of object, the AI produces an associative reading, and the listener decides whether that reading is meaningful, surprising or unconvincing. Tsiakas and Murray-Rust (2024) treat the interaction, rather than the model alone, as the unit of design. Curio is therefore better understood as a system that offers an interpretation for the listener to assess than as an autonomous recommendation tool.

This also changes the value of error. In a predictive system, an incorrect output indicates a failure to match the listener's preference. In Curio, a slightly inaccurate reading may be valuable if it remains recognisably connected to the object. André et al. (2009) describe serendipitous discovery as unexpected but retrospectively recognisable as connected to something real. The same balance is required here. The music must relate closely enough to the object to be understood as a response, while leaving sufficient distance for the listener to question or reinterpret it. An adjacent reading can give the listener something to compare with their partly inarticulate understanding of the object. The design hypothesis is therefore not that any mismatch is successful, but that the useful result lies between a predictable match and an arbitrary one.

Karmann's Paragraphica (2023) provides a useful precedent because it exposes aspects of an AI interpretation through a physical object to be interacted with, uncertainty should lie in interpretation, not in complete concealment of how the system operates. Repeated encounters may allow listeners to develop a sense of the system's tendencies without learning a fixed rule such as "this object always produces this kind of music." This extends the slow-technology concern with return and reflection (Hallnäs & Redström, 2001). Whether repeated use deepens the encounter or eventually makes the system predictable is an empirical question for the evaluation. The prototype should therefore be assessed not only by the accuracy of its recommendations, but by whether it produces surprising but recognisable connections, preserves listener agency and gives interpretation time to develop.

Design implication.

D30 Configure the AI to form an association from the object present in front of it rather than to predict what the listener is likely to want.

D31 Present each result as a provisional reading that the listener can question, reinterpret or redirect, so that responding to the interpretation does not require rejecting the track it produced.

D32 Keep the association recognisably connected to the object while leaving enough distance for the listener to question it, so that results are neither repetitive nor arbitrary.

D33 Design for repeated use, so that a listener can develop a sense of the system's tendencies over time without those tendencies hardening into a rule they can predict.

D34 Evaluate the system by whether its connections are surprising yet recognisable, whether listener agency is preserved and whether interpretation has time to develop, alongside the accuracy of what it recommends.

2.9 The Research Gap

Chapter 2 has shown that the central problem is not simply whether a system can recommend music accurately. It is how and when the direction of listening is decided. Existing systems generally begin from listening history, a declared query or a present signal translated into an affective, contextual or similarity-based representation. This makes recommendation possible, but places interpretation before encounter: the listener must describe a direction, or the system must infer one, before discovery can develop.

The literature contains several relevant precedents. Personal objects, photographs and tangible artefacts have been used to access, trigger or curate associated media, including music (Chiquet et al., 2006; Frohlich et al., 2004; Kwon & Lee, 2018). Sound Frame uses a sentimental object to trigger music, but the listener assigns the song or playlist to the object in advance; the object stores an existing association rather than helping the listener discover one (Chung, 2021). Well-Tempered Objects uses computer vision to interpret previously unseen everyday objects through properties such as shape, size and position, but generates sound for tangible composition rather than retrieving unfamiliar existing music (Rayzhekov & Murer, 2025). Image-based recommendation systems retrieve music from visual input, although that input is resolved into learned visual, semantic or affective representations (Hsia et al., 2018). Conversely, systems such as Musicream, OLO Radio and Musée support exploratory or unfamiliar listening, but rely on screen-based similarity, listening history or personal metadata rather than a listener-selected object (Goto & Goto, 2009; Jang et al., 2022; Odom & Duel, 2018).

Taken together, these precedents leave a specific research gap. Although personal objects have been used to trigger or organise media, and AI systems have interpreted visual inputs for recommendation or sound generation, the reviewed literature does not investigate personal objects as situated, open-ended prompts for AI-mediated discovery of unfamiliar existing music. A situated prompt is an object selected by a particular listener, in a particular moment, and brought into the interaction with a personal significance that the system cannot fully access. This gap is particularly relevant to Stimmung: the condition in which a listener senses a desire or direction in listening without being able, or needing, to formulate it as a mood, genre or track. Curio addresses this gap by exploring whether a personally meaningful object can provide a material starting point for navigating that uncertainty.

The object is therefore not simply another contextual input. Its personal significance motivates the listener's choice and shapes how they interpret the result, but it is not something the system needs to recover as data. Curio's associative AI responds to perceptible qualities in the object and produces a provisional association that guides the retrieval of existing music. Associative AI offers a possible connection that the listener may recognise, resist or develop through listening.

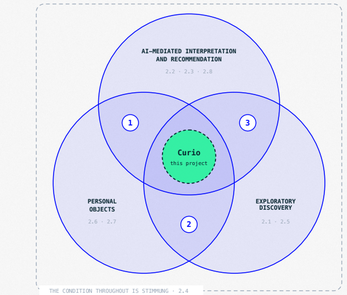
This relocates rather than removes agency; the listener does not begin by specifying a track, genre or mood, but decides what to bring, how to frame the encounter and how to respond to the system's reading. Stimmung is therefore not treated as a preference that already exists in a form the system can simply decode. It is an emerging orientation towards listening: something that can be sensed or pursued before it can be clearly stated. The object becomes part of an interaction through which that orientation may become clearer.

This leads to the project's research question:

How might a domestic, object-based music system use associative AI to help a listener navigate Stimmung and discover unfamiliar existing music without requiring them to declare or resolve their listening state before the encounter?

Curio is the design proposition through which this question is investigated. The device invites listeners to place personally meaningful objects onto a surface; AI vision produces a provisional associative reading, which guides the retrieval of unfamiliar existing music. The system does not claim to identify what an object means or to predict a listener's hidden preference. Instead, it tests whether an AI-generated reading can be sufficiently grounded for the chosen object to shape the encounter while remaining open to interpretation.

The central evaluation question is whether this arrangement can support a "nourishing" process of musical discovery. If different objects produce indistinguishable responses, the listener's choice has little effect. If the readings feel arbitrary, the encounter may lose its value. Evaluation must therefore examine whether object choices produce perceptibly different responses, whether listeners can find plausible relations between object, association and music, and whether they can accept, resist or redirect the system's suggestions. Curio's intended contribution is to investigate object-mediated recommendation as co-interpretation: a process in which AI offers provisional suggestions and listeners determine whether those suggestions open a worthwhile path through Stimmung and towards unfamiliar music.



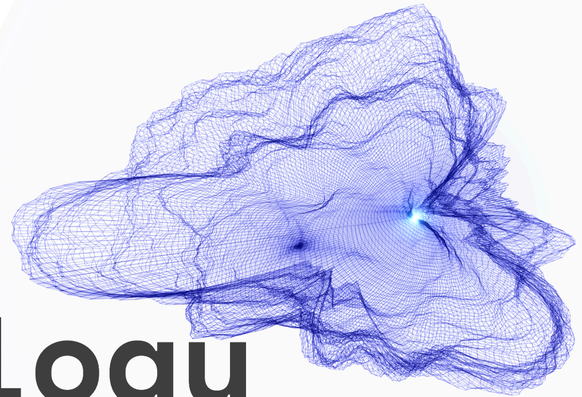
Literature Gap:

"personal objects as situated, open-ended prompts for AI-mediated discovery of unfamiliar existing music."

Right page:
Figure 2.4

Venn Diagram Highlighting the research gap identified.

Chapter 3



Methodology

This chapter defines Research through Design, explains why it was appropriate, and reflects on lessons learned from applying this methodology throughout the process.

What is RtD
How was RtD Applied

3.1 What is RtD

Right page:
Figure 3.1

A selection of research materials collected throughout the project.

Research Through Design (RtD) is a way of doing research through designing and making. A researcher builds an artefact to explore a question, then reflects on what the artefact reveals about that question. The artefact is therefore both the means of investigation and part of the research outcome. RtD produces situated design knowledge rather than universal laws; its contribution lies in explaining what was made then more importantly, why it was made and what was learned through the process (Gaver, 2012; Zimmerman & Forlizzi, 2008; Zimmerman et al., 2007).

//Why was it suitable

Research through Design was appropriate because the project investigated an evolving relationship among a listener, a personal object and associative AI, rather than evaluating a settled system. It asked how interpretive work might be distributed when listening intent remains unresolved. Qualities such as meaningfulness, traceability and listener agency could not be specified apart from interaction; they had to be made tangible through artefacts, encountered in use and examined through reflection. This aligns with RtD's production of situated and provisional design knowledge through making (Zimmerman et al., 2007; Gaver, 2012).

The artefacts were not implementations of a solution, but propositions about different ways of organising interpretation and discovery. Iterating and comparing them made it possible to examine what changed when more interpretive responsibility was given to the listener, the object or the AI. Interviews or measures of accuracy, novelty and diversity could have addressed narrower questions, but would not have revealed these differences.

The project did not aim to validate a finished product, establish a general theory of listener preference or make population-level claims. Its intended contribution was a set of grounded design propositions about the conditions under which object-mediated discovery might support a listener navigating Stimmung. RtD was suitable not because there was nothing to test, but because the relationship under investigation became visible through designing, experiencing and comparing alternatives. The resulting provocations are material arguments for examining and challenging the design space, rather than universal solutions (Gaver & Bowers, 2012)



3.2 How was RtD Applied

Figure 4.2

How design decisions flowed through the whole research project.

Research through Design (RtD) in this project was organised across three connected phases: grounding, exploration, and synthesis. This structure adapts Zimmerman et al.'s (2007) iterative design process of grounding, ideation, iteration, and reflection. Grounding established the context; exploration combined design implications with physical making and iterative refinement; and synthesis consolidated the learning. Reflective thinking was not confined to a single stage but occurred throughout the process, as emerging findings continuously raised new questions and reframing.

//Phase 1: Grounding

The literature review identified a research gap, which was investigated through fieldwork and a journey map in a record store. The resulting observations were treated as sensitising concepts (Blumer, 1954), acting as temporary guides to focus the research rather than fixed conclusions. Comparing these real-world observations with the literature revealed nine distinct dimensions of experiential loss, which defined the working question addressed by each design exploration.

//Phase 2: Exploration

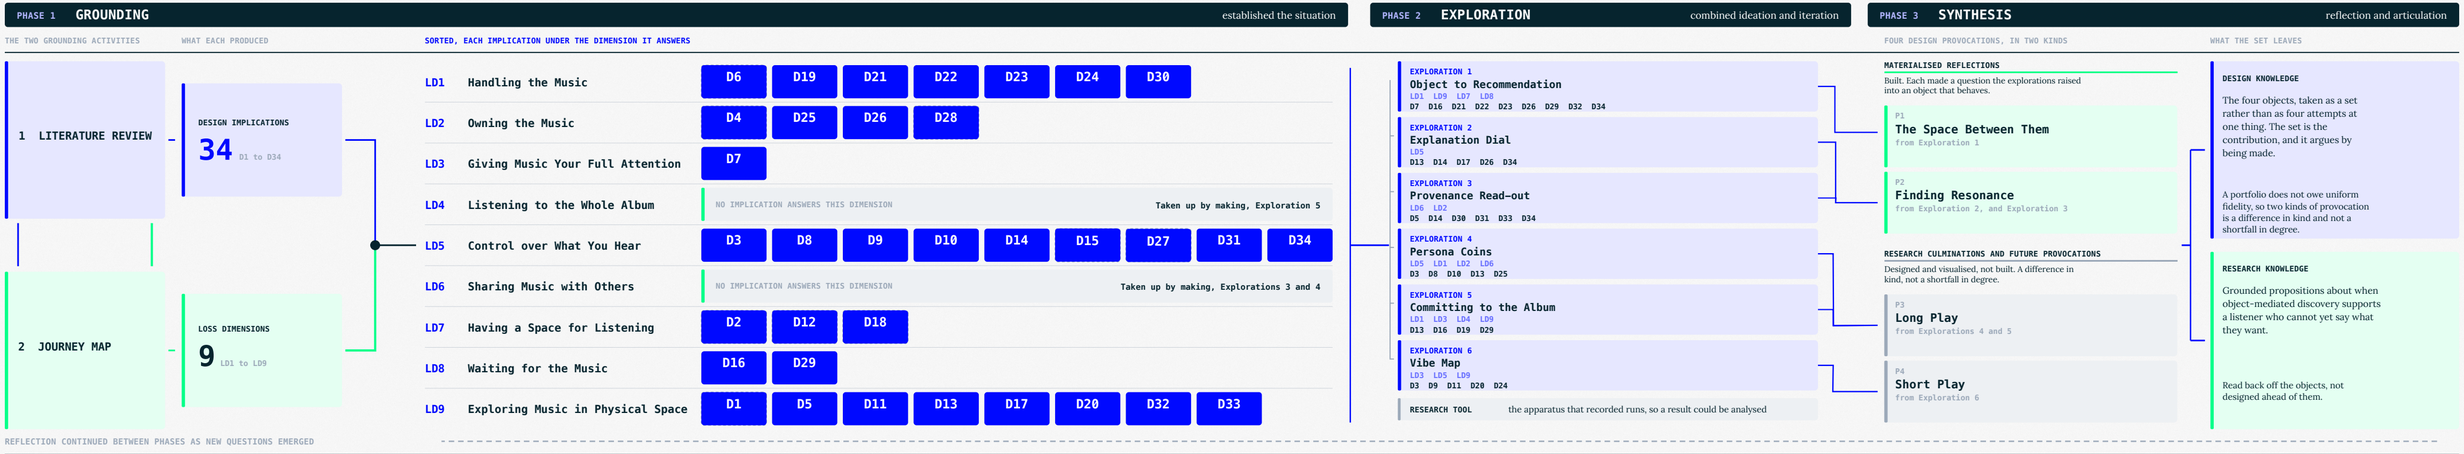
Six linked design explorations investigated these questions through iterative making and hands-on experimentation. Physical prototypes were built as research tools to isolate and study specific design qualities. Testing was done primarily through sustained, first-person use by the designer (autobiographical design), meaning that the focus was on acquiring deep design knowledge rather than validating general user needs.

//Phase 3: Synthesis

These explorations were synthesised into four distinct design provocations: the functioning research distillations of The Space Between and Finding Resonance, and the visualised future propositions of Long Play and Short Play. Although their physical formats differ, together they form an annotated portfolio (Gaver & Bowers, 2012) positioned along a placation and nourishment spectrum. They are presented as generative design suggestions rather than finalised, commercial solutions (Gaver, 2012).

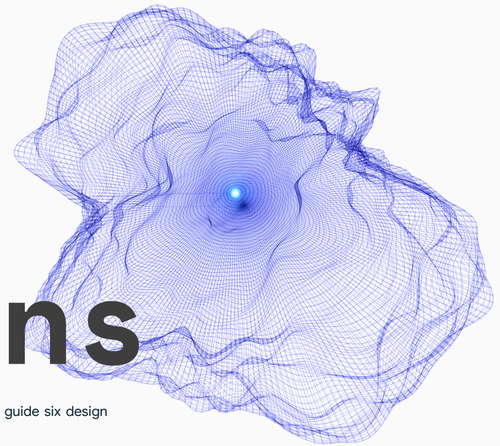
//Relationship Between Design Knowledge and Decisions

The process map [Figure 4.2] shows how literature-derived design insights (Dxx) and grounding activities generated provisional loss dimensions (Lxx). These informed the design implications and the working question investigated in each exploration. Prototyping, sustained personal use, and recorded observations provided the material for reflection, refining design knowledge and the core placation and nourishment criterion. These insights guided subsequent design decisions before being integrated into the final exhibition provocations.



Chapter 4

Explorations



This chapter synthesises existing literature and a journey map to produce 9 Loss Dimensions, which are then used to guide six design explorations into these areas of loss.

User Journey Map
Loss Dimensions
Research Tool
Exploration 1
Exploration 2
Exploration 3
Exploration 4
Exploration 5
Exploration 6
Tensions form the Explorations
Reflection and Synthesis

4.1 Research Process - User Journey Map

Right Page:

Figure 4.1
Image of the inside of record store the research was taken in:

Concerto Amsterdam
(Thompson, F. M.)

Existing research discussed in Chapter 2 examines vinyl and record stores largely through their social and cultural significance. No published study identified in this review presents an end-to-end map of the record-store experience as it unfolds for a user. Interviews and ethnographies explain why these spaces matter, particularly in relation to collecting and identity. Even when sensory or emotional details are included, they are usually used to support these broader cultural interpretations rather than examined as experiences in their own right (Whitehouse, 2023; Hracs & Jansson, 2020; Fernandez & Beverland, 2019; Guerra, 2022; Nieuwpoort, 2017).

The journey map was necessary because it offered a way to examine how the environment and medium of the record, record store, and record player shaped users' felt experiences. It made these specific experiential qualities visible and connected observations from the journey with the existing literature thus generating focused inputs for the design explorations that followed. It also provided a grounded basis for considering which qualities may have been weakened when music is accessed through streaming. The map does not claim to represent all listeners, measure emotional responses objectively, or demonstrate that the findings apply universally. Its purpose was therefore exploratory and design-oriented: to develop a situated account of how the experience felt and to identify qualities that could later be reintroduces or reinterpreted.

// Research Process

1. I first completed the journey independently, recording the different stages of my own experience.

2. I then repeated the journey with a participant. At each stage, both of us recorded notes about what we did, what we encountered, what we noticed in the environment, and how the experience felt.

3. After both journeys had been completed, the participant and I worked collaboratively on a mind map (Appendix B). We used our written notes alongside our memories of the experience to record observations, emotional responses, and insights from each stage. This process allowed us to compare what had been documented in the moment with what remained significant after the journey had ended. The resulting map is shown in [Figure 4.2] on the following page.

4. To support the recording of emotional responses, I used PrEmo, a pictorial emotion-measurement tool developed for communicating feelings through expressive characters rather than relying only on words(Appendix C). The tool allows participants to select one or more emotions and is suitable for qualitative interviews and collaborative sessions (Laurans & Desmet, 2017; Desmet, 2018). I chose it partly because I had used it before and was confident that it could support discussion with participants outside the design profession. Its visual format also helped us discuss subtle or mixed feelings without requiring the participant to use specialised emotional or design vocabulary.

5. This then opened up more focused research where existing literature could be compared the observations from the journey map. This created an iterative process in which experiential observations were tested against research.

6. The combination of the journey-map observations and the literature was used to develop new research buckets. These insights where then brought together and defined into nine loss dimensions, which became the main experiential inputs for the later design explorations.

//Participants **Consent forms = Appendix E*



The Researcher
Casual record collector and frequently visits record stores



The Particiapant
Has never bought a record and has only been in a record store a few times

The participants were myself, as the researcher, and one additional participant. Because the project was largely ethnographic in approach, I considered it important to participate directly in the journey rather than only observe another person. Experiencing and analysing the emotions from within the situation allowed me to examine qualities that might not have been visible through observation alone.

I regularly visit record stores and own a record collection. This gave me familiarity with the practices, objects, and cultural meanings associated with the medium. To provide a contrast, I selected a participant who had only visited a record store a few times and had never bought a record. This contrast was intended to place greater emphasis on the immediate experience of the environment rather than on established knowledge of record collecting or the cultural significance of vinyl.

The participant was not intended to represent all inexperienced listeners, just as my own account was not intended to represent all experienced collectors. Instead, the contrast helped make different aspects of the experience visible. My familiarity provided an informed account, while the participant's limited experience helped draw attention to elements that might otherwise have been overlooked or taken for granted.



4.1 Research Process - User Journey Map



researcher

participant

Left + Right Page:

Figure 4.2
Journey map of the whole record store experience.

// Important observations

1. Searching itself is the core act; the result is secondary to the experience of getting there.
2. Exploration should begin from an initial place of comfort, then offer exciting opportunities to branch out.
3. Anticipation is key to making an output feel rewarding.
4. The exploration space must be engaging enough to invite exploration, but fixed enough that exploring it feels low-cost.
5. Forcing a definitive final choice is key to making the outputs feel worthwhile.
6. Music exploration using the body creates a greater sense of agency and offers more inputs, making exploration feel more natural and less conscious.
7. Being able to focus in during exploration, and quickly focus back out, creates a more dynamic environment.
8. Visual exploration is more instinctive than using language.
9. Exploration in a physical space allows people to guide each other, creating new possibilities beyond solo exploration.
10. Exploration in a 3D space is heavily shaped by social cues: where people are standing, and which areas are untouched, guide different people through their own exploration journey.
11. Music listening in an open space can become a place for interaction.
12. Care and attention can be given to the listening experience by requiring precise or delicate interaction to play music.
13. Music exploration should happen alongside other experiences. It can enhance, and be enhanced by, daily life, lowering the barrier to engaging with music exploration as a separate activity.
14. The exploration and listening experience should be finite; when an end is reached, the listener should be asked a question.

Although the participants entered the record store from different starting points and with different expectations, their journeys showed a substantial degree of alignment. This consistency suggests that the experiential value of the record store is not solely dependent on nostalgia or enthusiast culture. Instead, its qualities may hold broader relevance for the design of music-exploration experiences.

Across both journeys, the emotional landscape was predominantly positive. Desire, fascination, joy, and satisfaction were the most frequently observed emotions. Desire generated forward momentum through the space, while fascination and joy reflected engagement in the present moment. Satisfaction typically emerged when exploration led to a meaningful outcome. Together, these findings show that the record store supports a rich and sequential experiential journey: one that moves from curiosity and exploration towards a personally meaningful sense of completion.

4.2 Loss Dimensions

The nine loss dimensions identify concrete experiential qualities that emerged through the journey-mapping process and were supported by the literature. They are not claims about vinyl's broader cultural significance. The dimensions are distinct enough to support separate design explorations, but they are not mutually exclusive. Several overlap. For example, handling a record may create anticipation, while a dedicated listening space may support sustained attention. This overlap is intentional, as it enables later design work to examine how these qualities interact, reinforce one another, or create tensions. Rather than prescribing particular solutions, the dimensions provide a framework for restoring or reinterpreting aspects of the vinyl listening experience.

Loss Dimension 1

Handling the Music

Handling music involved a careful, tactile relationship with the record. Searching could be more forceful, involving flicking through sleeves, picking up records, and turning them over, whereas playback required precision. Removing a record from its sleeve and placing it on the turntable demanded care, since careless handling could damage the record or the equipment. These actions slowed the transition into listening and gave it consequence. The physical object remained present throughout, making listening a deliberate activity rather than an immediate one.



Where the literature sits

- §2.5 Discovery in the Record Store
- §2.6 Physical Objects as an Interface

Loss Dimension 2

Owning the Music

A record collection could accumulate as a visible record of taste, memory, and changing interests. Its continued physical presence made it secure, revisit-able, and open to reflection. The collection showed both existing interests and gaps, making the process of building it tangible. Because records were acquired gradually, each addition could carry greater significance and strengthen the owner's attachment to the collection. Over time, the collection became part of the listener's identity and a durable personal archive.



Where the literature sits

- §2.5 Discovery in the Record Store
- §2.7 Objects as carriers of personal meaning

Loss Dimension 3

Giving Music Your Full Attention

Searching for music required focused attention and active judgement from the listener. Once a record began to play, there were few opportunities for interaction beyond stopping or adjusting it; changing tracks or selecting something else required a further physical action. This encouraged the listener to remain with the music, follow its development, and allow concentration to deepen. Listening could therefore become an activity defined by duration and commitment rather than background accompaniment.



Where the literature sits

- §2.1 Limits of Streaming Discovery
- §2.5 Discovery in the Record Store

Loss Dimension 4

Listening to the Whole Album

The physical format encouraged music to be experienced as an album rather than as isolated tracks. Moving between songs involved locating the desired position manually, while changing records required time and effort. As a result, track order, pacing, transitions, and side breaks shaped how individual songs were understood. The album could be experienced as a structured work with a beginning, development, and ending.



Where the literature sits

- §2.1 Limits of Streaming Discovery

Loss Dimension 5

Control Over What You Hear

Both record-store exploration and playback at home required the listener to make decisions in order to hear music. They had to determine what to seek, what to select, and what to play. Discovery was guided by personal curiosity and judgement rather than delivered automatically. Although recommendations, displays, and other people could influence these choices, the listener remained responsible for the path taken and the music encountered.



Where the literature sits

- §2.1 Limits of Streaming Discovery
- §2.4 Stimmung and Uncertain Intent

Loss Dimension 6

Sharing Music with Others

Physical music environments allowed other people to shape exploration, both directly and indirectly. Someone might influence discovery by placing a record prominently, browsing alongside another person, recommending an artist, or giving a recording as a gift. These encounters created opportunities for discussion, exchange, and shared attention. Because sharing music often required effort, intention, or knowledge, it could contribute to relationships between listeners, collectors, and music-makers.



Where the literature sits

- §2.2 AI in Music Recommendation
- §2.5 Discovery in the Record Store

Loss Dimension 7

Having a Space for Listening

A dedicated listening space gave music a recognised place within domestic life. The room could support particular routines, postures, and levels of attention, making listening an activity that one entered rather than sound that occupied any available space. Its continued presence also made the practice visible and easier to sustain over time.



Where the literature sits

- §2.7 Objects as carriers of personal meaning

Loss Dimension 8

Waiting for the Music

The slower process of obtaining and preparing a record created anticipation before any sound was heard. Travel, selection, and setup formed a period between wanting to hear something and hearing it. This time directed attention towards the forthcoming music and allowed expectation to build. The first sound could therefore feel like the arrival of an experience already imagined by the listener.



Where the literature sits

- §2.5 Discovery in the Record Store
- §2.8 Prediction to Interpretation

Loss Dimension 9

Exploring Music in Physical Space

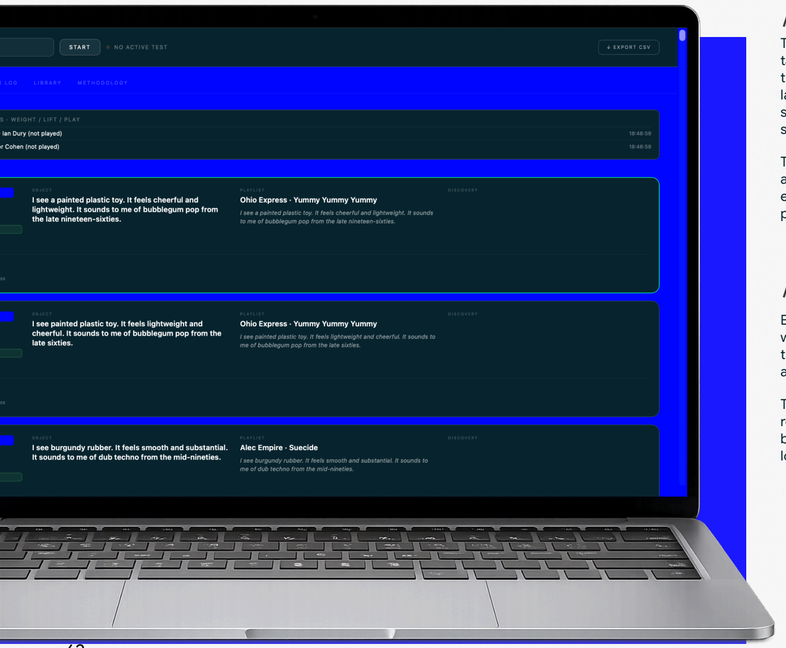
Physical environments enabled exploration through bodily movement, spatial memory, and material cues. Browsing shelves, moving between sections, and returning to records encountered earlier created a tangible system for navigation. Objects remained where they had been found, allowing discoveries to be revisited and supporting an unconscious map of the space. Partial visual cues, such as an unfamiliar cover, a worn sleeve, or a record placed at the front of a stack, could direct attention before the listener knew exactly what they were looking for. Repeated visits built spatial knowledge, making discovery feel gradual and earned.



Where the literature sits

- §2.4 Stimmung and Uncertain Intent
- §2.5 Discovery in the Record Store

4.3 Research Tool



// What it is

The Lab is a second window onto the same system. The listener's phone takes the photograph and sends it to a laptop, where the reading and the recommendation are worked out. The phone gets back a result. The laptop keeps everything else: the photograph, what each stage of the system produced, the exact instructions it was given, how long each step took, and what it cost.

The Lab is where that kept material is read. It holds a live view of scans as they run, a side-by-side comparison of any two, a searchable log of every scan ever made, a view of the listener's taste map, and the project's methodology notes.

// How it works

Each step of the system attaches a record of itself to its own output: which model ran, how long it took, how many words it consumed, and the exact instruction it was given. That record is removed before anything reaches the phone and kept in full on the laptop.

The split is the whole design. The listener sees a reading. The researcher sees how the reading was produced. The two are separated by the architecture rather than by discipline, so the record cannot be lost by simplifying the interface.



// Why it mattered to the exploration

1. It provided evidence against which self judgement could be assessed. Research through design places the researcher within the work, so the exploration was partly guided by the researcher's response to each output. The Lab recorded the system's actions alongside that response, allowing readers to evaluate the interpretation against the run that produced it.

2. It preserved authorship and accountability in a system developed with AI assistance. Because much of the code was written with an AI collaborator, there was a risk of producing a working system whose behaviour could not be fully explained. Recording each instruction, version and intermediate output made every behaviour open to inspection and justification. This was essential in a project focused on machine interpretation.

3. It preserved an unobstructed version for comparison. The design moved from a literal reading towards a more withholding interpretation. The effects of this abstraction could only be judged against the original.

5. It made documentation passive, allowing exploration to remain open. The researcher did not need to decide in advance what was worth recording or interrupt the process to take notes. Reflection could happen later using a complete record rather than memory. Several findings in this chapter were identified weeks after the relevant scans.

6. It made cost and duration explicit design constraints. Logged usage revealed a tenfold difference between live searches and cached results, which led the system to default to the cheaper option. A scan costing fifty pence and taking eighty seconds creates a different interaction and research approach from one costing five pence and taking fifteen seconds.

// How itwas applied

Exploration 1. Allowed the same object to be compared across prompt versions and across two different engines, which is what made the discriminability question answerable at all.

Exploration 2. Exposed a clear system logic behind abstract (poetic) outputs

Exploration 3. Allowed the sources behind each recommendation to be kept, so credited web results could be checked after the fact.

Exploration 4. Allowed the personas active for each result to be recorded, so the effect of adding or withholding a coin could be measured rather than estimated.

Exploration 5. Allowed for comparison across different weight readings to be mapped to different regions of the users listening history.

Vibe map. Wasn't used.

Left page:

Figure 4.3

Display showing "The Lab"

Figure 4.4

Diagram to show how information is exchanged within the system.

4.4.1 Exploration 1 - Intro



OBJECT TO RECOMMENDATION

//THE CORE CONCEPT

//Overview

This exploration emerged from a gap identified in the research: Streaming services often present music through a defined search or an inferred prediction, leaving limited room for discovery that begins without an explicit request. The exploration therefore examined whether an ordinary object, with no predefined relationship to music, could become an analogous starting point for a musical discovery.

The initial proposal used a camera. A participant would photograph a scene, and the system would interpret its content and framing to produce music. Developing the prototype revealed two problems. First, framing a photograph already involves an intention: by choosing what to capture, the participant decides what matters and may form an expectation of the response. Second, an image contains many competing details, whose relative importance may be interpreted differently by the participant and the system. The camera therefore made the input too deliberate and its interpretation too unstable.

The project consequently shifted towards objects already present in the room. A person does not usually pick up a ceramic mug because they have decided what they want to hear. An everyday object offered a less explicit point of departure while retaining a physical and open interaction. Together, the research and loss dimensions led to three design constraints:

1. The object should be something to search for, hold, and place, rather than a representation of a genre or musical category. Its physical presence and context should be able to influence the music without being reduced to a fixed label.
2. The interface should not reintroduce an explicit request. This excluded a query field and a confirmation step.
3. The system should remain as small as possible. If it performed too many operations at once, it would become difficult to identify which aspect of the interaction produced value.

The resulting core interaction was simple: choose an object, place it in front of the camera, press one button, and receive a track. This path remained stable while different approaches to interpreting the object and selecting music were tested.

The work then focused on two questions: how specifically did the system need to interpret the object, and how much technical complexity was necessary for users to experience a meaningful connection between the object and the music? The aim was not to identify an objectively correct track or rely on my own judgement of recommendation quality. Instead, the exploration examined whether users could experience the music as a potentially nourishing response to the object. The meaning of “nourishing” was intentionally left open for later experimentation.

The main result was not a measure of recommendation quality, but a shift in where the value of the experience appeared to lie: in the sequence of searching, choosing, handling, placing, waiting, and listening. This finding redirected the remainder of the project.

//Goal

To build the smallest working path from a physical object to a nourishing musical output, while enabling exploration of what makes the response feel like a good match.

The exploration is complete when the system can reliably turn an object into a playable track and users experience a connection between their search for an object and the music that follows. There is no predetermined “right” recommendation; what matters is whether users experience a meaningful or nourishing connection between the object and the music.

“Nourishing” remains intentionally undefined at this stage. It is a quality the system should enable, whose meaning will emerge through later experimentation.

Left page:

Figure 4.5

Image taken of the test's setup used during this exploration.

4.4.1 Exploration 1 - Input

// Loss Dimensions Exploration

Primary

1. LD1 Tactile Engagement

2. LD9 Embodied Exploration

Secondary

3. LD7 Dedicated Home-Listening Space

4. LD8 Anticipation

LD1 and LD9 are the dimensions this exploration was built to investigate. LD7 and LD8 run through the wider project to different degrees. They are included here because this exploration provides the first evidence for each.

These dimensions are treated as design questions rather than qualities assumed to be valuable in advance. The exploration tests whether searching, choosing, holding, placing, and waiting can make the final track feel more connected to the object. The aim is not to recreate vinyl or the record store, but to translate some of their qualities into a domestic object-to-music system.

LD1: Tactile Engagement

From LD1, I took the importance of having something physical to hold, place, and prepare for listening. With vinyl, the listener takes a record from its sleeve, holds it carefully, places it on a turntable, and sets it in motion. These actions are not simply preparation; they are part of the listening experience. They slow the transition into music, give it physical weight, and keep the object present throughout the activity.

This exploration translates that quality to an everyday object. Rather than handling a record, the listener holds an object and places it in front of the system. This also creates an important difference. A record is relatively legible: its sleeve, label, grooves, and turntable make clear what it is and how it is meant to be played. The object used here does not explain what music it will produce, and the system's interpretation remains hidden. The exploration therefore keeps the act of holding and placing but replaces a readable mechanism with an opaque one, testing whether the tactile act can still matter when the relationship between object and music is not immediately understandable.

LD9: Embodied Exploration

From LD9, I took the idea of searching through physical space for something worth playing. In a record store, discovery happens through movement. The listener looks across shelves, scans covers, handles records, and follows cues such as familiarity, newness, or intrigue. The search is not simply a delay before reaching the music; it is part of deciding what is worth hearing.

Streaming often reduces this process to a text field or recommender, where the time before a result is treated as friction to remove. This exploration moves the search back into physical space: instead of searching through a record store, the listener searches through their own room for an object to place before the system. The search therefore becomes part of the input, shaping both the object selected and the context given to the final output. The exploration tests whether that physical search contributes to the way the resulting track is experienced, rather than simply delaying access to it.

LD7: Dedicated Home-Listening Space

From LD7, I took the room rather than a fixed listening routine. This is the only exploration carried out entirely inside the home, and it uses ordinary domestic objects as its material. This is not incidental. The system depends on objects being within reach and carrying some personal or situational meaning. The home is therefore not just the setting for the interaction; it becomes part of the system's input, directing the listener back towards the room in which they already live and listen.

This is only a partial form of LD7. The exploration does not establish a dedicated listening room or a stable listening routine. Instead, it provides early evidence that an ordinary domestic space can become part of the listening experience.

LD8: Anticipation

From LD8, I took the interval between deciding to listen and hearing the first sound. In this exploration, that interval includes searching the room, choosing an object, placing it in front of the system, and waiting for the response.

At first, this delay was an accidental consequence of the prototype. It later became deliberate. The interval creates the possibility of anticipation, allowing the listener to wonder what the system might produce before the music begins. The exploration therefore tests whether waiting changes how the track is received, rather than assuming that delay is valuable by itself.

// Design Implications Applied

D1 Allow the listener's sense of what they want to develop through listening itself, rather than requiring them to name an artist, genre, mood or activity before the first track plays.

D6 Base the recommendation on the physical artefact present in front of the system rather than on the listener's accumulated listening history.

D7 Do not treat next-play accuracy or sustained attention as the sole objective, since holding a listener's attention and opening unfamiliar territory are distinct design targets that can pull against each other.

D11 Treat the listener's not yet knowing what they want as the condition the interaction begins from, rather than as a gap the system must close before it can respond.

D12 Allow the object, the music and the surrounding environment to take part in forming the listener's preference, rather than treating that preference as something already fixed inside the listener and waiting to be retrieved.

D18 Treat the setting as part of the discovery mechanism, since encounter-based discovery depends on a place that supplies objects within reach, time to dwell and existing personal associations, which in this project is the home.

D21 Make placing the object the complete input action, requiring no typing, tagging or confirmation, and keeping any screen used elsewhere in the system out of the initial choice.

D22 Keep the computation in the background so that choosing and placing an object is experienced as a physical gesture rather than as the operation of a digital system.

D23 Work with the variety of ordinary household objects rather than a designed set of tokens, so that differences in shape, weight, texture and material remain part of what the system reads.

D24 Let the object's perceived qualities shape the music that is returned, so that it contributes to the recommendation rather than only triggering one.

D25 Use objects the listener already owns rather than props supplied by the designer, so that the personal history the machine cannot access still enters the interaction through their choice.

D28 Do not assume that a given object carries the same meaning for every listener or across cultures, and avoid building fixed object-to-music mappings on that assumption.

D29 Give each result time to unfold, allowing space for listening and reflection rather than presenting recommendations as a rapid feed.

D30 Configure the AI to form an association from the object present in front of it rather than to predict what the listener is likely to want.

D34 Evaluate the system by whether its connections are surprising yet recognisable, whether listener agency is preserved and whether interpretation has time to develop, alongside the accuracy of what it recommends.

Newly applied

Re-applied

4.4.2 Exploration 1 - Development Process

1 2 3 4 5 6 7

1. The input shifted from the outside to ordinary objects

Scenes made framing part of the request, while objects offered a less deliberate starting point.

Shift. I began with a camera premise: a person photographed a scene, and the system interpreted its content and framing. Building and testing it exposed two problems. Framing already declared what the person considered important, and a scene held many competing details that the person and the system could interpret differently. Photographing objects in everyday settings, I found that individual objects produced more interesting associations than complete scenes, so I moved the starting point to ordinary objects already in the room.

Finding. The variable that mattered was not the camera itself but whether the first action required a musical decision. An object gave something concrete to interpret without making the interaction an explicit request. Objects became used and organised like a personal record collection.

theory development

L1 L9 D22

2. An opaque playlist tool was replaced with a controllable prototype

Building the app made the prompts, the intermediate reading, and the interaction sequence mine to set.

Shift. I first used ChatGPT's Spotify plug-in to create playlists from photographs. The results were generic, familiar tracks recurred, and I had little control over why they had been selected. The novelty faded quickly, making clear that the platform's recommendation system was contributing more than I could observe. To investigate, I built a small phone app with Claude Code, which let me control the prompts, change the information sent to the model, and inspect the intermediate reading before music was selected. I fixed the user-facing interaction: choose an object, place it in front of the camera, press one button, receive music. There was no query field, confirmation step, or menu. The first complete path ran from photograph to image reading to playable track.

Finding. The prototype became a way of examining the relationship between the object, its description, and the selected music. Fixing the interaction early gave later experiments a consistent starting point, while controlling the intermediate reading prevented the platform's recommendation system from becoming a hidden variable.

software development

D21 D22

3.The system was prevented from naming what it saw.

Naming produced literal associations, so the reading was restricted to qualities.

Shift. I tested different prompts for the reading stage and added an explicit instruction: do not name or describe what the object is. Instead, the system had to report qualities such as material, colour, wear, finish, shape, and atmosphere.

Finding. Allowed to name the object, the system anchored on the literal thing and often produced a direct association or a pun. Preventing it from naming the object forced the reading to work from qualities rather than categories. This became the main method for keeping the relationship between object and music open.

theory development

D26

4. The act of photographing became part of the interaction

Multiple views, changing light, and the countdown affected both the system's reading and the experience of capture.

Shift. I moved from one photograph to several, taken while the object was turned by hand on a countdown. Recognition improved, but capture also made the interaction feel deliberate: the countdown and the turning added a small ceremony not contained in the image data itself.

Repeated scans showed lighting and angle changing the system's reading. Warm and cool light pulled the description in different directions, and a change of angle altered which qualities became prominent. This also exposed a mismatch between my selection of an object and the system's view of it. I chose an object based on the whole thing, on what it was, what it meant, and why it interested me, while the camera sometimes gave the system only a partial view.

I also tested a version in which the reading became a catalogue search, letting the platform rather than the model select the track, then removed it. The catalogue's year data described the edition or reissue it held, not when the music was originally made. A five-year search window containing two landmark records returned only one track.

Finding. Capture was part of the interaction, not a technical step before it: the system responded to the object as captured, under particular lighting and from a particular angle. The catalogue also could not reliably supply the historical period the association needed, so that work had to remain with the model.

interaction development

L8 D16

5. Making the same object have the same output TEST 1

A deliberately awkward object set made the variable testable, and a visual fingerprint made objects recognisable across sessions.

Shift. I assembled an object set designed to challenge the system: objects describable in a short phrase, a sentence, or a paragraph; objects visually rich but with no clear referent; two near-duplicate pairs; and one object scanned repeatedly across a session under different angles and lighting.

This changed the main variable. I had initially focused on how literal an object was, but the more useful measure was how much language it took to describe one. The repeated scans exposed a further problem: the system could log the same object three different ways, because it was recognising its previous wording rather than the object itself. I replaced this with a visual fingerprint created from several photographs, which let the object be recognised across rotation and changing light.

Finding. The object set made descriptive complexity testable, while the visual fingerprint made comparison across sessions possible. If an object was going to carry meaning over time, the system had to recognise it independently of its changing description. This is where I settled on personal objects as the focus for input.

software development

D32

6. The exploration stopped competing with existing recommendation systems

External platforms produced stronger recommendations, shifting the contribution towards the physical interaction rather than output quality.

Shift. I ran the same object and task through a general assistant with music access and through the platform's own text-to-playlist feature. The external result was better than mine. Examining the difference showed that its advantage came from the platform's existing curation rather than from a better prompt.

Finding. The contribution of this exploration could not be recommendation quality. What remained distinctive was the interaction itself: searching the room, choosing and handling an object, placing it in front of the camera, waiting, and listening. A prompt could produce an answer, but it could not provide this physical sequence on its own.

theory development

D7 D34

7. The system was reduced to a repeatable sentence-to-track pipeline

A fixed prompt structure, two-stage reading, and single-track output made the relationship between object and music easier to examine.

Shift. The system had accumulated enough reading and selection logic that it no longer functioned clearly as a first step. I rebuilt it around one photograph, one sentence with five blanks, and one track. The first reading saw the photograph and filled the opening three blanks. The second received only the partial sentence, completed it, and named a track; it never saw the photograph.

A fixed frame replaced the open-ended prompt so the system filled defined gaps rather than writing freely. I also added two user-machine alignment elements, withheld from the recommendation engine, after earlier readouts varied in length and destabilised the hand-off between stages.

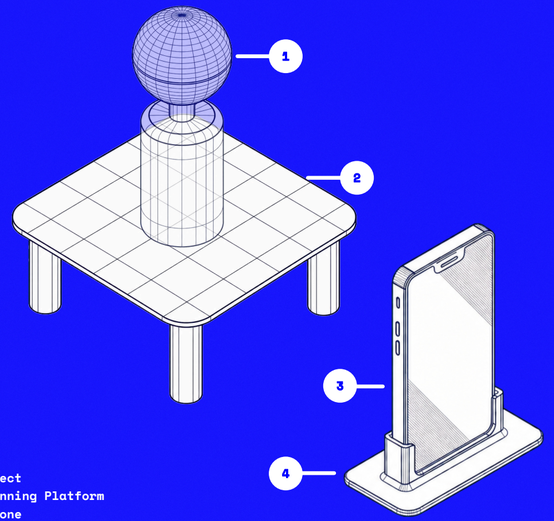
Moving from a playlist to a single track gave no objective measure of a good recommendation, but it made each response easier to assess as one encounter. I scanned the object set in a fixed order, noting expectation, interpretation, result, and feeling.

Finding. The readings separated by descriptive complexity: genre and period stayed stable across repeat scans, while feeling words changed more often. With the sentence as the only bridge between object and track, this version made the relationship more visible and testable than the larger system had.

software development

D29

4.4.3 Exploration 1 - Testing



1. Object
2. Scanning Platform
3. iPhone
4. Phone Holder

//Test setup

Eighteen objects from my own home, ranked before scanning on one axis: how hard I would find it to describe the object accurately. Each was scanned three times across three separate days in a shuffled order, and I scored every result one to five on how interesting I found it.

The system was the smallest version that worked. One photograph, one fixed sentence that filled in as it read, one record. A vision call named the object and two of its qualities; a second call received that half sentence alone, never the photograph, and completed it with a genre, a period and one named record.

No rubric was written, and only the ends of the scale were named. Defining interesting in advance would have decided the answer, and a definition was the thing I was after. The categories below were identified after scoring, not before.

//Test justification

In a research-through-design process, testing is justified when making alone cannot resolve a consequential design question. Here, testing was needed to articulate an implicit criterion of quality: different objects produced qualitatively different outcomes, but I could not yet explain why. This criterion informed later decisions, including which objects to select for the provotypes and which behaviours the system should adopt.

//Goal

To formulate an explicit, operational definition of what constitutes an interesting result and to further develop what "interesting" might mean in this context. The test would be complete when the objects could be classified as successful or unsuccessful, with that classification justified.

//Key Insights

1. Objects not used solely for their function, and ambiguous in themselves, produced the most interesting results by a clear margin. They had enough about them to be interesting, in surface, wear and form, and were ambiguous enough that I could not pre-declare what should come back before I scanned them.

2. Every failure in the set was the same failure. In each case I arrived at the scan already holding a settled reading, whether that the object meant nothing, that it was too specific to be moved, or that it was simply complete. Interest survived only where I had a reading and could not finish it, which makes this a property of the relationship between the listener and the object rather than of the object alone.

3. Objects that were too abstract failed, and they failed on my side rather than the system's. I could form no association with them at all, so there was nothing for a musical result to be relative to. The reading was not wrong. I had no way to relate it to the thing in my hand.

4. Objects carrying a specific personal meaning failed through a mismatch of resolution. The system read them generally, as paper and pigment, while I read them specifically, as a particular day and a particular person. The result answered the machine's reading, so it landed as the system missing the point. This is why a photograph, a drawing or a note is not the kind of personal object the first insight names.

5. Writing made the association too literal. Where text was legible the record followed the words rather than the object, and the labelled and written objects scored lowest in the set. The marathon medal is the counter case: its lettering has rubbed illegible, and it scored among the highest.

6. Highly functional and easily described objects were not specific enough to associate with. A short phrase covered each of them completely, and once that phrase was available both the system and I stopped at it.

A Curio is a materially expressive personal object whose significance exceeds its practical use. Its visual attributes carry meaning without representation or inscription. Interpretively open and resistant to fixed definition, it excludes representational objects such as photographs, drawings, and written notes. The term derives from curiosity, denoting an object valued for its rarity or interest rather than its function.

Left page:
Figure 4.6

Diagram of the test setup for Test 1.

Left + Right page:
Figure 4.7
All the object tested in this exploration, ordered with their mean "interesting score' printed below.
Full results can be seen in Appendix D

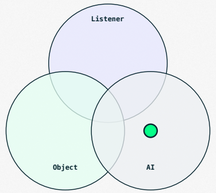


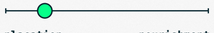
mean score	2.1	2.9	3.9	3.5	4.1	1.3	3.2	3.4	2.3	3.9	3.8	3.7	3.1	3.4	1.1	1.9	2.0	2.0

*full results can be seen in Appendix D

4.4.4 Exploration 1 - Output

Where does the agency sit





/The listener

The listener makes one decision: selecting an object. This replaces a declared search query and begins discovery in the domestic environment rather than in a catalogue. They then place the object, wait, read a sentence and hear a record. None of these later steps accepts input: they cannot correct the reading, request another, change its register or indicate that the result has missed.

The listener's agency is therefore entirely front-loaded into the choice of object. This is a deliberate baseline, establishing one end of the axis explored later: physical participation may make recommendation more situated, but physical input is not the same as ongoing agency.

//The object

The object functions simultaneously as query, material prompt and reason to search the room. Its personal history remains with the listener, while the system responds only to its visible and material qualities. This preserves associative openness but exposes a gap between lived meaning and machine reading.

Nothing is assigned to it in advance without the input system controls, the same object read twice would not necessarily produce the same words. The object initiates the process without determining its outcome. The object should therefore be understood as an occasion for interpretation, not as a stable source of truth.

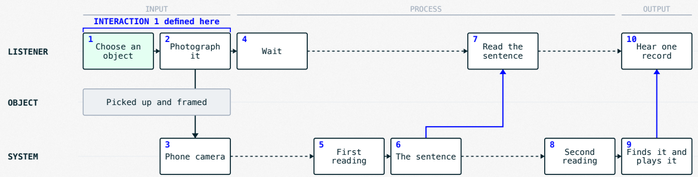
//The AI

The AI creates a linguistic bridge rather than a direct object-to-music relation. It uses a fixed sentence:

"I see ___. It feels ___ and ___. It sounds to me of ___ from ___."

The first reading describes the image by filling the material and qualities. A second reading receives only that half-sentence and completes it with a genre and period, then names one record. The photograph is withheld from this second reading, making the sentence the route from object to music rather than a caption beside it.

The AI is interpreter and decider at once. It determines how the object is described and what that description means, with no point at which the listener can intervene. The platform is only asked to find and play the named record; its recommendation engine is never called. The fixed, one-way pipeline is deliberately useful as a baseline, but its principal failure is not simply an inaccurate recommendation. It is the system's unwarranted confidence when the image is inadequate./The listener



WORKED EXAMPLE
A painted plastic toy, photographed on a phone.
"I see a painted plastic toy. It feels cheerful and lightweight. It sounds to me of bubblegum pop from the late nineteen-sixties."

Ohio Express – Yummy Yummy Yummy

// Key insights

1. The primary experiential value lies in searching for an object rather than receiving a high-quality recommendation. Searching was enjoyable in itself: it brought neglected objects back into attention and involved moving them around the house to arrange them beside the device. This shifts the design focus from recommendation accuracy towards situated, spatial and domestic exploration. However, this finding is based mainly on the author's experience and requires validation with other listeners.

2. Objects provide an indirect but intentional route into music. They are encountered through material, colour, texture, wear, scale and context. Their ambiguity makes them useful for associative recommendation because they allow several meaningful musical directions rather than demanding one correct answer.

3. The object reaches music through language. The system translates material perception into a sentence, and that sentence guides the musical selection. Preventing the system from naming the object keeps the recommendation associative rather than literal, but it can also make the musical connection harder to understand. The listener's response therefore depends heavily on the output and how the connection is presented.

4. Subtractive prototyping isolated the minimum causal chain from object to music. One photograph, one generated sentence and one record make the system's central proposition clear. Early testing produced surprisingly nourishing responses, although this does not yet show whether the associations remain convincing beyond the novelty of the interaction.

5. Capture and waiting have the potential to become part of the experience rather than mere latency. Handling the object, the countdown and the staged sentence make the transition to music perceptible. At present, however, the initial search and final output feel more nourishing than the process between them. Developing this middle stage is therefore a key challenge.

6. Some parts of an object reading are stable, while others change between scans. Genre and historical period are usually consistent, but texture and mood can change with lighting, angle and session. A visual fingerprint can recognise the same object, while the written description may drift because it is influenced by previous interpretations. This can create useful ambiguity, but it also raises problems of reliability and trust.

// Key questions

1. How can the system give the listener enough agency to question, correct or revise its interpretation without removing surprise?
   Answered in Exploration 5

2. What should the system reveal about its hidden interpretation, and when, so that the recommendation feels grounded without becoming literal or over-explained?
   Answered in Exploration 2

3. How might the process in-between initial interpretation through to output become as nourishing ?
   Answered in Exploration 6

Left page:
Figure 4.8

Diagram showing where agency sits by the end of this exploration

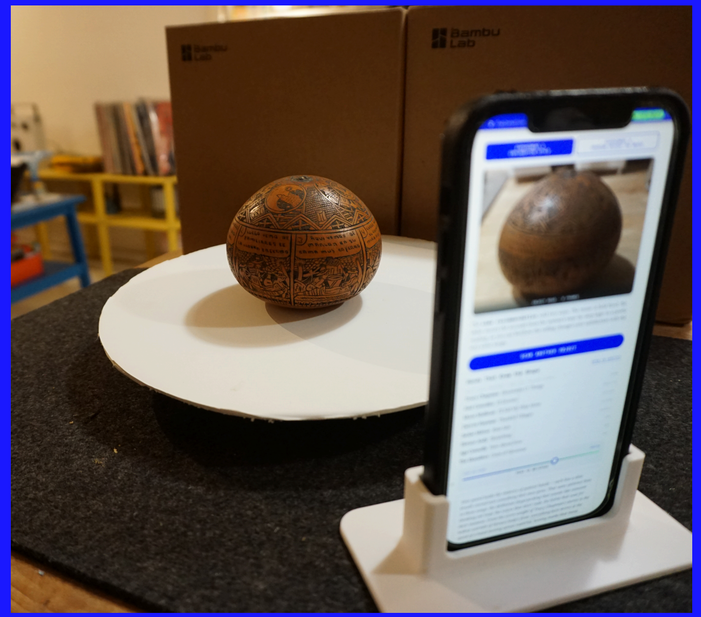
Left page:
Figure 4.9

Diagram showing where how nourishing the experience was for the listener by the end of this exploration.

Left page:
Figure 4.10

Diagram showing the system architecture developed by the end of this exploration

4.5.1 Exploration 2 - Intro



EXPLANATION DIAL

//EXPLORATION INTO TRANSPARENCY

//Overview

This exploration began where Exploration 1 left off. After an object was scanned, the system returned one fixed sentence. This kept the interaction simple, but concealed both how the machine had read the object and how that reading informed its recommendation. In this exploration, the system interpreted each object through 10 lenses. All readings were passed into the recommendation process, making its internal logic and read-outs more complex. This raised the central design question: should the object's interpretation shape only the output song or, be translated into an explanation for the user, or do both and to what degree?

Alongside these 8 readings, each recommendation produced a reasoning chain connecting the object to the track. The way a recommendation was explained appeared to affect how satisfied the user felt with it. A reasoning chain was not necessarily transparency; it could also be persuasion depending on how literal or poetic the read out was. This created a tension with Exploration 1, which treated the relationship between object and music as something to be felt rather than fully explained. The question became whether explanation would provide the desired transparency to user understand the system or encourage them to just accept it.

To separate these effects, the exploration used two tests. The first kept the music fixed and changed only how the recommendation was explained. The second changed the machine's interpretive stance to test whether it altered the selected music. Testing was also carried out with another user, comparing their desired output with the original user's response so that the findings did not rest solely on one person's opinion.

The testing revealed a clear and well-structured explanation can make a recommendation feel more convincing, even when the user may otherwise have questioned it. Providing more reasoning did not automatically give the user more agency; instead, it could steer them toward accepting a choice while still seeming to offer control.

This raised questions over where agency should sit in the interaction, what should be known to the listener and what should be only known to the AI system.

// Goal

Determine how much of the machine's internal logic should be made clear to the user, and how different levels of explanation affect the user's perceived resonance with the output.

Left page:
figure 4.11

Image taken of the test setup used during this exploration.

4.5.2 Exploration 2 - Input

Loss Dimensions Exploration

Primary

1. LD5 Control over What You Heard

LD5: Control over What You Hear

From LD5, I took the listener's agency over the journey through music and over the choices made along the way. In the record-store experience, the listener decides where to look, what to pick up, and which object to select. The reasons for that selection may not be fully articulated, but they are held by the listener as tacit knowledge. This knowledge can come from anywhere eg. memory, curiosity, association

This exploration asks how a system should respond to that tacit knowledge. When the AI analyses an object, it cannot access all of the personal knowledge involved in the listener's selection. Instead, it interprets the object through ten dimensions. These readings make the system's interpretation present, but they also raise a question about how much of that interpretation should be revealed. Should the system explain what it thinks the object means, leave space for the listener's own understanding, or remain partly obscure?

The record store represents agency over the journey and selection, even when the reasoning behind a choice remains tacit. The record represents a musical choice that arrives with an account of why it was made. The record player represents control over playback. These are not three separate losses, but three positions from which LD5 can be understood.

This exploration focuses on the relationship between the listener's tacit knowledge and the system's visible response. The musical recommendation may remain the same while the system changes how its interpretation is communicated. The question is whether making the system's reasoning more explicit supports the listener's agency, or whether it replaces the listener's own understanding with a machine-generated account. It therefore examines how obscure the system should remain when responding to knowledge that the listener possesses but has not expressed.

// Design Implications Applied

D13 Support gradual and reversible tuning of the result, so that the listener can adjust their direction over several moves rather than receiving a single recommendation to accept or reject.

D14 Give the listener feedback on how their input has been read without presenting that reading as the one correct interpretation, so that the ambiguity remains something they can work with.

D26 Present the system's response as an association rather than a correct reading, treating the distance between what the object means to the listener and what the machine can perceive as something the interaction makes use of.

D7 Do not treat next-play accuracy or sustained attention as the sole objective, since holding a listener's attention and opening unfamiliar territory are distinct design targets that can pull against each other.

D11 Treat the listener's not yet knowing what they want as the condition the interaction begins from, rather than as a gap the system must close before it can respond.

D12 Allow the object, the music and the surrounding environment to take part in forming the listener's preference, rather than treating that preference as something already fixed inside the listener and waiting to be retrieved.

D34 Evaluate the system by whether its connections are surprising yet recognisable, whether listener agency is preserved and whether interpretation has time to develop, alongside the accuracy of what it recommends.

● Newly applied
● Re-applied

4.5.3 Exploration 2 - Development Process

experimented with output medium
paper vs screen

1 — 2 — 3 — 4 — 5 — 6

1. I applied 8 reading lenses to the object

Ten reading lenses and a rotating base made more aspects of the object available to the system, but also made its interpretation harder to inspect.

Shift. Building on the first two explorations, which analysed objects through different aspects, I developed ten reading lenses. To make the process repeatable, I built a rotating base that captured five photographs at intervals rather than relying on hand movement. Together, the different views and lenses produced a fuller object profile. The prompt outputs became richer and more poetic, often surfacing qualities that earlier tests had missed.

Finding. The system became more expressive because it had more of the object to work with, but this also made it harder to identify what input was producing the musical output.

software development
hardware development

2. Richer readings changed how I received the music

The change affected not only the music selected but also my judgement of the recommendation.

Shift. More poetic responses were initially more interesting and playful to use. They provided new ways of thinking about objects that I had not previously considered. When I had a strong personal association with an object, however, a detailed but misaligned interpretation was harder to dismiss than a simple or literal one. As the language became more poetic, the system appeared to take on greater interpretive agency. It became unclear what had been read from the object and what had been added by the system.

I began to distinguish between nourishment and placation. A nourishing reading opened further attention towards the object and the music. A placating reading made the recommendation feel justified before I had decided whether I accepted it.

Finding. A richer explanation could make the interaction more engaging, but it could also make the system's interpretation feel more authoritative and less likely to align with my own.

theory development L5 D26

3. The ten reading lenses became weighted variables

The language used to evaluate the object introduced greater detail, but weighting this detail added complexity without producing a clear change in the selected music.

Shift. The poetic evaluation introduced linguistic qualifiers such as "very", "quiet", "light", and "heavy". I explored whether these terms could be translated into weights for each of the ten reading lenses, allowing the balance between them to alter the selected track. In theory, this offered more precise control over the relationship between the object reading and the music.

In practice, changing the weights produced little noticeable difference in the output.

Finding. The weighting system added logic without adding a clear change in the result. This suggested that the system could become over-engineered: each additional layer might make its operation harder to understand without making the recommendation more responsive. Transparency therefore depended not only on exposing the system's reasoning, but also on limiting the hidden logic between the object and the track.

software development D13

4. Explanation and selection were separated

Fixing one output while varying the other made it possible to distinguish changes in reception from changes in recommendation.

Shift. I tested the system through repeated self-use. For each run, I scanned an object, read the explanation at different slider positions, rated the result, and recorded what changed. The first experiment held the music constant and varied only the explanation. The second varied the interpretive stance to test whether it changed the music selected.

The poetic explanation produced the strongest, and most troubling, effect. A tidy account could make me accept a track that I might otherwise question. It also became repetitive: the argument followed a predictable structure, and each stance change required a new pipeline run of roughly a minute. Poetic explanations were less consistent, but most effective for objects that were difficult to describe.

Finding. Separating explanation from selection made the effects of each easier to examine. It also showed that explanation could influence how I received a recommendation independently of the track itself.

interaction development L5 D17

5. Experimentation with the materiality of the output

The format and amount of an explanation changed the authority it appeared to have.

Shift. I tested explanations as printed slips vs a screen output. A printed slip made the result feel more magical allowing for a desire for agency removal, a screen based one ment I felt like I desired more control.

Finding. A printed slip made the result feel more magical allowing for a desire for agency removal for the user, it felt like it took on a different roll to a traditional recommendation system. This highlighted the need to consider the output presentation as a traditional screen based one ment increased my desire for control and understanding of it.

theory development
dead end

6. Longer-term testing shifted the criterion from clarity to nourishment

TEST 2

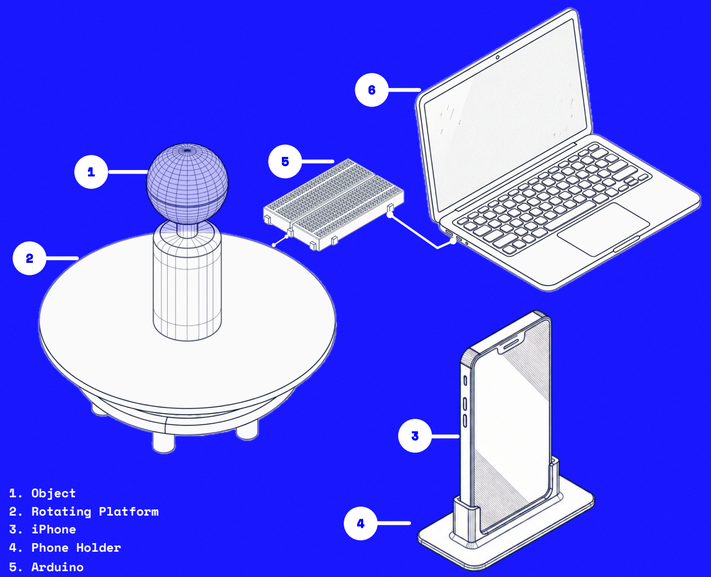
Brief demonstrations produced positive first reactions, but did not show whether the interaction remained meaningful over time.

Shift. When I showed the system to other people, their initial reactions were often that it was interesting or "cool". Repeated self-testing over a longer period, however, made the novelty wear off and reduced my interest in showing it. This pointed towards the need for longer-term, ethnographic testing with other listeners rather than relying on short demonstrations.

Finding. Explanation can placate as opposed to inform, a convincing account can make a recommendation feel justified while narrowing the listener's role to accepting or rejecting the system's interpretation. The exploration therefore tested control over how a recommendation was presented. The project shifted from transparency of the output being a core focus towards generating a nourishing journey towards said output .

theory development D34

4.5.4 Exploration 2 - Testing



1. Object
2. Rotating Platform
3. iPhone
4. Phone Holder
5. Arduino
6. Laptop

//Test setup

The researcher and one participant from outside the project, where tasked to work through their own possessions with the system running on their phone. They chose what to scan, scanned it themselves, and moved a control between two readouts of telling the same result. The session was held at the point in the development where the exploration's central idea was forming and I needed a perspective that was not my own.

The system was the version the exploration had reached. A scan returned a songs and an account of them, and the control moved that account between two registers. At the literal end it broke the object into the eight aspects the profile stage reads, era, region, texture, form, wear, energy, associative features, production, and traced each song back through them. At the poetic end it returned a title and two sentences that set a scene and disguised the system logic. Within a single scan the two registers described the same songs, so anything that moved between them came from the telling.

The format was deliberately loose as the goal was to get the thoughts of a reflective partner who was not part of the project, the participant was not briefed on the ideas I was exploring. The session functioned as a co-creative session where observations and thoughts where noted by the researcher and reflected upon after the fact. The key insights below are a synthesised version of the discoveries from this testing.

//Test justification

In a research-through-design process, testing is justified when making alone cannot resolve a consequential design question. Here the question was whether the register in which a result is accounted for changes what that result is taken to be worth, and it was the one question my own method was structurally unable to answer. I had written both registers, built the system that chose the songs, and knew before reading either telling which end of the control it came from. A reader who knows the manipulation cannot be moved by it, so the effect I most needed to observe was the one effect I had disqualified myself from having.

The session was kept loose for the same reason no rubric was written for Test 1. A controlled comparison would have required me to fix in advance what counted as a better account, and that judgement was precisely what I lacked and was trying to develop. What the session produced was therefore a set of directions rather than a measurement, and it is not offered as evidence about listeners in general. It informed later decisions directly: whether explanation was still the right variable to build on, and what form the next provotype had to take.

//Goal

To establish whether the account a system gives of its choice changes what a listener takes that choice to be worth, and to develop a reading of what explanation is actually doing to the listener who receives it. The test would be complete when the difference between the two registers could be stated as a mechanism rather than a preference, and when that mechanism could be used to decide whether control over the account belonged in the listener's hands.

//Key Insights

1. The way the system explained itself changed how much the participant liked the result, even though the result had not changed. Same songs, different reception. That puts the explanation inside the recommendation rather than around it, and it is what the rest of the exploration rests on.
2. The literal read-out gave the participant something to argue with. Clear steps meant they could check the reasoning, and checking it meant finding the parts that did not hold up. Being transparent did not make them more willing to accept the choice. It gave them the means to reject it.
3. The poetic version avoided that by never making a claim. It set a scene instead of an argument, so there was nothing to check and nothing to disagree with, and the result came through intact. But acceptance is not interest. Either way the conversation stayed on the reading rather than on the music it was supposed to be recommending.
4. The poetic version stopped working after a few objects. Once the participant could see how it was put together it began to feel formulaic, and by the last object it read as a gimmick. Any generated voice that always takes the same stance will do this, so it is a limit of the material rather than of this particular prompt.
5. Which version they wanted changed from object to object, and I could not predict it. Some objects they wanted read plainly, others they wanted left open. No single setting would have been right across the whole set, which is the strongest argument in the session for letting the listener control how much the system shows.
6. The user became bored whenever the answer was simply handed over, in either version. The interaction only held them when there was something left for them to do. That sits awkwardly against insight five, because the listener needs control over how much is shown, but a control they can settle is one they will use to make the question go away. So the openness cannot be a setting on a screen. It has to be built into the thing itself, which is what Provotype 2 was made to test.

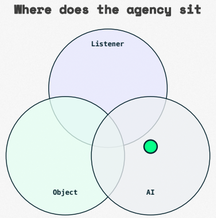
The session settled that the machine's account is part of the recommendation rather than a wrapper around it, and that no register of it works while the listener is left with nothing to do.

Left page:
figure 4.12

Diagram of the test setup for Test 2.

Participant Consent Forms an be found in:
(Appendix F)

4.5.5 Exploration 2 - Output



//The listener

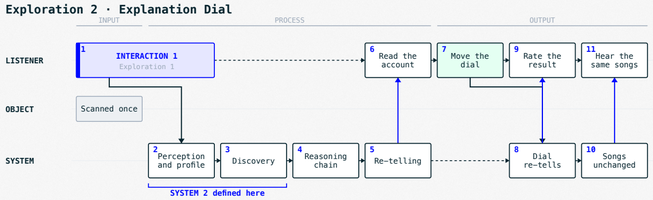
At the literal end they're handed a reasoning chain, and reading it is work: judging whether each step holds, supplying the interpretive energy to cross the jumps between object and track. Real agency, because there's something to be right or wrong about. At the poetic end there are no steps, so no step can be refused, and the only decision left is whether to accept that the writing caught a feeling they couldn't name themselves. A verdict on the whole, given once. Ends on your line: the system behaving more like a person reduces what the listener can do, because fluency closes the gaps judgement would have had to cross.

//The object

The ten aspects look like the object gaining ground, and that reading is wrong. Each aspect is abstract, an abstract reading has to be interpreted before it does anything, so naming more qualities creates more interpretation to be done and the system is the only party positioned to do it. The object supplies more material and less determination.

//The AI

moves opposite to the listener. At the literal end it gives up nearly everything: it still decides how the object becomes the output, but holds nothing over how that output is received, because the chain is on the page and can be walked through and disagreed with. At the poetic end it gains that second authority too, since with no exposed reasoning there's nothing against which a different reception could form. Perception stops being something the listener assembles and becomes something the account governs.



//Key insights

1. A clear and coherent explanation can settle the listener into accepting a recommendation while appearing to give them control, showing that transparency is not necessarily the same as agency.
2. Literal and poetic explanations produce different forms of agency because a literal reasoning chain exposes steps that can be checked or rejected, whereas a poetic account closes those gaps and makes the listener judge the account as a whole.
3. The ten reading lenses give the system more interpretive authority rather than giving the object more agency; they supply more material, but because each lens is abstract, the system determines more of the object's meaning.
4. The material form of the explanation changes its perceived authority, with a printed slip making the recommendation feel finite, magical, and easier to surrender to, while a screen encourages scrutiny and a desire for control.

//Key questions

1. What form of explanation allows the listener to test, refuse, or extend a recommendation without the system's account becoming the account?
   Answered in Exploration 6
   Answered in Exploration 5
2. How can agency move upstream from judging the explanation to shaping what the system selects, especially when the listener's reasons for choosing an object remain tacit?
   Answered in Exploration 4

Left page:
Figure 4.13

Diagram showing where agency sits by the end of this exploration.

Left page:
Figure 4.14

Diagram showing where how nourishing the experience was for the listener by the end of this exploration.

Left page:
Figure 4.15

Diagram showing the system architecture developed by the end of this exploration.

4.5.6 Stimmung Concept Development

The Journey

Before Exploration 2, Stimmung described a listener's developing sense of what they might want from music before they could explain it clearly. It did not refer to a mood that the system could simply detect. Instead, it located experience between the listener, object, environment, and sound (Gumbrecht, 2012; Heidegger, 1927/1962; Thonhauser, 2020). Not knowing was therefore treated as a valid starting condition rather than a problem that had to be solved before interaction could begin.

Exploration 2 initially treated not-knowing as something the AI could capture and shape through explanation. I assumed that a sufficiently detailed account of the object might influence how the listener understood the musical result and make the recommendation feel resonant. In practical terms, the exploration made the explanation adjustable: the listener could move between different explanatory registers while the musical recommendation remained fixed. This made it possible to observe whether explanation changed the listener's judgement rather than the music itself.

Test 2 challenged this assumption. It did not establish a single ideal level of explanation. Instead, it suggested that the usefulness of an explanation depended on the object, the listener, and the degree of control available within the interaction. Different objects called for different explanatory approaches, while different listeners preferred different amounts of detail. These observations are treated as an exploratory design finding rather than a general claim about all listeners. They indicate that explanation should remain adjustable and responsive to the relation between listener, object, and system.

Explanation did, however, provide a logic that listeners could begin to understand over time. This made the system's behaviour more legible, not because it delivered a final justification, but because it allowed the listener to recognise how the system worked and gradually develop a way of responding to it.

This led to a revised application of Stimmung. It is no longer only the uncertain orientation that exists before the musical output. It is also the space between the AI's reading and the listener's judgement. The listener should not simply receive an interpretation and decide whether to accept it. They should be able to understand how the system works and use that understanding to influence what happens next.

The agency at stake is therefore not simply the ability to see the system's reasoning. It is the ability to use that reasoning to question, revise, and redirect the system. In Curio, explanation becomes part of the mechanism through which the listener can shape the recommendation itself. This places LD5, Control over What You Hear, at the centre of the exploration. Control includes not only accepting or rejecting a finished recommendation, but influencing the logic through which it is produced.

The development of Stimmung therefore shifted the project away from asking the AI to find the correct account of the listener's desire. Instead, the listener should be given agency within the process through which that desire is explored. This raised a further design question: how can the system keep this space open without making the interaction feel arbitrary or unclear?

The Shift

The concepts now have different roles. Stimmung names the experiential condition: an emerging, relational sense of what the listener may want. Ambiguity names the design quality that keeps this condition open to interpretation. Nourishment and placation are the project's analytic labels for what the interaction does with that openness. Nourishment sustains the listener's ability to attend, question, and adjust, while placation closes uncertainty by making the result feel settled.

Test 2 therefore did not make ambiguity an unrelated new topic. It showed why the design needed it. If the meaning and useful level of explanation cannot be fixed in advance, the interaction must allow the listener to encounter, interpret, and adjust the system's reading rather than receive one final explanation. Ambiguity becomes the design condition through which the revised concept of Stimmung can remain active.

Gaver et al. (2003) treat ambiguity not as incomplete information, but as a quality of the relationship between a person and an artefact. Ambiguity can encourage users to interpret a system for themselves and develop a more personal relation to it. Productive ambiguity does not mean removing all explanation or making the system arbitrary. It requires enough structure for the listener to recognise and question a relationship while leaving it open to reinterpretation and adjustment.

Bilgic and Mooney (2005) distinguish between explanations that promote acceptance and explanations that help users judge the quality of a recommendation. Their study found that a socially persuasive explanation could cause users to overestimate the quality of an item, while explanations grounded in the item's features or in the factors that influenced the recommendation supported more accurate judgements.

This distinction clarifies two roles for explanation in Curio. Explanation as framing acts on the listener by shaping how they hear and evaluate an already selected track. It can make a weak object–music relationship appear convincing. This is close to what the project calls placation: uncertainty is settled by making the result feel correct.

Explanation as steering acts on the system's subsequent behaviour. The listener can use the AI's reading to alter the relationship between the object, the interpretation, and therefore the music produced. Agency therefore lies not simply in understanding an explanation, but in acting on that understanding, influencing what happens next.

Together, these ideas shift the interaction from placating to nourishing. A nourishing interaction does not simply produce a satisfying recommendation. It sustains attention, allows disagreement, supports adjustment, and gives the listener time to develop resonance. Nourishment is therefore a design condition that gives the listener enough time and control to develop a judgement before the interaction settles.

Resolution therefore describes the system settling the uncertain relation on the listener's behalf, while exploration includes both the physical search for an object and the interpretive space that follows the AI's reading. Later explorations will therefore test whether listeners can learn and influence the system's logic, sustain attention, and develop resonance without the AI settling the interpretation on their behalf.

4.6.1 Exploration 3 - Intro



PROVENANCE READ-OUT

//EXPLORATION INTO RECOMMENDATION

//Overview

This exploration began with a practical frustration: I was attempting to use Spotify data in my own recommendation system, but Spotify made only limited data available to developers. I initially considered replacing it with open-source alternatives such as MusicBrainz and the Million Song Dataset. However, exploring these alternatives led me to recognise that simply changing the data source would not address my deeper concern about how recommendations were formed and whose knowledge they represented. The decision to bypass platform-provided data therefore emerged through the exploration rather than being fixed at the outset. I instead began developing a recommender shaped by arguments written by enthusiasts and special-interest communities, with those sources made visible and credited.

This also returned me to a frustration with Spotify’s recommendation system. Although its recommendations may draw on the behaviour of real people, I could not see how they had been made, whose knowledge they drew upon, or what reasoning informed them. I wanted to explore a form of curated collection closer to the kind of recommendation that might emerge in a record store, rather than from an opaque model of what someone like me had played. The exploration developed from Exploration 2, which made the machine’s reasoning visible, but not the origins of its recommendations. Here, I wanted the full chain; search, source, recommendation, and playlist; to remain visible from the outset, before later explorations introduced more ambiguity.

Unlike the previous two explorations, which returned one song at a time, this exploration tested whether individual recommendations could form a playlist. Each object now produced several candidate tracks rather than a single result. This created greater variation, but also increased the number of misses: the more tracks generated for each object, the more likely it was that some would be weak matches. Some tracks did not fit the object or the rest of the playlist, while others were interesting without being appealing. The system could produce variation, but not necessarily a satisfying sequence.

The question of what counted as a good output therefore became a limiting factor; variation alone could not resolve it. This led directly to Exploration 4, in which I provided the system with more context before it made recommendations and found a more satisfactory direction.

The transparent, human-sourced approach also introduced practical limits. Gathering, connecting, and presenting sources required a more complex recommendation process, while improvements to search and recommendation increased token use and cost. Live searches were particularly slow, often taking tens of seconds or longer to produce a result. What initially created anticipation gradually became boredom, as waiting began to interrupt listening and make further experimentation difficult. At this point, continuing was no longer viable at the scale needed to compare approaches.

The exploration was nevertheless worthwhile. It exposed problems that single-song recommendations concealed, including variation, sequencing, fit, source attribution, and the influence of explanation on perceived credibility. I ended the exploration before formal testing at scale, not because the approach had no value, but because its central promise had reached a practical limit. It became a productive endpoint rather than an unfinished version. Two further consequences informed the later work. First, attempting to avoid a platform’s hidden recommendation model showed how difficult it is to access and replace the knowledge that such platforms organise. Second, deciding how much of a recommendation to disclose turned ambiguity from an accepted condition of the system into a design problem.

//Goal

The goal of this exploration was to develop a recommendation system that bypassed platform data and used credited recommendations from enthusiast communities to build traceable playlists from objects, while assessing the costs and limitations of this approach.

Left page:

Figure 4.16

Screenshot of the system readout taken during this round of testing.

4.6.2 Exploration 3 - Input

Loss Dimensions Exploration

Primary

1. LD6 Sharing Music with Others

Secondary

1. LD2 Owning the Music

Primary: LD6 Sharing Music with Others

Before streaming, music recommendation was often a situated social act. A friend might suggest an album, or someone working in a record store might recommend one. In both cases, the recommendation had a recognisable source. The listener could understand who was speaking and judge the suggestion in relation to that person's knowledge and taste.

Recommendations now also appear through Reddit threads and specialist forums. Streaming platforms draw on user activity, including listening patterns, skips, saves, and collective behaviour. However, they rarely identify the people or sources behind a recommendation. Although a suggestion may be shaped by many users, the listener cannot locate it to a particular person or source. This makes the recommendation feel impersonal.

This exploration responds by attaching each recommendation to the specific human writing from which it originated. It tests whether making the source visible can restore some of the social meaning that is lost when recommendations become detached from identifiable people.

The record store provides the clearest model for LD6. It places recommendation within a relationship between the listener and another person whose knowledge and taste can be evaluated. The listener remains free to accept or reject the suggestion, while the source of the recommendation remains visible.

Secondary: LD2 Owning the Music

LD2 forms a secondary layer of the same problem. When streaming recommendations are detached from identifiable sources, they carry less of a story about how music entered the listener's life. The resulting collection may therefore contain less personal history, because fewer recommendations retain a story of where they came from.

The system's source record partially addresses this loss by preserving the origin of a recommendation. It can show how a piece of music entered the listener's experience.. The record store connects recommendation and acquisition: someone may suggest a record, after which the listener can take the physical object away. In this case, acquisition is vested in the personal object.

// Design Implications Applied

D5 Support familiar listening while also making less visible parts of the catalogue reachable, expanding the range of what the listener encounters without resorting to random novelty.

D8 Use listening history as one source of context that informs the system's response, rather than as the factor that determines what the listener is able to encounter.

D10 Treat whatever the system infers about the listener as a provisional working representation rather than a settled basis for selection, so that it remains open to revision once listening has begun and does not permanently bound what music the listener can be offered.

D15 Keep interpretive authority with the listener and let discovery arise through material, spatial, social and chance encounters rather than through a single computed judgement.

D33 Design for repeated use, so that a listener can develop a sense of the system's tendencies over time without those tendencies hardening into a rule they can predict.

D14 Give the listener feedback on how their input has been read without presenting that reading as the one correct interpretation, so that the ambiguity remains something they can work with.

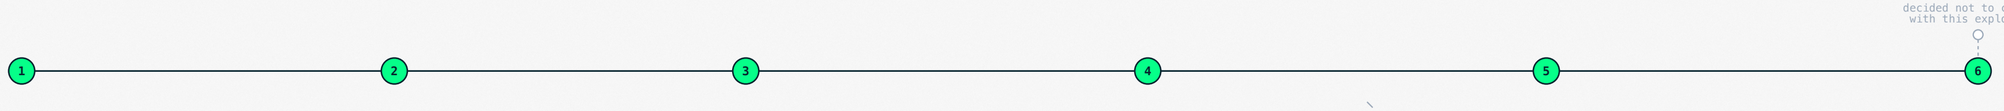
D26 Present the system's response as an association rather than a correct reading, treating the distance between what the object means to the listener and what the machine can perceive as something the interaction makes use of.

D34 Evaluate the system by whether its connections are surprising yet recognisable, whether listener agency is preserved and whether interpretation has time to develop, alongside the accuracy of what it recommends.

Newly applied

Re-applied

4.6.3 Exploration 3 - Development Process



1. Human sources turned object-based discovery into a traceable playlist

Shift. The discovery engine was changed from producing a single recommendation to producing a playlist. The prompt specified five live searches RateYourMusic, a sounds-like subreddit, Discogs lists, Bandcamp Daily, and a territory-specific source because enthusiast communities were considered to hold forms of adjacency knowledge that catalogue interfaces and similar-artist endpoints did not. Each search, source, and artist-discovery route was logged.

Finding. Source selection became central to the exploration because it kept recommendations connected to the people and communities whose knowledge had informed them. The playlist also made variation and sequencing visible: a track could be compelling alone but fail within the sequence.

software development

L6 D30

2. Varying popularity exposed the limits of search

Shift. The search was instructed to look for known, forgotten, adjacent, and new music, using both popular and less popular results to introduce variation into the playlist. This included searching for obscure recommendations in comment sections and selecting relevant entries from lower down in ranked lists, rather than relying only on the most visible results.

Finding. The search was effective at identifying prominent results, but struggled to find genuinely obscure material. It often selected the top results from a search or ranked list instead of following less visible recommendations in comment sections or choosing lower-ranked entries that made sense in context. The categories helped describe the intended variation, but did not provide a reliable way to produce it. Measures such as cultural distance, and a table based on energy and era, were removed because they did not clearly explain why a playlist failed.

software development

D5

3. An audit revealed that source attribution was not actually working

Shift. Two search architectures were considered. Live search performed new web searches for each scan, while cached search reused previously collected relationships. Live search often took more than a minute to produce a result and was more expensive, but its results were judged more interesting because it could draw on fresh and less predictable sources. Cached search was faster and cheaper, but less effective at finding obscure material.

A later code audit checked whether the system matched the written method. It found that the default cached route did not use web search when generating recommendations. The interface could therefore show no credited sources, even though it appeared to support provenance. Live search was then forced on the server whenever provenance was being tested, so the condition could not be bypassed.

Finding. Before the correction, the provenance test was not valid: the system was not reliably generating the sources it was supposed to display. The audit therefore changed the interpretation of the earlier result from an apparent design finding to an implementation failure.

software development

4. Transparent reasoning became part of the recommendation

Shift. Each recommendation was initially accompanied by an arrow chain from object cue to track. This was intended to make the system's reasoning visible.

Finding. The chain was not neutral transparency. By giving the recommendation a coherent account, it also made the recommendation more convincing, which could steer listeners towards accepting it. Greater disclosure was therefore not automatically a way of supporting agency.

theory development

D31

5. The design shifted towards preserving ambiguity

Shift. The full reasoning chain was removed from the listener's screen after testing showed that too much explanation reduced the value of the final result. The level of detail made the recommendations feel more like the output of a search query than an interesting or unexpected playlist.

Finding. Complete reasoning chains made the results less interesting and caused the sources used to justify recommendations to become repetitive. Reducing the visible explanation preserved more uncertainty and made the recommendations feel less mechanically retrieved. Transparency was therefore treated as a design choice with trade-offs, rather than as an automatically beneficial property.

theory development

D14

6. Cost and speed made provenance a productive dead end

Shift. A walkthrough assessed whether the system could be developed and used at the required scale. Web searches were the slowest and most demanding part of a pipeline that already took tens of seconds. Further testing of live and cached search was also becoming too expensive.

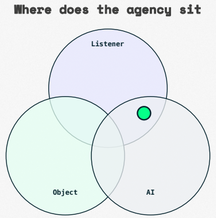
Finding. Development therefore stopped pursuing the comparison as its central question. The focus shifted towards producing a lower-cost output with less capable models, rather than continuing to investigate whether live, human-source search produced better recommendations. Provenance was stopped as a productive direction for daily use because the cost and speed of testing began to displace the core design question.

dead end

theory development

L6 D33 D34

4.6.4 Exploration 3 - Output



The listener

The listener's agency is limited to choosing and scanning an object, then receiving the resulting playlist. Source attribution makes the recommendation more legible, but it does not make the sources actionable: the listener cannot follow a source, change its influence, reject it, or request recommendations from elsewhere.

Listening history was also disabled by me and enforced at the server. This was necessary to test human-sourced recommendation, but it meant that the listener could not change the condition shaping the exploration. The system therefore disclosed more about its sources without giving the listener greater control over how those sources affected the result.

The object

The object has limited agency within the system. It initiates the interaction and provides the qualities from which the search cue is produced, but it does not determine the recommendation directly. Its personal meaning remains inaccessible to the AI, and the object cannot alter the sources, searches or playlist once the interaction begins.

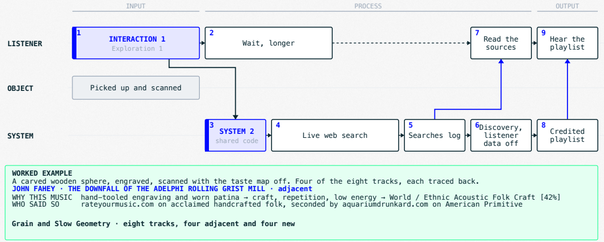
The object therefore acts as a starting condition rather than an active participant. It influences what the system searches for, but the main interpretive and selection work is carried out by the AI.

The AI

The AI holds most of the system's operational agency. It translates the scanned object into a written cue, selects and queries external sources, interprets the returned material, and assembles the playlist. It also determines which parts of this process are shown to the listener.

The system initially displayed a full reasoning chain, but this made recommendations more persuasive without giving the listener more control. When live search was not correctly enforced, the system could also appear to provide provenance without actually generating new sources. The later audit showed that agency depended not only on the interface, but on which search architecture was active.

This exploration therefore produced an uneven distribution of agency: the listener chose the starting object, the object constrained the initial cue, and the AI controlled interpretation, search, selection and explanation. Source disclosure made the AI's agency more visible, but did not redistribute it.



// Key insights

1. Attribution changed how a recommendation was received more than what it recommended. Naming the human source did not make the results more accurate, but it made them feel more accountable and made some weak recommendations easier to stay with. The source provided standing rather than evidence of quality. This partially restored the social character of record-store recommendation, but did not create actual sharing, exchange or ownership.

2. Knowing where a recommendation came from built trust, but complete transparency shifted attention from the music to the system. A source label helped the listener understand the result, while a full reasoning chain invited criticism at every stage and encouraged the explanation to do the persuading. The design problem is therefore not how much the system can reveal, but what it should reveal before the system itself becomes the subject of attention.

3. The provenance trail already existed in the system's records, but it was not yet an interaction. Searches, sources and artist-discovery routes were logged, then reduced in the interface to a fixed list. A more useful trail may need to be navigable, allowing the listener to follow or weight particular sources rather than simply receive more explanation.

4. Listening history acted as an unstable design input rather than neutral personalisation. When used, it tended to produce results that were either safe and familiar or unexpectedly disconnected. Removing it widened the basis of recommendation, but did not establish that history should always be excluded. It suggests instead that personal data should be provisional, inspectable and open to listener control.

5. Producing a playlist exposed problems that single-track recommendation concealed. More tracks created variation, but also increased weak matches and sequencing failures. Instructions to find known, forgotten, adjacent and new music did not reliably produce genuinely obscure material. Variation alone was therefore not enough; the system also needed to produce a sequence with sufficient fit between the object, the tracks and the listener's developing attention.

6. Trying to bypass Spotify showed that the platform's limitations should become part of the scope rather than only an obstacle. Reconstructing the knowledge organised by the platform required too much time, money and infrastructure for one person to sustain. The next exploration should therefore work within Spotify's constraints and ask what forms of source visibility, listener control and recommendation change are possible without rebuilding the platform's data environment.

// Key questions

1. How can provenance become a navigable form of listener agency rather than a fixed explanation? What should the listener be able to follow or influence, and how can this happen without making the system's logic more prominent than the music?

   Answered in Exploration 6

2. How can the social aspects of music recommendation be explored locally rather than digitally? Can recommendation be situated in shared spaces, objects or encounters with other people instead of being represented only through online sources and written attribution?

   Answered in Exploration 4

Left page:
Figure 4.17

Diagram showing where agency sits by the end of this exploration.

Left page:
Figure 4.18

Diagram showing where how nourishing the experience was for the listener by the end of this exploration.

Left page:
Figure 4.19

Diagram showing the system architecture developed by the end of this exploration.

4.7.1 Exploration 4 - Intro



PERSONA COINS

//EXPLORATION INTO DATA AGENCY

//Overview

This is the only exploration in which I designed the input object itself. Earlier explorations attempted to translate personal objects from the physical world into usable data. Here, I reverse that process: I convert data from the listener's Spotify history into an object that can be manipulated, with the aim of giving the listener greater agency over how that data is interpreted and used.

The idea developed from a frustration with recommendation systems, which often struggle to make effective recommendations for listeners with varied tastes. Previous explorations showed me that a recommendation is more likely to be valued as a discovery when it departs from familiar music while retaining a recognisable connection to the listener's existing taste. Unexpected recommendations are only useful when they feel intentional rather than random. The system must therefore develop a clear understanding of the listener's taste before it can make a meaningful deviation from it.

I initially developed the listener's taste map as an on-screen interface. The taste map scans the listener's liked songs on Spotify and organises them into nine regions representing different areas of their listening history. These regions are arranged like a musical solar system: regions closer to the centre represent the listener's current interests, while regions further towards the edge represent music they enjoyed in the past but listen to less frequently now. The map therefore makes an otherwise hidden and constantly changing taste profile visible and easier to understand.

As development progressed, I became aware of limitations in keeping this model on screen. The act of clicking through the regions offered only a limited form of interaction. It did not encourage sustained experimentation with the system over time and created a barrier to repeatedly engaging with it. I therefore wanted the model to become more than a visualisation presented by the system. I wanted the listener to interact with it more frequently and to take a more active role in shaping the recommendations it produced.

I therefore manifested the digital taste map as a physical object: a set of persona coins, each representing one region of the listener's taste history. This physical manifestation made the listener's musical history feel tangible. Rather than existing only as an abstract profile on screen, it became a collection that could be built, handled, and expanded. Diversifying one's taste could also become rewarding, since discovering music beyond familiar areas could give the listener something new to add to their collection.

Placing a coin makes that region available to the recommendation system, while leaving it out withholds that region from the recommendation input. Combining several coins allows the listener to bring different parts of their taste together and create a temporary profile for each recommendation. The physical interaction gives the listener greater control over how the system interprets their data, making them a shared author of the recommendation outcome rather than a passive recipient of it. Because the coins are physical, they can also be lent to another listener, shared between listeners, or used by more than one person at a time. This makes the taste profile visible, adjustable, and shareable.

//Goal

To give listeners greater agency over how their personal data is interpreted and used by the recommendation system, and to examine how this affects the recommendations it produces. I

Left page:
Figure 4.20

Image taken during development of this exploration.

4.7.2 Exploration 4 - Input

Loss Dimensions Exploration

Primary

1. **LD5: Control over What You Hear**
2. **LD1: Handling the Music**
3. **LD6: Sharing Music with Others**

Secondary

1. **LD2: Owning the Music**

LD5: Control over What You Hear

In the record store, searching for music involved a self-directed journey. Discovery could follow personal curiosity and judgment. Suggestions from staff, friends, or other listeners could inform the process, but the listener remained responsible for the music they encountered.

The exploration aims to address this loss by giving the listener control over the input to a recommendation rather than its final output. The system still selects what plays, but the listener can decide which parts of their listening history influence that decision. Exploration 3 showed that removing the listener's history could produce interesting but arbitrary results. The coin keeps the history available while allowing the listener to choose which parts are active.

LD1: Handling the Music

The interaction with a record was an intentional and tactile relationship. Selecting a record, handling it, and preparing it for to play required attention. Furthermore the system only new what the user supplied two it allowing users to feel agency over vast a vast and varying collection.

The exploration aims to address this loss by giving part of the listener's digital history a physical form. The listener can hold the coin, place it, arrange it, remove it, or withhold it from the system. This transfers the deliberate handling of a record to the information shaping the recommendation. Unlike a record, however, the coin does not reveal what music it will produce, and the system's interpretation remains hidden. The exploration therefore tests whether tactile engagement can still matter when the relationship between the object and the music is opaque.

LD6: Sharing Music with Others

People once explored music collaboratively in record stores or at home, browsing physical collections together. Because these collections were limited by space, discovery was shaped by what was available in the shared setting, making it both personal and social.

Digital libraries remove these spatial limits, but they can make listening histories less tangible and harder to share intentionally. This exploration addresses that loss by turning part of a listener's recommendation history into a portable, shareable object: a coin that can be passed to another person and incorporated into their recommendations. In this way, multiple users can engage with and contribute to the exploration process.

LD2: Owning the Music

Records offered a way to make a person's musical taste physical and therefore visible. Taste is inherently internal and difficult to externalise; it becomes legible to others through the objects, choices, and arrangements that represent it. A record collection could make this personal history visible as something that could be seen, handled, revisited, and shared.

Ownership is therefore not only about possession. It also involves having an accessible and stable record that a person can return to and choose when to engage with. A collection can become part of someone's identity because it reflects their individual taste, interests, and history.

The persona coins explore this aspect of ownership. Each coin gives a name and physical form to a particular part of a listener's musical identity. Because personal taste differs from person to person, the coins can become unique objects that reflect something specific about their owner. They do not restore ownership of the music or recommendation data itself, but they offer a more tangible way to recognise, arrange, and share a personal history of listening.

// Design Implications Applied

D4 Treat each skip, save or repeated play as a provisional signal rather than as evidence of durable taste, and give the listener a way to correct or reverse the personalisation that results.

D9 Preserve the listener's effort and agency in finding music, leaving room for choosing, comparing, interpreting and dwelling rather than reducing discovery to passive delivery.

D16 Preserve the physical actions that slow the interaction enough for attention to form, allowing the listener to linger over a result, compare it with another, backtrack or change direction at their own pace.

D27 Allow the listener to choose an object instinctively without having to justify it, and to reject a result, retry or choose another object without that mismatch being treated as a failure.

D5 Support familiar listening while also making less visible parts of the catalogue reachable, expanding the range of what the listener encounters without resorting to random novelty.

D8 Use listening history as one source of context that informs the system's response, rather than as the factor that determines what the listener is able to encounter.

D10 Treat whatever the system infers about the listener as a provisional working representation rather than a settled basis for selection, so that it remains open to revision once listening has begun and does not permanently bound what music the listener can be offered.

D13 Support gradual and reversible tuning of the result, so that the listener can adjust their direction over several moves rather than receiving a single recommendation to accept or reject.

D14 Give the listener feedback on how their input has been read without presenting that reading as the one correct interpretation, so that the ambiguity remains something they can work with.

D33 Design for repeated use, so that a listener can develop a sense of the system's tendencies over time without those tendencies hardening into a rule they can predict.

● Newly applied
● Re-applied

4.7.3 Exploration 4 - Development Process

1 — 2 — 3 — 4 — 5 — 6 — 7

1. Liked songs became the basis of the listener model

Shift. Spotify did not provide access to a complete listening history, but it did provide the listener's saved or liked songs. I used this as the starting point for the model. Unlike listening history, liked songs exclude music that may have been played casually, repeatedly by someone else, or left running in the background. They form a collection of music that the listener had, at some point, chosen to keep.

Finding. Knowing that recommendations were grounded in my own saved music increased my confidence in the system. The data source was limited, but liked songs proved an interesting backbone for a recommendation system and supplied the first stage of agency to the system as these songs require intentional selection

software development

D8

2. A fixed snapshot became a renewable collection

Shift. The liked-song library produced a timely snapshot of the listener's taste rather than a continuously updated profile. Initially, I treated this as a weakness: the map would become less representative as listening habits changed, requiring the listener to update it over time.

Finding. During development, this limitation became valuable. Updating the map did not feel like correcting an outdated system; it felt like adding a new part to a personal collection. When a new persona emerged, it marked a change in taste that could be acknowledged and kept.

theory development

L2 D10

3. Liked songs were grouped into an on-screen persona map

Shift. I used the liked-song history to form groupings of music, positioned as personas on a map. This made the listener model visible and interactive rather than leaving it as an invisible part of the recommendation system. Initially, the map was intended to remain on screen as a way to improve recommendations.

Finding. Seeing the categories defined in this way changed my relationship with the system. Rather than accepting the full map as a fixed profile, I wanted to decide which parts of it should influence a recommendation.

software development

4. Tick-box selection revealed the limits of a screen-based interaction

Shift. I first enabled this choice through tick boxes, allowing personas to be included or excluded after the map had been generated. Although this made the listener model adjustable, the interaction felt limited. The choice could only be made once, which created pressure to select the "right" configuration and reduced the sense of exploration.

Finding. The interaction needed to support temporary, reversible choices rather than a single configuration decision. This led to the shift from on-screen selection to physical coins.

dead end

theory development

D13

5. Coins were intended to make persona composition physical

Shift. I designed a board with four NFC-reader slots so that several persona coins could be active at once. Each coin represented a territory within the listener's map, allowing the listener to combine and withhold parts of their taste while the object continued to provide the immediate recommendation input. Hardware limitations meant that the final prototype used only one physical reader: the remaining coins were managed on screen. I also explored a flip interaction to reset the coin or to use it as a negative input ie. play the opposite to (X) but hardware limitations made realisation difficult.

Finding. The intended interaction was physical composition, rather than profile switching. The delivered prototype demonstrated this as a hybrid interaction: one coin could be physically loaded at a time, while broader composition meaning that the broader composition remained screen-based.

interaction development

hardware development

L1 D3

6. Deterministic weighting made coin composition testable

Shift. I built a listener-side weighting module that mapped each coin to taste-map territories through fixed genre and artist affinity rules, without a model call. The same coin configuration therefore always produced the same active territory set. With no coins, the system used the full map; one coin narrowed it to its associated territories; and additional coins widened or redirected the active map before recommendation.

Finding. Deterministic blending made the coins' effects inspectable. Any change in a recommendation could be linked to the chosen configuration, rather than to variation in the generative model.

software development

7. Making the data physical introduced new questions about access and agency

Shift. Turning personas into coins created an interesting new interaction, but it also made the system more convoluted. It required external objects: how would a listener receive their initial coins, and how would they obtain new ones when their taste map developed? The focus of experimentation also shifted. Instead of trying different personal objects, the user began experimenting with combinations of coins.

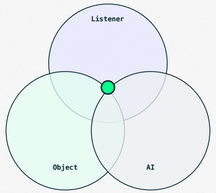
Finding. The coins made the listener's data tangible and gave them agency over it, but this should not displace the project's core interaction with personal objects. The tension between data agency and interactional simplicity is developed further in Exploration 6.

theory development

L6 D25

4.7.4 Exploration 4 - Output

Where does the agency sit





//The listener

The listener's agency begins with the data used to construct the taste model. Rather than using the complete listening history, the system draws from liked songs: music the listener has actively chosen to save. This makes the initial dataset more intentional than a record of every play, which could include background listening, accidental plays or music selected by someone else.

The listener then chooses which parts of this dataset should influence the recommendation. Selecting a persona coin activates one region of the taste map, while combining or withholding coins changes the profile available to the system. The listener therefore shapes the conditions of the recommendation without selecting its final result. Agency is moved from accepting or rejecting an outcome to deciding what the system is allowed to interpret.

//The object

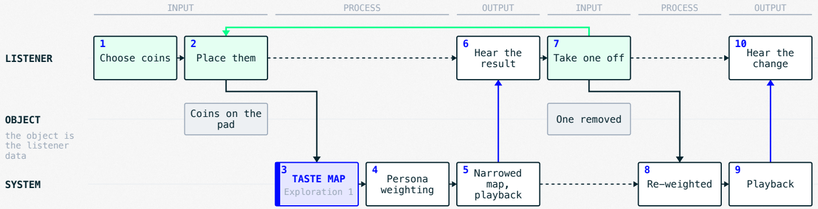
The persona coin gives physical form to a part of the listener's liked-song data. It does not contain a fixed record or determine what will be played. Instead, it makes one area of the listener's taste available to the AI.

The object therefore functions as a means of selecting data rather than as a source of meaning. Placing, combining or removing coins allows the listener to compose a temporary taste profile. This makes the listener model visible and adjustable, while keeping the relationship between the selected data and the final recommendation open.

The coin also changes what it means to handle music. The listener is not handling the music itself, as they would with a record, but the information used to find it. This gives physical form to an otherwise hidden part of the recommendation process.

//The AI

The AI determines how the listener's selected data becomes a recommendation. Each coin is mapped to specific taste-map territories through fixed genre and artist affinities. The same coin configuration therefore produces the same active territories, making the effect of the listener's input more visible and testable.



Key insights

1. Using liked songs rather than the complete listening history gives the listener agency before the recommendation is generated, because the starting dataset is made up of music they have actively chosen to keep. A complete listening history may include incidental, passive or accidental plays, whereas liked songs provide a more intentional account of taste. This suggests that personalisation should distinguish between behaviour that was merely recorded and music that the listener deliberately claimed as part of their collection.

2. Treating a listener's collected songs as a set of bounded personas changes musical data from an invisible profile into a collection of distinct objects. Each persona gives a recognisable part of the listener's taste a physical form, making the data feel tangible from the outset. Discovering music beyond familiar areas can then produce a new persona to add to the collection, giving exploration a reward beyond the recommendation itself.

3. Selecting parts of the taste model allows the system to move beyond familiar music without losing the connection that makes a recommendation feel intentional. Removing the listener's history altogether can produce interesting but arbitrary results. Persona selection keeps the history available while allowing the listener to choose which areas should guide the next recommendation.

4. The move from tick-box selection to physical coins showed that the form of a control affects the kind of agency it supports. Tick boxes allowed the listener to adjust the profile, but made the decision feel fixed and encouraged them to search for the correct configuration. Physical composition suggested a more temporary and exploratory relationship, in which personas could be combined, withheld and revisited.

5. Making personal data physical can make personalisation more attentive, but it can also compete with the personal object that gives the interaction its associative and nourishing qualities. The coins make the hidden data behind a recommendation available for handling, but they risk becoming the main subject of the experience. Future development should therefore use data objects to support the personal object rather than replace it.

6. Allowing each persona to carry equal weight in the recommendation strengthened the listener's agency over the system. No single area of taste could dominate the result, which made it easier for less prominent or more niche interests to influence the recommendation. Combining personas therefore became a way to redirect the system towards overlooked parts of the listener's taste rather than reinforcing only their most frequent listening.

//Key questions

1. How can control over the listener's data be introduced without distracting from the established, nourishing benefits of the personal object as the input?

Answered in Provocation 6

2. How can persona coins be updated alongside changes in the listener's liked-song data, and how should each coin represent its associated music so that the listener can distinguish one persona from another without consulting the system?

Answered in Provocations

3. Is the persona the best way to divide a listener's musical taste? The current model makes taste tangible and manageable, but it may impose fixed categories on interests that overlap, change or resist clear separation.

Open Question

Left page:
Figure 4.21

Diagram showing where agency sits by the end of this exploration.

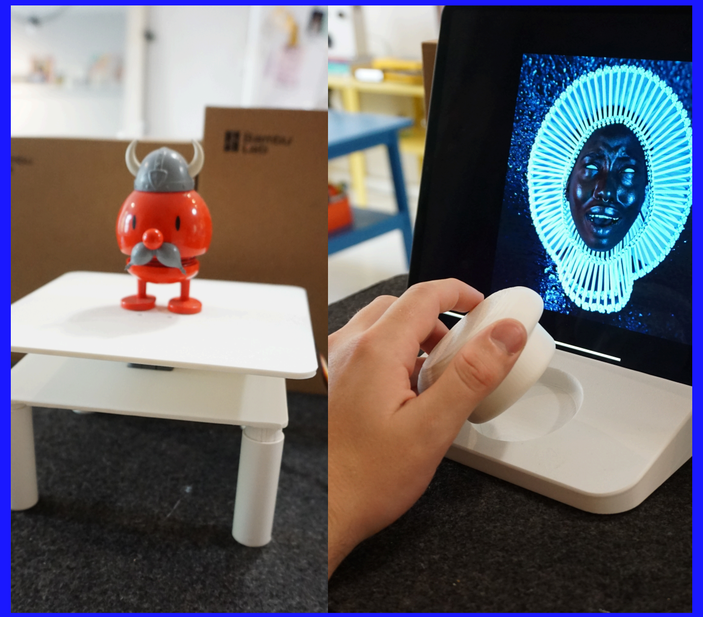
Left page:
Figure 4.22

Diagram showing where how nourishing the experience was for the listener by the end of this exploration.

Left page:
Figure 4.23

Diagram showing the system architecture developed by the end of this exploration.

4.8.1 Exploration 5 - Intro



COMMITTING TO THE ALBUM

//EXPLORATION INTO PASSIVE LISTENING

//Overview

This is the only exploration in the project where the unit of music is the album rather than the song. It combines two related probes to test how changing the output from a song to an album affects the interaction.

The weight probe reframed the object's input mechanism. Earlier explorations interpreted objects through several material qualities; here, the system read only weight. This shifted more interpretive and directional work to the listener, who could choose, compare, and adjust objects to control the direction of exploration. Weight was something listeners could literally judge, but its exact value remained ambiguous. The interaction therefore began generally and was refined through experimentation with different objects. A greater weight caused the system to look further into the listener's past listening, while a lighter object surfaced more recent songs. Developed from Exploration 4's listening-history feature, weight became a way to search through that history.

The probe tested how physical objects could guide recommendations, while also exploring whether the complexity that had built up in the system was necessary. Earlier object readings involved eight different aspects; reducing the input to weight made the relationship between the object and recommendation more direct.

The interaction was initially active, requiring listeners to move around their homes, experiment with objects, and compare how their weight affected the recommendations. Once a promising direction had been established, the interaction could become more passive: listeners could place an object down, make adjustments, and leave the system to run. This was initially engaging, but the object eventually felt largely unimportant because the resulting interaction was less interesting than in previous explorations.

Using albums emerged from the time invested in the initial input. If listeners spent time experimenting with objects and refining a direction, the result needed to offer a greater payoff than a single song. Selecting a whole release gave the gesture more consequence and reduced the desire to keep changing direction once a promising result had been found.

The lift probe developed from this shift. Once weight had been used to refine the direction, the system would slowly cycle through recommended album artwork based on the previous object scan and the listener's listening history from Exploration 4. Rather than continually searching, the listener could wait for an album that felt worth choosing. Lifting the "stopper" would start the album currently on screen, while replacing it would return the system to browsing. This created a clear transition from actively tuning the recommendations to passively browsing and committing to an album.

Together, the probes test how the listener moves from active physical exploration to slower browsing and selection. The weight probe tests gradual guidance through the listener's listening history, while the lift probe tests whether, after tuning the direction, they are willing to pause the search and commit to an album. The shift from songs to albums tests whether a more substantial output gives the gesture, and the time spent tuning it, greater value.

//Goal

Explore the effect of putting more interpretive agency on the listener and how this changes their interaction with the object and the system.

Left page:
Figure 4.24

Two images taken during development of this exploration.

left image = weight as an input.

right image = lift the stopper to play the music.

4.8.2 Exploration 5 - Input

Loss Dimensions Exploration

Primary

1. **LD4: Listening to the Whole Album**

Secondary

1. **LD1: Handling the Music**
2. **LD3: Giving Music Your Full Attention**
3. **LD7: Having a Space for Listening**

Together, these dimensions are translated in different ways. LD4 is the main aim: making the album something the listener commits to. LD1 gives this commitment a physical gesture. LD3 supports a slower process of tuning and listening. Loss Dimension 7 gives listening a physical place in the home.

LD4: Listening to the Whole Album

From LD4, I took the idea of listening to an album as a whole, rather than as separate songs. In a record store, records are chosen as complete objects and played in order. Streaming often makes it easier to choose and change individual songs.

With a record and record player, playing an album takes time. You have to choose the record, put it on, and start it. This makes it harder to change your mind during listening. You also cannot easily jump between songs, because their position on the record is not always obvious. This encourages commitment to the whole album.

I tried to reflect this through the lift-to-play interaction. Once the listener lifts the stopper, the displayed album begins playing as a whole. They cannot skip tracks; they can only play, go back, or search again. The lift is intended to make choosing an album feel like a physical commitment.

LD1: Handling the Music

From LD1, I took the physical process of browsing, choosing, and preparing music to play. In a record store, records are handled, removed from shelves, carried, and prepared before they can be heard. Streaming makes this process almost instant and easy to reverse.

The weight interaction gives the listener a physical way to adjust the system before playback. The lift gesture turns the moment of choosing a record into the act of removing the stopper and committing to the album on display.

Because the physical sensor was never connected, this remains a design idea rather than a tested feature.

LD3: Giving Music Your Full Attention

From LD3, I took the idea of spending time with music before and during playback. Record stores allow time for browsing and choosing, while a record player extends this through the preparation and length of the album.

This exploration creates a slower move from searching to listening. The listener can make small or large adjustments, find a position that feels right, and then leave the system to browse or play. Playing the full album supports listening beyond the first track.

This does not prove that the listener gives the album their full attention. It creates conditions that may support more focused and longer listening.

LD 7: Having a Space for Listening

Having a space for listening meant that music had a clear place in the home. A room could support particular routines, positions, and ways of paying attention. Listening became an activity that someone entered, rather than music simply filling any available space. The space also made listening visible as part of everyday life.

The lift gesture requires the listener to physically move to one place to use the system, similar to using a record player. This gives the interaction a designated home. It also creates physical distance between the listener and the ability to immediately change the music, which may make it easier to stay with the chosen album.

// Design Implications Applied

D3 Keep the listener's control available throughout the session rather than only at its start, so that they can pause, redirect, compare alternative directions or revise the exploration after playback has begun.

D19 Let an object's handling, weight, texture and manipulation contribute to how the listener interprets what they hear, without encoding a fixed correspondence between particular objects and particular music.

D32 Keep the association recognisably connected to the object while leaving enough distance for the listener to question it, so that results are neither repetitive nor arbitrary.

D5 Support familiar listening while also making less visible parts of the catalogue reachable, expanding the range of what the listener encounters without resorting to random novelty.

D16 Preserve the physical actions that slow the interaction enough for attention to form, allowing the listener to linger over a result, compare it with another, backtrack or change direction at their own pace.

D29 Give each result time to unfold, allowing space for listening and reflection rather than presenting recommendations as a rapid feed.

D33 Design for repeated use, so that a listener can develop a sense of the system's tendencies over time without those tendencies hardening into a rule they can predict.

Newly applied

Re-applied

4.8.3 Exploration 5 - Development Process

1 — 2 — 3 — 4 — 5 — 6

1. Briefly testing memory-based object meaning

Shift. I briefly experimented with associating objects with personal memories. The listener described a memory, and I used a chatbot to translate it into a musical direction. This produced a more solid and specific desired output for the listener than the recommendation AI did.

Finding. This approach became a dead end because it required too much declarative work. The listener had to explain what the object meant before the system could respond, making the input too fixed and reducing the openness of the physical interaction.

dead end
theory development
D27

2. Choosing weight as a physical input

Shift. I assessed several physical aspects of objects but chose weight because it could be measured definitively by the system. At the same time, weight could only be known by the listener through physically picking up the object. Exploration 1 had shown that this act of handling and judging an object was central to the interaction.

To test the MVP, I shifted to using kitchen scales as a simple stand-in for the sensing system. This allowed me to explore how different weights could guide recommendations before developing a dedicated physical mechanism. A heavier object was intended to move further into the listener's past listening, while a lighter object surfaced more recent music.

Finding. The object could guide the recommendation without determining its meaning. However, exact weight was difficult to know, so the interaction began generally and was refined through comparing and experimenting with different objects.

theory development
L1 D19

3. Using weight to guide exploration

Shift. Weight became a way to move through the listener's listening history. A heavier object was intended to take the system further into the listener's past listening, while a lighter object surfaced more recent music. This allowed listeners to control the direction of the exploration through small or large adjustments, while keeping control available throughout the session.

The early interaction was active. Listeners had to move around their homes, handle objects, and compare how their weight affected the direction of the recommendations. Once a direction had been refined, the interaction could become slower and more passive, allowing the system to browse while the listener waited, listened, backtracked, or changed direction.

Finding. Weight became useful when its relationship to the music was sufficiently legible for the listener to learn through repeated use. When the relationship was unclear, the interaction felt arbitrary. The proposed mapping was therefore kept recognisable but open, rather than assigning particular objects to particular kinds of music.

interaction development
L7 D3 D32

4. Changing the output from a song to an album

Shift. The time invested in handling objects and refining the direction made a single-song result feel insufficient. I therefore changed the output from a song to an album. A whole release offered a more substantial outcome and made the listener's input feel more worthwhile.

This drew on the idea of listening to an album as a complete object. Once playback began, the listener could not skip between tracks. They could play the album, go back, or search again.

Finding. The album became the unit of commitment rather than the individual track. This gave the process of tuning the recommendation a clearer purpose and greater consequence.

theory development
L4

5. The introduction of lifting a "stopper" to as a commitment to play music

Shift. The weight interaction allowed the listener to refine the direction of the exploration. The lift interaction then provided a separate moment for stopping the search and committing to a result. Album artwork slowly cycled while the listener browsed. Lifting the stopper started the album currently displayed, while replacing it returned the system to browsing.

This separated two actions: weight was used to guide the exploration, while lifting was used to commit to listening. It also gave the interaction a physical place in the home and made choosing music resemble the handling and preparation involved in playing a record.

Finding. The lift created a clear transition from active searching to slower browsing and listening. It supported the idea of giving an album time to unfold, although it could not prove that the listener maintained their full attention.

interaction development
L3 D16

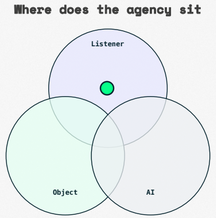
6. Prototyping the proposition and identifying its limits

Shift. I removed weight as a direct input and instead connected the system to the object-reading code developed in the previous exploration. This code generated album suggestions from the object's reading. I then experimented with how long each recommendation remained on screen, allowing recommendations based on older personas to linger for longer. The physical scale and lift sensor were not connected, so this tested the timing and logic of the browsing experience rather than the complete physical system.

Finding. Varying the dwell time made the interaction more engaging. Longer-lasting recommendations gave the listener time to notice patterns and gradually understand how the system worked through repeated use. This created a slow conversation between the listener and the AI: the system revealed its tendencies over time, while the listener learned how to interpret and respond to them.

software development
D29 D33

4.8.4 Exploration 5 - Output



//The listener

The listener holds interpretive agency in the weight interaction. Rather than describing what an object means, they physically weigh and compare objects, using their own judgement to steer the recommendation system. A heavier object moves further into past listening, while a lighter object surfaces more recent music. The listener therefore interprets the relationship between the object and the recommendation through repeated experimentation.

The stopper gives the listener complete agency over whether playback begins. They can allow albums to continue cycling or lift the stopper to start the album currently displayed. Once playback begins, the album plays as a whole rather than as a sequence of individually selected tracks.

//The object

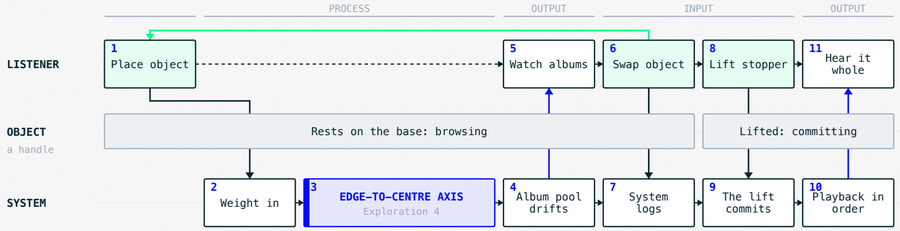
The object has less agency in this exploration than in previous ones because most of its material qualities are removed from the interaction. The system reads only weight; its shape, texture, colour, and personal associations do not affect the recommendation.

Weight can influence how far the system searches through listening history, but it does not express a particular musical meaning. The object therefore acts as a parameter for steering rather than as an interpretive participant. Its role is to make adjustment physical while leaving the listener to decide how its effect should be understood.

// The AI

The AI holds operational agency but not interpretive agency in the weight interaction. Because the input is a factual measurement, the system does not need to interpret what the object means. It applies a predefined relationship: heavier objects search further into the listener's past, while lighter objects search more recently.

The AI therefore calculates and presents recommendations, but the listener remains responsible for interpreting whether the result is useful and how to adjust the next input. In the later prototype, object-reading code reintroduced more interpretive work for the AI, while longer dwell times allowed the listener to learn its tendencies. The weight system itself, however, keeps interpretation with the listener rather than the AI.



//Key Insights

1. Asking users to interpret and adjust through physical action gives them more agency than asking them to state a preference directly. Uncertainty becomes part of a nourishing interaction, provided they can learn how their actions affect the result.

2. Simplifying an input can make a system easier to understand and control, but it can also reduce the expressive agency of the object. Future designs need to balance clarity with the richness of material qualities.

3. Requiring the participant to physically move towards a system can support both passive and active engagement. The system can remain available while they are elsewhere, then become an active encounter when they choose to approach, adjust, and commit.

4. The outcome needs to match the time and effort invested in reaching it. When participants spend longer interpreting, adjusting, or waiting, a more substantial result can make that effort feel purposeful.

5. Slowness can help participants build an understanding of how a system behaves. Waiting is not empty when the system's responses allow people to recognise patterns and develop their own way of reading it.

//Key questions

1. How can a design open opportunities for slow interpretation of the systems logic while still being legible enough for participants to influence its direction in the short term?
   Open Question
2. Should future versions keep interpretation with the listener through factual physical inputs, or shift it towards the AI through richer readings of the object?
   Answered in Provocations

Left page:
Figure 4.25

Diagram showing where agency sits by the end of this exploration.

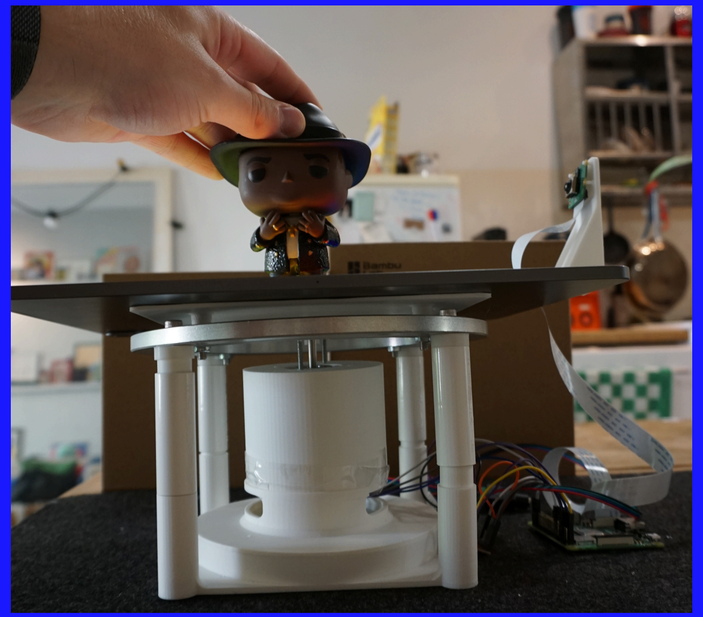
Left page:
Figure 4.26

Diagram showing where how nourishing the experience was for the listener by the end of this exploration.

Left page:
Figure 4.27

Diagram showing the system architecture developed by the end of this exploration.

4.9.1 Exploration 6 - Intro



VIBE MAP

//EXPLORATION INTO AMBIGUITY AND STIMMUNG

// Overview

This is the only exploration in which the object changes its role after being read. Elsewhere, the system reads the object and returns a result; here, it reads the object once, and the object then becomes the instrument the listener uses to navigate the recommendation space.

The exploration began as a development of learnings from previous explorations: it became apparent that the system needed to hand agency back to the listener once it had produced its first response, letting the listener dictate the direction of the exploration that followed. In the earlier system, an object went in, the machine searched on the listener's behalf, and an output came out. The physical object gave the appearance of agency, but it did not let the listener take part in the path through the recommendations. The result was a trigger, not navigation.

The answer was to stop handing over a destination and lay the recommendation process out as a navigable field, generated from the eight aspects of the scanned object. This formed a "vibe map": an area where the listener could explore the machine's associative logic between musical output and the readings of the eight aspects, rather than a display of their own taste. The taste map stayed inside the recommendation system, shaping which music was returned from a given position in the field.

The field held no labels or named destinations. Here, inspired by the card game Dixit, ambiguous visual and musical changes allowed the listener to move through the space and shift how much weight the system gave to one of the 8 aspects read from the object, steering which reading it attended to without selecting an aspect directly or being told what the object meant.

This extended the project's work on Stimmung and agency: Stimmung as the emerging, relational sense of what the listener may want, and ambiguity as what keeps that relation open to adjustment. Agency here lay in shaping the relationship between object, reading, and music without ever being handed a final interpretation.

Conceptually, this let the listener travel through the space rather than receive a destination. The exploration was intended to investigate how the AI and the listener could collaborate in sense making, with the listener able to steer the exploration journey after establishing its starting point.

//Goal

Explore how the AI and the listener can collaborate in sense making, allowing the listener to steer their exploration journey after establishing the starting point.

//Caveat

This exploration was never fully realised, held back by the time and cost required to realise it rather than by the idea itself. However, the attempt to build it forced the concept to be fully theorised: the mechanism, the interaction logic, and the reasoning behind withholding a destination were all worked through in detail.

Left page:
figure 4.28

Image taken of the physical rig developed during this exploration.

4.9.2 Exploration 6 - Input

Loss Dimensions Exploration

Primary

1. **LD9: Exploring Music in Physical Space**

Secondary

1. **LD5: Control over What You Hear**
2. **LD3: Giving Music Your Full Attention**

LD9: Exploring Music in Physical Space (Primary)

A record store offers a series of small, contained discovery moments. A visual cue, such as a genre divider, handwritten sign, or album cover displayed at eye level, draws the listener towards one section. Then from this the user has an even smaller space to narrow there search down into when then digging through the specific crate chosen. This builds on the house-to-record-store comparison from Exploration 1, but focuses specifically on the structure of searching through a crate rather than browsing an open space. The vibe map works in a similar way. It is a limited unlabelled field where users explore the visual cues of the object chosen.

LD5: Control over What You Hear (Secondary)

Vinyl listening depends on deliberate action. Nothing plays or changes until the listener chooses to make it happen, whether by lifting the arm, pressing a button, or turning the record over. The exploration carries this sense of control into the scanned object. Instead of using a separate input, the listener handles the object they have already chosen to move through the vibe map. Its position and movement influence which aspect is emphasised and which direction the music takes. Control therefore remains connected to a physical, felt action rather than a touch screen based menu selection.

LD3: Giving Music Your Full Attention (Secondary)

Listening to an album often involves knowing that another track is coming without being shown exactly what it is. The listener relies on the album's sequence and their own memory, which creates a balance between anticipation and surprise. The exploration uses a similar approach by avoiding text-based information. The listener has to hear the music to work out what it is before deciding whether to continue in that direction or adjust course. The object, the system, and the listener respond to one another. The listener learns through listening, then acts on that understanding. This keeps attention focused on the music rather than on reading a written queue.

// Design Implications Applied

D17 Let the listener encounter the object and hear the music before any artist, genre or mood description frames it, and make that information available afterwards so it extends the encounter rather than preceding it.

D20 Let the listener's interpretations accumulate and shape the musical possibilities that follow, so that ambiguity opens a path through the catalogue rather than producing unrelated results.

D31 Present each result as a provisional reading that the listener can question, reinterpret or redirect, so that responding to the interpretation does not require rejecting the track it produced.

D8 Use listening history as one source of context that informs the system's response, rather than as the factor that determines what the listener is able to encounter.

D11 Treat the listener's not yet knowing what they want as the condition the interaction begins from, rather than as a gap the system must close before it can respond.

D15 Keep interpretive authority with the listener and let discovery arise through material, spatial, social and chance encounters rather than through a single computed judgement.

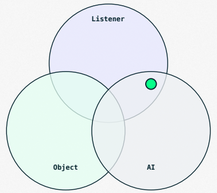
D19 Let an object's handling, weight, texture and manipulation contribute to how the listener interprets what they hear, without encoding a fixed correspondence between particular objects and particular music.

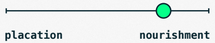
D24 Let the object's perceived qualities shape the music that is returned, so that it contributes to the recommendation rather than only triggering one.

● **Newly applied**
● **Re-applied**

4.9.3 Exploration 6 - Development Process

Where does the agency sit





Left page:
Figure 4.29

Diagram showing where agency sits by the end of this exploration.

Left page:
Figure 4.30

Diagram showing where how nourishing the experience was for the listener by the end of this exploration.

//What was built

I specified and assembled the hardware for this exploration: a single-board computer, a camera module, a stepper motor, a motor driver and a magnetic encoder. A second camera, positioned above the object, read its position on the field.

Assembling these parts suggested a use for the motor beyond rotating the object during scanning. It could also give the listener something to push against. I assigned that resistance an action: saving a track. This gave the system an output that lasted beyond the session, allowing what the listener discovered to accumulate into a collection rather than exist only while it was playing.

Integrating the components into a system that responded quickly enough for movement to feel like call and response rather than a laborious request took longer than the exploration allowed. I stopped at that point. The rig, the saving mechanism and the field are carried forward in Provotype 4 as a designed future rather than a working device.

//Building the field

The rest of the exploration focused on the "Vibe Map", and this is where most of the thinking took place.

The exploration had already established that the destination should be withheld: the system would read the object once, produce an abstract starting point, and then hand over navigation to the listener. This left the question of how the listener could navigate this abstract space in a way that gave them enough of an idea to know the direction of travel, but not so much that their action was implicit.

The development led me to externalise the eight qualities extracted from the object and visualise them in a way that the listener could actively navigate. This presented the listener with a fixed framework for exploration. A field with nothing on it feels arbitrary, making movement meaningless. What was needed was something ambiguous but still responsive. I built it in four stages.

The circle is divided. The scan returns

of the object, with each reading taking up an equal portion of the field. Direction selects a reading, while the centre represents a position that none of them leads towards.

Each of the eight slices of the field produces a prompt based on the AI's reading. This prompt, in turn, generates a related picture. The combined effect is an abstraction of the object's qualities, visualised. A primary inspiration for this development was the board game Dixit, in which players have to make meaning from abstract visuals.

The boundaries are dissolved. The final layer blurred the edges between the slices until no line indicated where one reading ended and the next began. The listener still has enough to draw from for movement to feel meaningful, and the conversation between the music output and the visual input allows the listener to explore and learn over time. The system logic creates a continuous conversation between the listener and the AI.

The listener's map is a bounded space, meaning the listener can explore all that the space has to offer. It is not infinite, encouraging new objects to be used and new experiences to be had.

//Closing insights

The design target became legible but unreadable: enough form to invite movement, but not enough to reveal where that movement leads.

Withholding a destination is only bearable if movement produces an immediate response. Otherwise, the ambiguity feels like system failure rather than space to be explored.

Stage 3



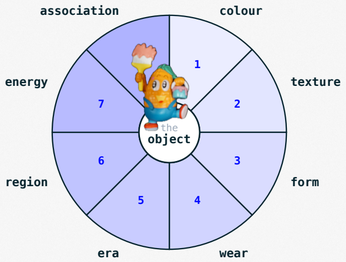
An ambiguity layer is placed on top if the initial output

Stage 2



Each section is rendered based off its own sections prompt.

Stage 1



The base is split into 8 sections based of the 8 qualities extracted from the object by the AI.

4.10 Tensions form the Explorations

The initial explorations were first guided by an overly simplistic question; how might qualities lost from the record-store experience be reintroduced into streaming? However, the aim was not to recreate the past. I was designing for a nourishing listening experience, which meant recognising and retaining the qualities that already make streaming valuable.

Each exploration brought several of these qualities into a streaming context at once, testing how they might be recombined within a digital service. This meant treating streaming as more than a medium from which something had been lost. It offers immediate access, scale, flexibility, and assisted discovery. These are genuine benefits, not problems to be designed away.

Reintroducing qualities from the past therefore changes the balance of the experience. It may create greater attention and engagement, while also narrowing access, slowing discovery, or reducing control. These trade-offs form the central design tension.

The project seeks a nourishing listening experience that supports attention and engagement without giving up the reach and flexibility of streaming. The past is therefore a source of qualities to work with, not a template to reproduce, and inconvenience is not treated as valuable in itself.

//Where next

The explorations did not resolve these tensions. Provotype 3 and 4 aims to materialise these tensions in the form of Design Provocations .

Provotype 3 begins from the physical side. The object and the search for it remain, but they lead into a catalogue with no fixed edge. Provotype 4 begins with streaming's condition of unlimited availability and introduces slowness, sequence, and looking back.

//Tensions uncovered

Although each exploration worked with multiple losses, together they revealed a set of competing pairs. Setting the two sides against each other revealed five tensions.

1. Weightlessness against presence

This tension sits beneath the others because it concerns the properties of the medium itself. Streaming can be shared, stored, and accessed without taking up physical space. Those same qualities make it difficult for the music to feel located in a particular place.

A record takes up space, which is inconvenient, but this is also what makes it available to be handled, displayed, and encountered.

Streaming: weightlessness	*Physical music: presence*
Easily shared	Tactile
Takes up no space	Can be displayed
Quick to reach	Gives the music a place in space

2. Everything against enough

Streaming provides access to a catalogue larger than any shop and larger than any listener can comprehend. This expands what can be heard and makes a broader and more diverse listening experience possible.

What it removes is the the focus, a shelf is small enough to be seen and understood. Its finitude makes selection feel like a decision as all potential options are visible and navigable. The value of the record store was not simply its smallness, but the way its limits made choice manageable.

Streaming: abundance	*Physical music: finitude*
More is possible	Manageable to navigate
Broad and diverse access	Possibilities can be seen
Nothing has to be ruled out in advance	Selection can feel deliberate and owned

3. Speed against dwelling

Streaming lowers the cost of exploration. Nothing has to be committed to before it is heard, and moving on costs almost nothing.

The same low cost removes the pause between one record and the next. Exploration 5 showed that the effort of walking across the room to retrieve an object could create a more deliberate relationship with listening. Slowness is not valuable by itself; it matters when something is made of the time it creates.

Streaming: speed	*Physical music: dwelling*
Exploration costs almost nothing	Slowness leaves room for reflection
Nothing has to be committed to in advance	Each move involves a small commitment
Rapid movement between possibilities	Decisions become noticeable

4. Steering against surrender

Both streaming and physical records give listening a shape. The difference is who controls that shape and when.

Streaming keeps control in the listener's hands throughout the experience. A record asks for a decision at the beginning and then takes over. Its sequence and duration have already been defined by the artist. Control exists on both sides, but one is exercised continuously while the other is given up after the initial decision.

Streaming: steering	*Physical music: surrender*
Personalised sequencing	A fixed listening source
Rapid variation	A clear beginning and end
Control remains with the listener	The listener can follow the artist's journey

5. Being shown against going looking

Both recommendation and self-directed searching can lead to music the listener did not previously know. The difference is who performs the work of getting there.

Streaming makes discovery happen without the listener having to set out to find anything. This is why more discovery can happen more often. By contrast, searching in a record store is slower and less reliable, but the search itself can create engagement. Exploration 1 suggested that the investment may lie not only in what is found, but in the act of finding it.

Streaming: assisted discovery	*Physical music: self-directed discovery*
New music can be encountered without seeking it	The listener has agency over the decision
Discovery is frequent and low effort	Choosing creates investment
The system performs much of the search	The search becomes part of the experience

4.11 Reflection and Synthesis

Left page:

Figure 4.31

Diagram showing the transition from Explorations to Provotypes.

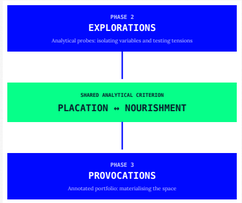
Right page:

Figure 4.32

Diagram showing the lineage from the six design explorations to the four synthesised provotypes.

The six design explorations detailed in this chapter did not function as linear steps toward a single product, nor were they disconnected prototyping exercises. Rather, within a Research through Design (RtD) framework, they operated as analytical filters.

When read chronologically, the progression of these explorations reveals a clear conceptual shift in this research: the transition from optimising system legibility to designing for productive ambiguity.



//The Emergence of the Placation–Nourishment Axis

The shared analytical criterion that governs this development, placation versus nourishment, was not defined ahead of the making; it was generated through it.

Exploration 1 initially set out to evaluate if personal objects could evoke personally resonant music associations. However, testing revealed that the experiential value did not reside in the recommendation's accuracy, but rather in the slow, effortful physical gesture of choosing and placing the object. Exploration 2 and Exploration 3 exposed a critical tension regarding transparency. Laying the machine's reasoning out step-by-step did not empower the listener; instead, it behaved as an tool for persuasion. A tidy, logical machine justification settled the listener so effectively that it narrowed their role to a passive acceptance or rejection of the algorithmic verdict.

This realisation reorganised the approach to the project ambiguity moved to a crucial design resource to preserve within the interaction. Hiding a deterministic recommendation pipeline is manufactured opacity, but keeping the machine's reading "felt but not legible" creates a generative gap. For this ambiguity to be navigable rather than arbitrary, it requires a structured, physical interaction space to anchor it.

//The Lineage of the Provocations

The four provocations serve two related but distinct purposes. The first two materialise earlier stages of the research process. Rather than presenting resolved outcomes, they make visible the questions, tensions, and uncertainties that emerged through the explorations. In this way, they invite viewers to wrestle with some of the same issues that shaped my own research and design process.

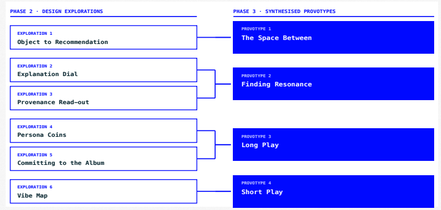
The second two materialise the learning that emerged from this process. They present conclusions from two different angles, demonstrating that the framework does not lead to one fixed interpretation or outcome. Instead, it can be realised in multiple ways. Making these interpretations tangible gives the viewer something concrete to engage with and gives me something material to reflect upon. This reflects the principles of Research Through Design, where knowledge is generated through making, doing, and reflecting on what has been made.

Provotype 1: The Space Between emerged from Exploration 1. It translates the raw, unguided capture–read–match–music pipeline into an intentionally literal interaction. The phone camera sees the object plainly and resolves it to a direct match.

Provotype 2: Finding Resonance emerged from Explorations 2 and 3. It combines the disclosure dial from Exploration 2 with the provenance tracing developed in Exploration 3. These ideas are materialised through a fork-shaped physical control, allowing the user to move between literal, step-by-step reasoning and more poetic, atmospheric forms of explanation.

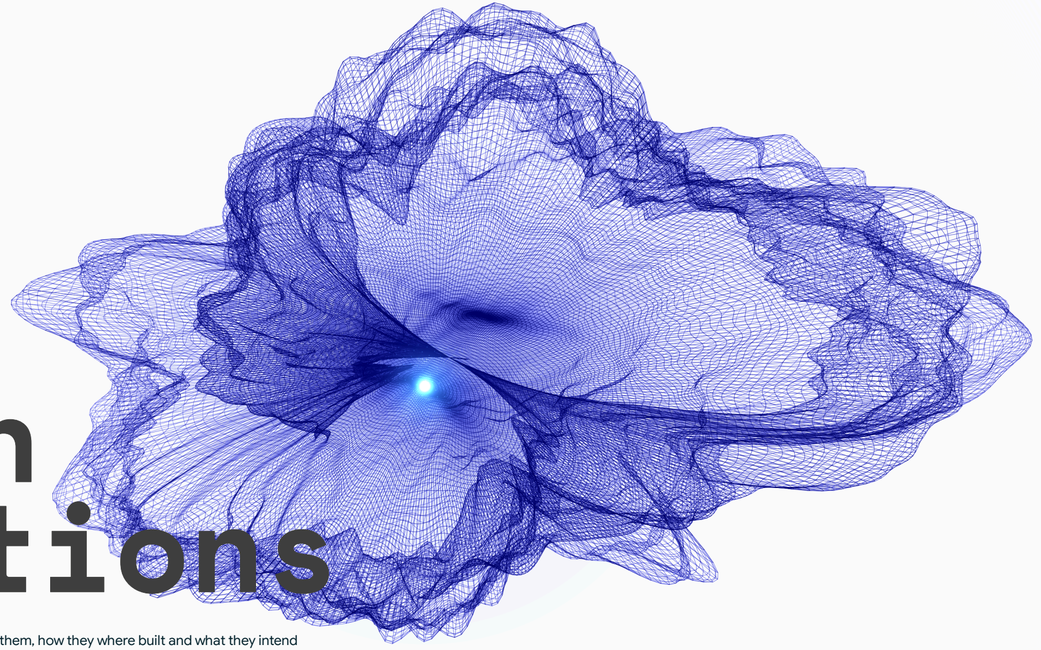
Provotype 3: Long Play synthesises the control of input data materialised as Persona Coins from Exploration 4 with the physical commitment and slow selection process mechanism of the stopper from Exploration 5.

Provotype 4: Short Play materialises the continuous, navigable "vibe space" developed in Exploration 6. Rather than directing the listener towards a single musical destination, it projects a responsive, label-free map of qualitative regions based on the object's features. The listener must continuously shape their own path of discovery by moving the physical object across the surface.



Chapter 5

Research Provocations



This chapter presents the materialisations of my research process, the inspirations behind them, how they where built and what they intend to provoke.

Provocation 1 - Introduction
Provocation 1 - Genesis
Provocation 1 - Concept Development
Provocation 1 - Anatomy
Provocation 1 - Interaction
Provocation 1 - Conclusion
Provocation 2 - Introduction
Provocation 2 - Genesis
Provocation 2 - Concept Development
Provocation 2 - Anatomy
Provocation 2 - Interaction
Provocation 2 - Conclusion

5.1.1 Provotype 1 // INTRODUCTION

Right page:

Figure 5.1

A Photograph of a user interacting with the final version of my provotype "The Space Between"

//THE SPACE BETWEEN

Look at your thumbs.
Small input devices for systems designed for output.
Think of how much of the world around you is interacted with solely through these input devices.
But what if the output is unknown? Furthermore, what if the output isn't the only goal?
This is often the experience when discovering music.
When discovery is reduced to a gesture on a screen, something gets flattened.
The distance between questioning and finding disappears.
And perhaps something is lost there.
This provocation begins with this not knowing.

We surround ourselves with objects; objects whose purpose is to say something externally about our internal selves. Music has always played a role in this. Our internal selves latch onto certain people who speak, in some way, to something within us, and through their songs these feelings become tangible.

So what if, instead of telling a search engine what we want, or letting an algorithm prescribe an answer to us, we gave it something from our world?

An object.
A memory.
A fragment of ourselves.
What if this abstraction was a space to sit in? What if meaning could be made in collaboration with the system, and what if the output became a byproduct of this meaning-making?

This provotype represents the first step in my exploration and asks what happens to the experience of searching for music when discovery leaves the screen and enters our lived environment, where exploration is enacted by one's whole self, through one's whole body.



5.1.2 Provotype 1 // GENESIS

//Key Inputs

Key Finding

1. From Exploration 1

"The primary experiential value lies in searching for an object rather than receiving a high-quality recommendation. Searching was enjoyable in itself: it brought neglected objects back into attention and involved moving them around the house to arrange them beside the device. This shifts the design focus from recommendation accuracy towards situated, spatial and domestic exploration."

2. From Exploration 1

"Objects provide an indirect but intentional route into music. They are encountered through material, colour, texture, wear, scale and context. Their ambiguity makes them useful for associative recommendation because they allow several meaningful musical directions rather than demanding one correct answer."

3. From Exploration 1

Capture and waiting have the potential to become part of the experience rather than mere latency. Handling the object, the countdown and the staged sentence make the transition to music perceptible.

4. From Exploration 2

"A clear and coherent explanation can settle the listener into accepting a recommendation while appearing to give them control, showing that transparency is not necessarily the same as agency."

Key Theory

1. AI magic, opacity and honest interpretation, §2.8 (Lupetti & Murray-Rust, 2024; Campolo & Crawford, 2020)

Chapter 2 frames AI as an interpreter rather than an oracle. Hidden mechanisms can give AI an unwarranted authority and place its decisions beyond scrutiny. The Space Between stages this failure: the listener provides a meaningful object but receives only a finished result, with the system's interpretation concealed. It restores physical input while leaving the AI's operation unaccountable.

2. Stimmung and productive ambiguity, §2.4 (Gumbrecht 2012; Gaver et al., 2003)

Chapter 2 frames Stimmung as something felt but not fully nameable, and ambiguity as valuable when it leaves space for personal interpretation. The Space Between closes this space by converting the listener's unformed feeling into a single recommendation with no opportunity to question or reinterpret the connection. It therefore demonstrates placation rather than nourishment, a distinction developed by this project.

Inspiration

1. Karmann's Paragraphica (2023). This AI-based, lens-free camera converts contextual data into a written description and then generates an image from that text. I adapted this process in Curio by using the blank callouts used in this project as the prompt structure, I found that in the project mentioned this provided a engaging way of showing the systems thinking to the listener (Karmann, 2023).

2. The record-store search. I adapted the spatial exploration of browsing a record store into Curio by using visual cues to guide users toward personal objects rather than fixed album choices.

Tensions

Weightlessness against presence: The owned object restores physical presence, gesture, and intention within a weightless digital system.

Speed against dwelling: Searching for the object and waiting for the result slow the encounter, creating space for interpretation.

Steering against surrender: Once the object is placed, control passes from the listener to the system.

Being shown against going looking: The listener actively searches for an object, but the system determines and presents the music.

5.1.3 Provotype 1 // CONCEPT DEVELOPMENT

//Project Mood board

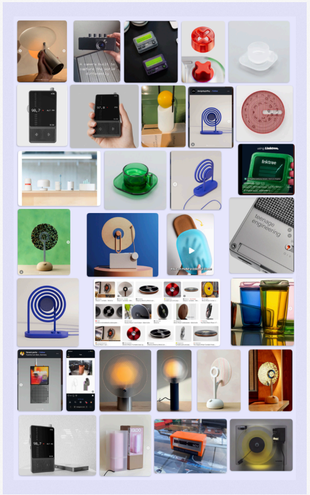


Figure 5.2

Collection of design inspiration used in ideation.

//Development



Figure 5.3

Some sketches of initial concepts, an AI render to act as guidance for 3D modelling, a combination of final pieces and test pieces from the materialisation phase.

//Final Concept



Figure 5.4

Image of the final materialised concept

5.1.4 Provotype 1 // ANATOMY

Symbolism / Metaphor

Components

1. The Vinyl Disc

This component has been designed to be the proportions of 12 inch vinyl record, to reflect where the inspiration for the search aspect of this provotype came from.

It has been made out of semi translucent PETG, this has been done so that the object that sits behind it is slightly obscured from the viewer reflecting the obscured reading the device takes from the object.

The centre label has been replaced by a viewing hole. Traditionally the centre label would denote a graphic and description of the record, as in this device the object serves this purpose it occupies this space

2. The Object Stand

The section the object sits on is only big enough to put personal objects on it reflecting insights from test 1 about what objects should be used in the interaction

3. The Phone Holder

This covers most of the users screen only leaving space for the output reading of the object to appear. Interaction is therefore focused on the object and the read out from the system.

Functional Parts

4. Users phone

5. Casing part 1

6. Casing part 2

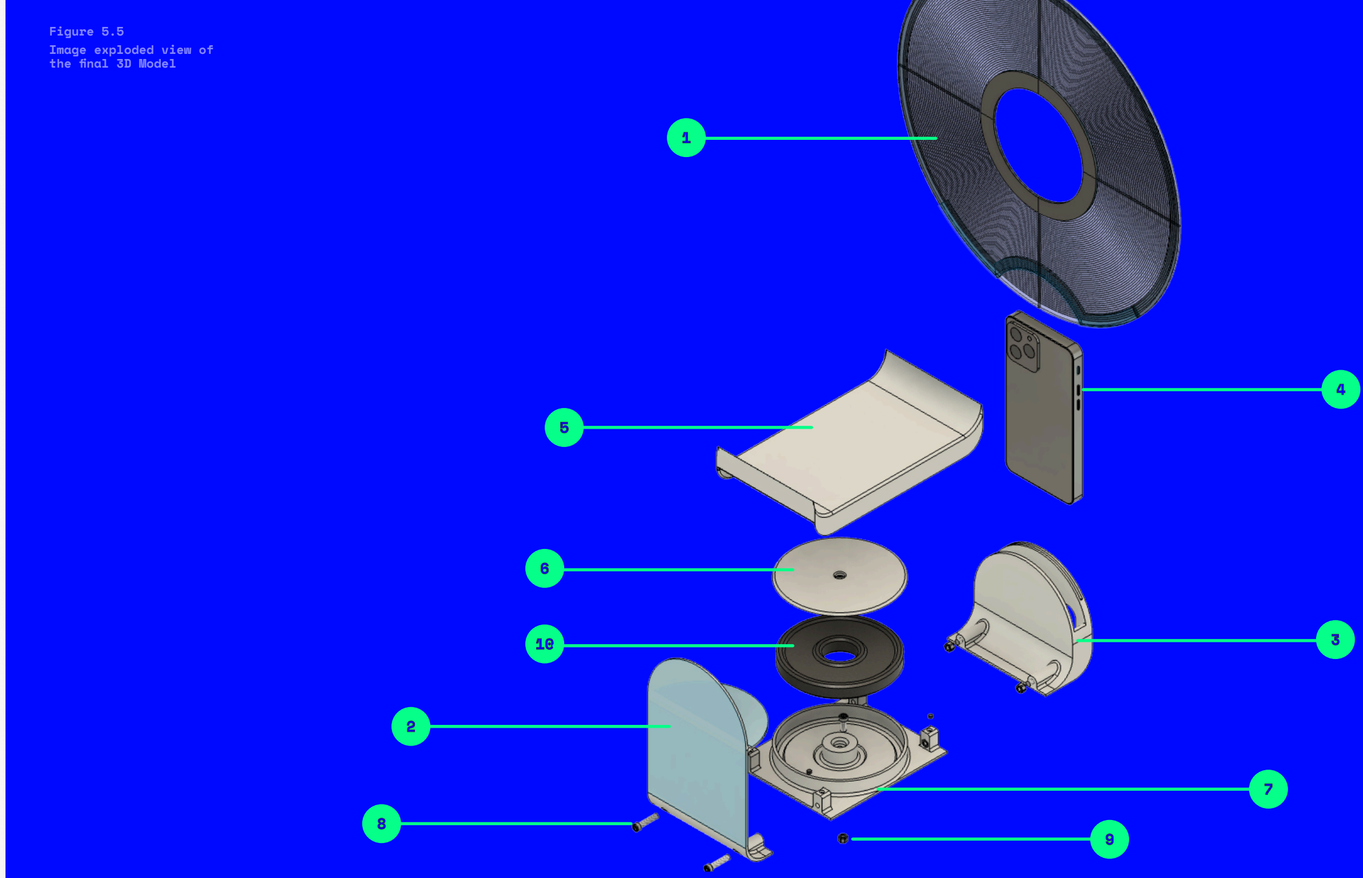
7. Casing part 3

8. M4 screws

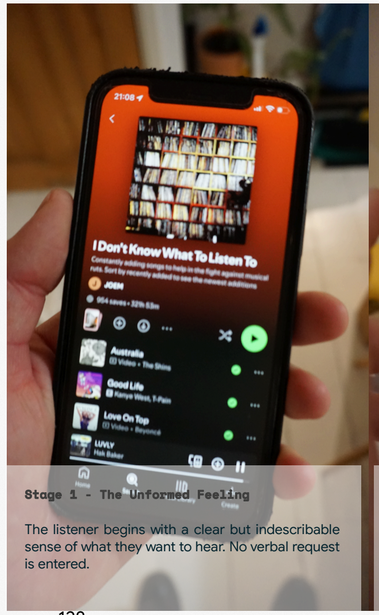
9. M4 nuts

10. 0.5kg Weight

Figure 5.5
Image exploded view of the final 3D Model

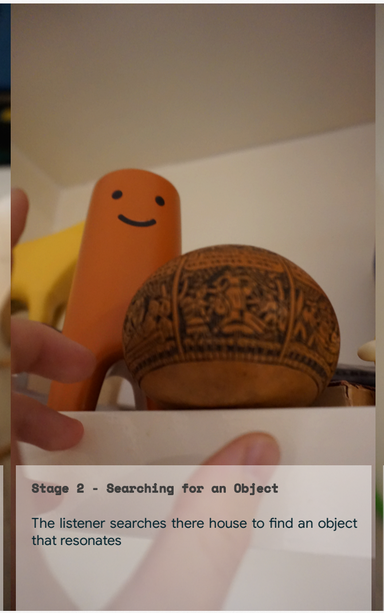


5.1.5 Provotype 1 // INTERACTION



Stage 1 - The Unformed Feeling

The listener begins with a clear but indescribable sense of what they want to hear. No verbal request is entered.



Stage 2 - Searching for an Object

The listener searches there house to find an object that resonates



Stage 3 - Selecting the Object

The listener chooses an object through association rather than rational explanation.



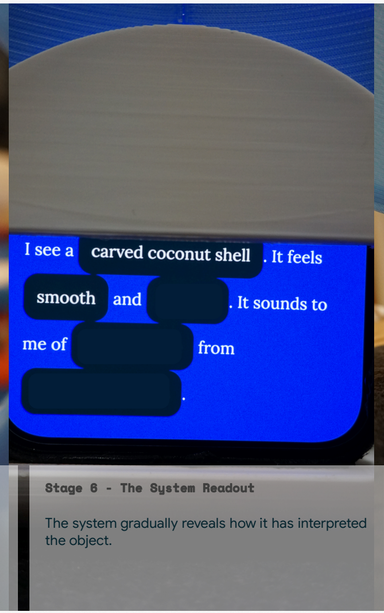
Stage 4 - Materialising the Request

The listener places the selected object in front of the camera. This gesture transforms an intangible feeling into a physical input.



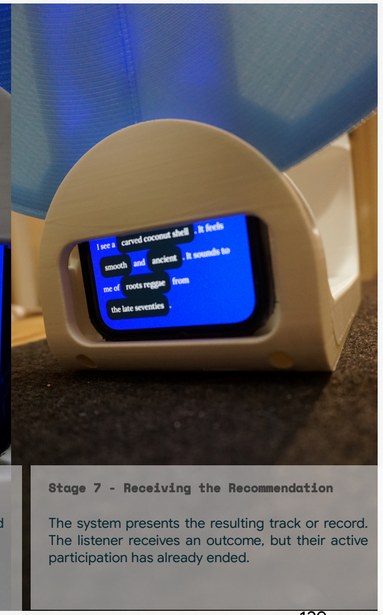
Stage 5 - The Artificial Wait

After the object is placed, the system begins processing. The listener cannot intervene or clarify their intention.



Stage 6 - The System Readout

The system gradually reveals how it has interpreted the object.



Stage 7 - Receiving the Recommendation

The system presents the resulting track or record. The listener receives an outcome, but their active participation has already ended.

5.1.6 Provotype 1 // CONCLUSION

Right page:

Figure 5.6

A render of the materialised concept.

Provotype 1 is a research artefact, not a proposed solution. It materialises the project's starting position, making its assumptions visible so they can be examined and challenged. Its argument comes from two places.

The literature review established that contemporary recommendation is organised around an explicit request input by the listener. Interpretation happens before the encounter, leaving the listener to accept or reject a completed result. Against this model, Stimmung named the condition the project was concerned with. A sense of what you want may be clear while still resisting description. Treating that sense as an incomplete query reduces it to the nearest available label.

Exploration 1 tested whether an ordinary object could work as an ambiguous stand-in for this request. Selecting an object to manifest this unknowable feeling, holding it, and placing it in front of an associative AI system initiated a process of discovery without requiring anything specific to be named. This highlighted that the value of the interaction is in what happens prior to interacting with the device. Searching the house for an object was more rewarding and important to the interaction than the accuracy of the track that followed, shifting the focus of the project from recommendation accuracy to a nourishing exploration experience. The provotype introduces an artificial waiting period for the tracks to surface and a system readout that slowly reveals itself. This was not latency to be tolerated but an attentive part of the interaction that gave space for meaning to be made between the user's interpretation and the AI's.

The Space Between stages these findings as a controlled demonstration. Once the object has been placed, every subsequent decision belongs to the system. The listener cannot question or correct the process, highlighting a gap that later explorations would explore. This constraint is deliberate, and the AI's output is formulaic and consistent. It provokes the questions: How should this unresolved meaning be resolved? Who should resolve it; the AI, the listener, or the object? Where should this agency lie?

The central finding is that placing something you own in front of a camera restores the gesture, intention, and materiality that streaming removed, but not what that gesture carried. Physical input is not the same as agency. The listener's participation ends when the object is placed; everything that follows is delivery.

The provocation therefore demonstrates the placating side of music exploration, which my project aims to transform into nourishment.



5.2.1 Provotype 2 // INTRODUCTION

Right page:

Figure 5.7

A Photograph of a user interacting with the final version of my provotype "Finding resonance"

//FINDING RESONANCE

We are in control of our own actions and choices.

Take out your phone and check your screen time from the past week. Are we in control of our own actions and choices?

We are beholden to the technologies we use and the algorithms behind them. The systems that surround us influence us without us even knowing it. If we accept that we are being influenced, do we want to understand how that influence takes place? And does understanding its mechanisms give us agency, or simply make the influence more visible?

Finding Resonance begins with this contradiction.

Place a personal object on the platter. As the system reads it, slide the tuning fork to find resonance with its interpretation.

This provocation has two modes. In one, the listener changes the system's reading of the object, moving between the literal and the poetic. This reading then influences the song produced by the system.

In the other, the song remains fixed. The listener instead controls how the system presents its reasoning. In this mode, the reading influences the user's perception.

The tuning fork therefore offers two kinds of control: influencing what the system produces, or influencing how the system persuades us to accept it. This provotype asks whether the unknown is meant to be presented or discovered, and what the costs of either path are.



5.2.2 Provotype 2 // GENESIS

//Key Inputs

Key Finding

1. From Exploration 2

"A clear and coherent explanation can settle the listener into accepting a recommendation while appearing to give them control, showing that transparency is not necessarily the same as agency."

2. From Exploration 2

"Literal and poetic explanations produce different forms of agency because a literal reasoning chain exposes steps that can be checked or rejected, whereas a poetic account closes those gaps and makes the listener judge the account as a whole."

3. From Exploration 2

"The user became bored whenever the answer was simply handed over, in either version. The interaction only held them when there was something left for them to do."

Theory

1. Recommender explanations and judgment, §2.8 (Bilgic & Mooney, 2005)
Chapter 2 distinguishes explanations that promote acceptance from those that support critical judgment and user control. Provotype 2 applies this distinction through the fork-shaped slider: the listener can move towards poetic explanation, which may settle them into accepting a recommendation, or literal explanation, which exposes the reasoning and allows them to steer the musical result.

2. Stimmung and productive ambiguity, §2.4 (Gumbrecht & 2012; Gaver et al., 2003)
Chapter 2 treats Stimmung as felt but not fully nameable, and ambiguity as valuable when it keeps interpretation open. Provotype 2 applies this as bounded ambiguity, offering different explanatory registers without presenting one as the correct reading. The listener is given enough structure to question and reinterpret the AI's response.

3. Attunement and resonance, §2.4 (Gumbrecht & 2012; Heidegger, 1927/1962)
Chapter 2 frames Stimmung as a relational atmosphere that is sensed rather than specified as a query. Provotype 2 materialises this process through the tuning fork: moving between explanatory registers lets the listener search for the level at which the AI's interpretation resonates with their unformed feeling.

Inspiration

1. Odom and Duel's OLO Radio (2018). This physical radio turns listening-history data into a slow, open-ended interaction through a motorised slider. I adapted this approach in Finding Resonance by giving the AI's interpretation a physical form that could be adjusted by hand (Odom & Duel, 2018).

2. The record store. In the record store, resonance is found through searching, browsing and following associations rather than through a linear journey. I carried this into Finding Resonance by asking the listener to search for the point at which the AI's interpretation becomes meaningful.

3. The tuning fork as metaphor. A tuning fork establishes resonance in music, while also making vibration perceptible through the body. Finding Resonance brings these meanings together by allowing the listener to tune both the musical outcome through use of there body.

Tensions

Weightlessness against presence: The physical object and platter give the AI's interpretation a tangible form.

Speed against dwelling: The slow rotation and adjustable readout encourage the listener to stay with the interpretation rather than immediately accept the recommendation.

Steering against surrender: The slider can either let the system persuade the listener or allow the listener to influence the recommendation.

Being shown against going looking: The system presents an interpretation, but the listener must search through its levels to find a relationship that resonates.

5.2.3 Provotype 2 // CONCEPT DEVELOPMENT

//Project Mood board

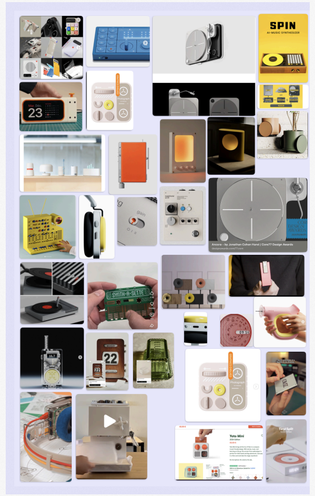


Figure 5.8

Collection of design inspiration used in ideation.

//Development

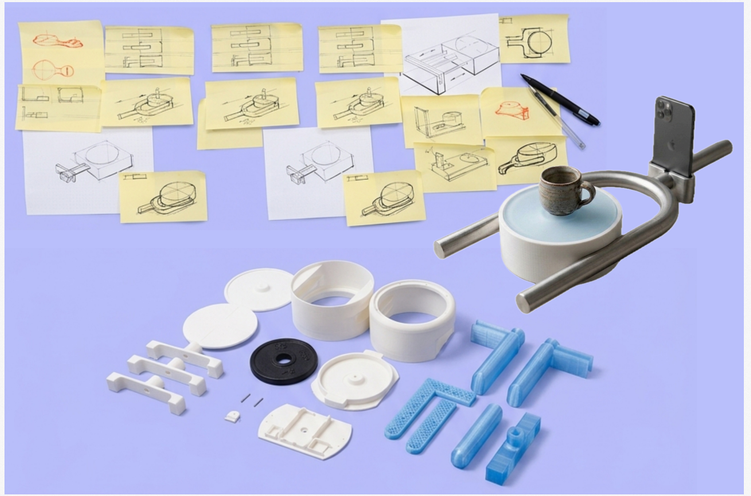


Figure 5.9

Some sketches of initial concepts, an AI render to act as guidance for 3D modelling, a combination of final pieces and test pieces from the materialisation phase.

//Final Concept



Figure 5.10

Image of the final materialised concept

5.2.4 Provotype 2 // ANATOMY

Symbolism / Metaphor

Components

1. The Tuning Fork

This component is the core metaphor of the provotype. A tuning fork establishes a reference pitch through vibration and can also make vibration perceptible through the body. Positioned between the object and the listener, it connects the source of the AI's interpretation with the person judging whether it resonates. As the user slides the fork it changes the system's explanatory register, the user manipulates the system till they find "resonance' with it

2. 10kΩ Slide Potentiometer

This component requires precise and deliberate movement. Its action recalls the careful handling of a vinyl player and the faders used on mixing desks to shape sound.

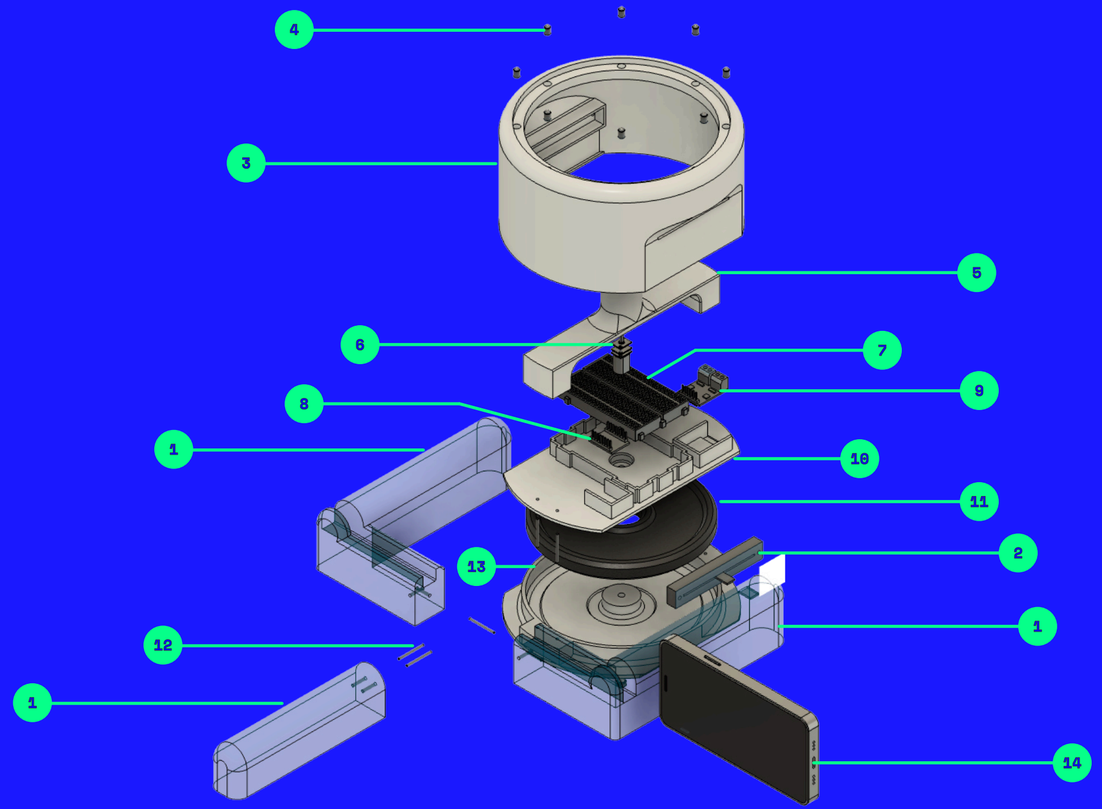
3. Main Body

The main body has a rotating top reminiscent of a vinyl player. The object is turned slowly and read from multiple sides, this creates a ritual waiting, making the AI's processing visible through the mechanical movement.

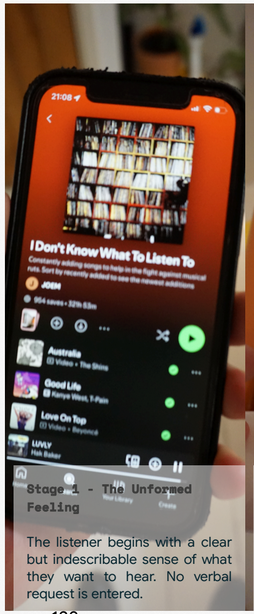
Functional Parts

4. Bearing
5. Motor holder
6. Motor
7. Breadboard
8. ESP32-C3
9. L9110S driver
10. Component holder
11. 1kg weight
12. Pins
13. Base
14. Phone

Figure 5.11
Image exploded view of the final 3D Model



5.2.5 Provotype 2 // INTERACTION



Stage 1 - The Unformed Feeling

The listener begins with a clear but indescribable sense of what they want to hear. No verbal request is entered.



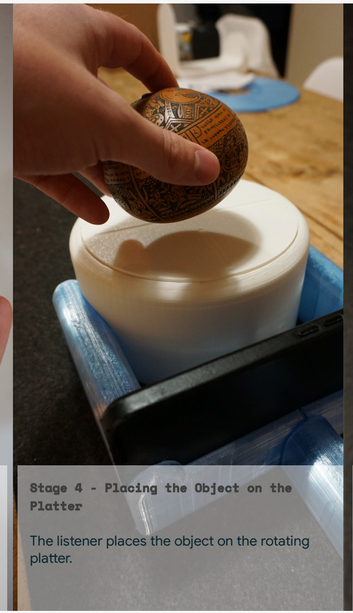
Stage 2 - Searching for an Object

The listener searches there house to find an object that resonates



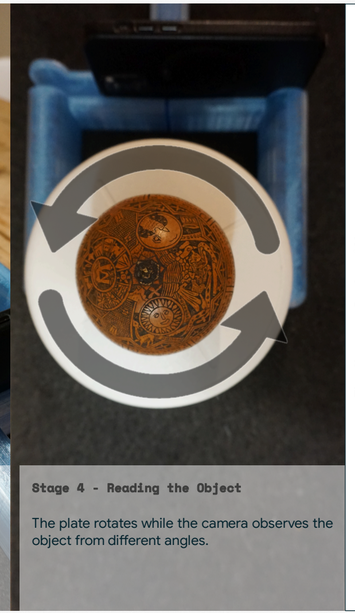
Stage 3 - Selecting the Object

The listener chooses an object through association rather than rational explanation.



Stage 4 - Placing the Object on the Platter

The listener places the object on the rotating platter.



Stage 4 - Reading the Object

The plate rotates while the camera observes the object from different angles.

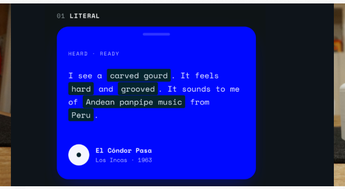
The Users selects either

Path B - Recommendation changes



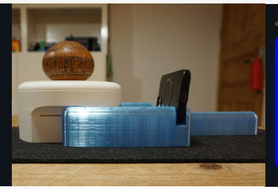
Stage 5 - First Position

The slider is moved and the system produces a different recommendation.



Stage 6 - Changing Register

The listener moves through different positions, each producing a different musical result. The slider becomes an active means of shaping the outcome.





Stage 7 - Finding Resonance

The listener searches for a result that resonates with their original feeling.

Final Condition

The listener helps shape the recommendation through interaction. This path demonstrates nourishment through participation and shared meaning-making.

Path A - Explanation changes



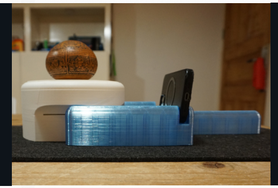
Stage 5 - First Position

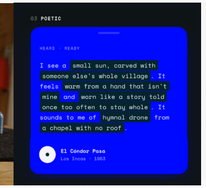
The slider is moved to reveal the system's detailed reasoning. The recommendation remains fixed.



Stage 6 - Changing Register

The listener moves towards an atmospheric description of the same result. The music does not change; only the language surrounding it shifts.





Stage 7 - Finding Resonance

The listener searches for an explanation that makes the fixed recommendation feel convincing.

Final Condition

The listener accepts the recommendation because the explanation has settled or persuaded them. This path demonstrates placation through convincing interpretation.

5.2.6 Provotype 2 // CONCLUSION

Right page:

Figure 5.12

A render of the materialised concept.

Provotype 2 began with a question about system transparency. Provotype 1 had concealed almost everything: the listener received a single sentence but could not see how the object had been interpreted or why that interpretation had produced a particular record. The apparent correction was to reveal more. Exploration 2 made the system's reasoning visible through an adjustable readout while keeping the music fixed, so that only the explanation changed. Exploration 3 extended this approach by tracing each recommendation back to the human sources that had supported it.

Both explorations pointed towards the same insight: disclosure did not necessarily return control to the listener. A clear, well-ordered chain of reasoning made a recommendation more convincing, generating acceptance of recommendations the listener might otherwise have rejected. Attribution behaved similarly. Naming a source made a weak match easier to accept without making the match any stronger and, most importantly, without bringing it any closer to the indescribable feeling the user initially held. Explanation was not neutral; it performed persuasive work. As it became more complete, attention shifted away from the initial, unknowable feeling the listener brought with them and instead towards the system producing it.

This changed how the project understood Stimmung. Until this point, Stimmung had described the listener's unformed sense of what they might want, while the design task was understood as finding a way to answer that sense. If the object could be interpreted accurately enough, the uncertainty would be resolved. The explorations showed that no level of explanation output by an AI system will necessarily be accurate. It may prove satisfying to the listener, but the feeling that follows is one of placation rather than nourishment. Not-knowing therefore ceased to be a problem for the system to close. It became a condition the interaction needed to hold open. Ambiguity shifted from a defect to a resource, provided it remained sufficiently bounded for the listener to recognise and question a relationship rather than simply receive one.

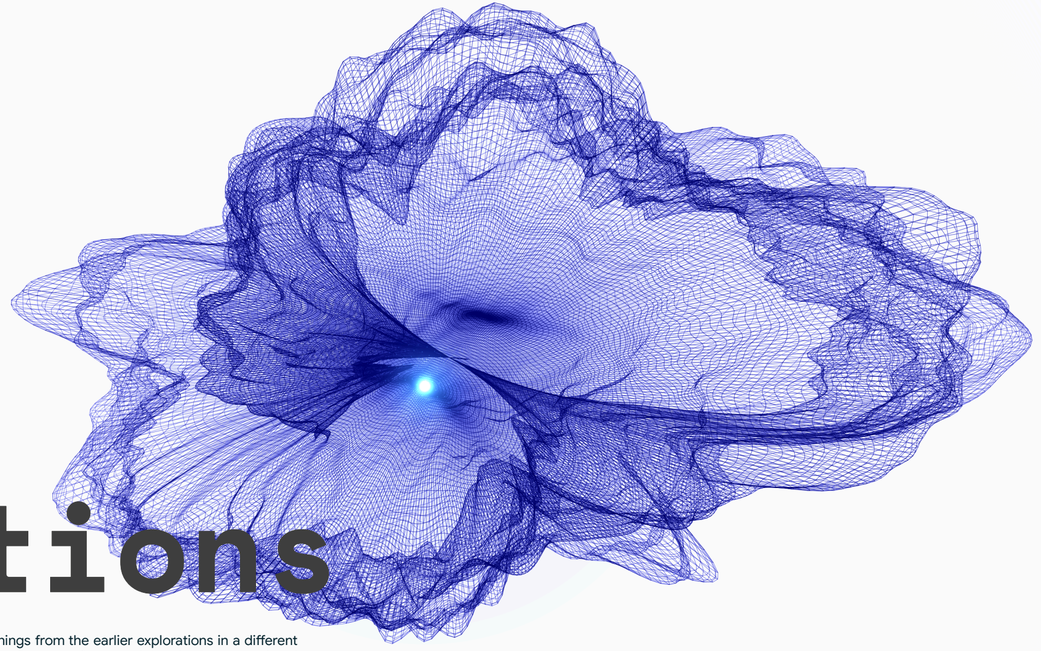
Finding Resonance makes this design knowledge tangible. An object rotates slowly on a platter while a camera reads it from multiple angles. A fork-shaped slider sits between the listener and the object. The system can function in one of two ways. Firstly, moving the slider through its positions retells the same result at different levels of interpretation, from the system's step-by-step reasoning to a description of atmosphere. The tuning fork functions as both metaphor and mechanism. The listener is not given a single correct register; they search for the level at which the AI's account resonates, highlighting that their resonance with the outcome can be influenced by the information presented to them.

Secondly, the listener interacts in the same way, but the system's output affects the musical recommendation. This gives the listener greater agency over the output by allowing them to find resonance with it. Explanation is therefore not simply a window onto the machine. It can either steer the system or settle the listener. The fork places both possibilities in the listener's hand.

This is the core of the provocation and demonstrates the project's central axis: the system can either influence the user to find resonance with a result, asking little of them in return, or it can ask the listener to participate in the process of meaning-making and therefore require more of them. This forms the two sides of the nourishment versus placation argument. Future provotypes explore how this interaction between object, listener, and AI can be applied differently when the output medium changes, such as from an album to a song, because each medium requires different things from the listener during the interaction.



Chapter 6



Design Provocations

This chapter presents two culminations of the research conducted, each utilises the learnings from the earlier explorations in a different way to produce a distinct outcome. Together, Long Play and Short Play demonstrate how the same framework can be interpreted and realised through different forms.

Ideation & Concept Generation
From Framework to Hardware
System Architecture
Technical Build Up
Object Collection
Final User Test

6.1 Two Economies of Commitment

The two final provotypes should not be understood as alternative interactions for two media formats, but as responses to different levels of listener commitment. An album and a song are not simply different kinds of content: they ask different things of the listener. An album demands a substantial investment of time and attention, whereas a song requires far less. This difference in commitment should shape how discovery is structured.

// Provotype 3: Long Play

Long Play treats the album as a high-commitment listening unit.

Because an album asks for time, the provotype concentrates commitment into a single decisive gesture and withholds everything before it. The listener sets up the conditions of discovery once, then leaves them alone. The screen proposes albums slowly and in silence, and the system offers no answer until the listener takes one.

Setup is not a declaration. The listener taps their phone against a reader, handing over their listening history, and the system reads that history into a set of listening personas. It shows them, one at a time, and the listener finds an object of their own that fits each one. Each object is scanned and bound to its persona. The listener controls what is handed across, never the reading taken from it, and answering a persona with something already owned is a choice from a shelf rather than a stated want.

The bound objects are then arranged on a circular base, and their positions carry the weighting. An object near the centre counts strongly, one at the rim faintly, and the recommendation comes from the relation between them rather than from any single one. Weight here means statistical importance and nothing else; no mass is measured and nothing is weighed. Position holds a persona and a proportion at the same time, so a taste can be given as a subtle mixture.

While the objects sit still the system holds one fixed reading. It shows an album, then slowly the next, silently, with no label and no account of what changed. Albums it holds in higher confidence in stay on screen longer and come back around, so duration carries the uncertainty that a named reason would otherwise resolve. Interpretation happens once per arrangement rather than continuously, and moving an object is what causes the system to read again.

The gesture that commits is not the lifting of a record. The listener lifts a cork ball off the speaker, and the album on screen plays whole and in order. It is the only moment in the interaction that resolves anything. When the album ends the music stops and the drift resumes. Over time, as the listener's listening diversifies, a new persona area opens and a new object is found and scanned for it, so the set is not fixed at setup.

Visually, the provotype refuses to generate an abstract interface for the album, and it prints no genre or mood read-out either. The album sleeve is already an authored surface that speaks about the music before the listener can describe what they are looking for, and its artwork, typography and material presence perform the pre-declarative work of making an album judge-able. Because that bounding is inherited rather than generated, the provotype can spend its design attention on arrangement, proportion and commitment instead of replacing the sleeve with another representation. What it concedes in the same breath is that the coupling between position and persona is arbitrary and swappable.

In this provotype, the system's agency forms a quiet plateau. It interprets, withholds, and allows the listener's attention to build. At the lift, the only stage with no inference in it at all, that agency collapses into playback.

//Provotype 4: Short Play

Short Play begins from the opposite condition: the song is a low-commitment listening unit.

A single track does not justify the same physical or cognitive effort as an album. Asking the listener to perform a slow, decisive selection for every song would put the cost of the interaction above the value of the result. Short Play therefore places no commitment at the point of arrival. The value sits in the movement through the possibilities.

The object is set into a bay, where it turns and is read once, in full, from every side. It is never read again. The system then builds a place from that reading: 8 soft regions drawn from material, colour, finish and associative readings combine and blending into one another with no boundary, no text.. Regions are stable and rankings are not, which is why one reading becomes a field rather than a list.

After the render the system stops deciding. The listener moves the object and lets it sets and a song begins to play. There is no confirmation step, the lister is guided to explore based of the visual cues seen before them and the music that those cues represent.

Moving the object moves the music not the image the image stays fixed creating a defined field to explore and retrace per object. The pool is frozen at the scan, so the field keeps its coordinates . When the object stops, the music holds where it rests.

Here the object changes role from the thing the system reads; here it is read first and then becomes the thing the listener steers with enhancing the experience with the continued tactility of the object used. This is a role the object gains rather than one it loses, since the personal object as input stays constant across the set.

Short Play therefore supplies an associative visual reading that the single songs don't tend to possess, there was nothing prior in a streaming context that would suggest the "vibe" of a song the only visual cues they where presented with where that of the album they where a part of, making it an unsuitable visual cue. The map supplies it, but keeps it open and navigable rather than turning it into a fixed claim about what the song is.

	Long Play	*Short Play*
Listening unit	Album	Son
Cost	High temporal commitment	Low temporal commitment
Main gesture	Discrete: Arranging to Intentional: Lifting	Continuous: sliding and wandering
Visual strategy	Uses the album sleeve	Builds a coordinate map
AI agency	Sustained before selection, then resolves	Strong at the start, then recedes
Listener agency	Builds gradually, then commits	Remains active throughout
Where the nourishment sits	In the arrival of the output	In the process of exploration

6.2 Provotype 3 - Long play



Onboarding

Interaction. The listener taps their phone against a reader, handing over their listening history. The system reads it into a set of personas, presents them one at a time, and the listener scans an object of their own for each.

Affordance. Answering a persona with something already owned makes the choice one from a shelf rather than a stated want. The listener controls what is handed over, but not how it is read, so setup supplies information without becoming a declaration.

1



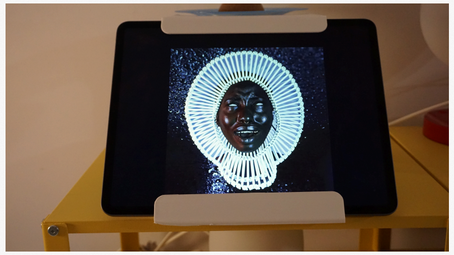
The placing

Interaction. The listener arranges the bound objects on the circular base, placing some near the centre and others at the rim. Nothing plays.

Affordance. Distance from the centre sets the weighting: an object near the middle counts strongly, while one at the rim counts faintly. Position holds both a persona and its proportion, allowing taste to be given as a mixture rather than a single dominant fact.

Note. Weight here means statistical importance and nothing else. No mass is measured and nothing is weighed.

2



The drift

Interaction. The listener leaves the arrangement alone. The screen shows one album, then slowly another, in silence, without a label or explanation of what changed.

Affordance. A named reason invites acceptance, so the screen shows the consequence of the arrangement rather than an explanation of it. The reading remains fixed while the objects sit still, so a long drift is not the system watching.

3



The dwell

Interaction. The listener still does not touch the objects. Albums the system holds with higher confidence stay on screen longer, and the albums it is most confident about return more often.

Affordance. Duration carries the uncertainty that a named reason would otherwise resolve. An even rotation would leave no uncertainty to feel, so confidence takes a form the listener can sense without being given a number.

4

Long play



The lift

Interaction. The listener lifts a cork ball from the speaker, and the album on screen plays in full and in order. When it ends, the music stops and the drift resumes.

Affordance. One gesture commits, and it is the only moment in the interaction that resolves anything. Concentrating commitment into a single decisive act matches the album's demand on time and attention.

5



The rearrangement

Interaction. The listener slides one object closer to the centre or swaps the positions of two objects. The system reads the arrangement once, and different albums surface when the drift resumes.

Affordance. Interpretation happens once per arrangement rather than continuously, so moving an object is what prompts the system to read again. Nothing is added or removed; the arrangement itself is the input.

6



The unlock

Interaction. As the listener's listening diversifies, a new persona area opens. They choose an object of their own, scan it, and add it to the set.

Affordance. A set fixed at setup would freeze a taste that can change. The system notices the new persona, but the listener still supplies the object, so the set grows in the same way it was made.

7



6.3 Provotype 4 - Short play

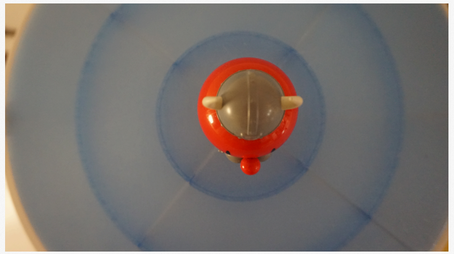


Onboarding

Interaction. The listener taps their phone against a reader, handing over their listening history.

Affordance. The recommendation system now has a basis to work from, in-order to bring unknown music it must first develop an idea what is know to the user.

1



The read

Interaction. The listener places the object in the middle of the screen, where it turns and is read once, in full, from every side. It is never read again.

Affordance. The regions remain stable, but the music within them is not reduced to a fixed ranking. One reading therefore becomes a field rather than a list. Reading once and never again keeps the terrain still under the hand that is about to move through it.

2



The render

Interaction. The system builds a place from that reading: eight soft regions drawn from the object's material, colour, finish and associated qualities blend into one another without boundaries or text.

Affordance. A song has no visual of its own, so the visual field has to be constructed. Softness and the absence of labels keep the reading navigable rather than turning it into a fixed claim about what the object is.

3



Entering

Interaction. The listener moves the object from this central position and music begins to play.

Affordance. Position is the query and the object already has one, so nothing needs to be declared before the first track.

4

Short play



Travelling

Interaction. The listener moves the object across the field, and the music moves with it. The image stays fixed. As the listener moves so does the music playing (with a 1 second delay)

Affordance. A still image creates a defined field to explore and retrace, one field per object. The listener follows the visual cues in front of them and discovers what those cues produce through the music.

5



Doubling back

Interaction. The listener returns the object to a place it has already visited, and the same music is there.

Affordance. The pool is frozen at the scan, so the field keeps its coordinates. If the map reshuffled under the hand, it would create an endless search space and therefore preventing a sense of completion and experimentation by using new objects.

6



Resting

Interaction. The listener stops moving the object, and the music holds at that position continually playing songs that are triggered by this recommendation. The user can always move to explore more.

Affordance. Allows for active searching then the reward of a personalised stream of nourishing recommendations.

7



Chapter 7

Discussion

This chapter makes three arguments drawn from learnings from the RtD process and suggests a transferable that can be applied by other design researchers.



The Case for Using the Past as Design Material
The Case for Personal Objects as Inputs
The Case for Effective Ambiguity
The Transferable Method

7.1 The Case for Using the Past as Design Material

Not Designing Nostalgia

A primary challenge of this project was to look back at physical, record-based music discovery without falling into the trap of uncritical, "rose-tinted" nostalgia. Looking back at a superseded technology can easily be mistaken for a sentimental preference. However, in my process, the justification for studying past practices is methodological rather than sentimental. Experiential qualities that operated silently in an older medium often become read as nostalgia once it has been withdrawn by a newer one. While physical record shops were the dominant means of finding music, nobody had to explicitly name what tactile engagement, spatial movement, or album commitment were doing for the listener, because there was no alternative against which to compare them. Therefore, the proper object of study in this thesis is not record-based practice itself, but the qualitative dimensions of experience that became absent during the transition to frictionless digital streaming. The purpose of looking back is not to restore obsolete physical artefacts, but to recover lost affordances and translate them into a contemporary context.

Recovering Affordances through empirical observation

To address a gap in the literature, where analyses of record stores often rely on interviews focused primarily on their surrounding social and cultural factors, this project approached the physical record store through systematic empirical journey mapping. These observations were then cross-read against established sociological and cultural studies of music consumption to identify robust, translatable experiential qualities, connecting to what people seek from music exploration, rather than to what they want from a record-store experience itself. By combining these experiential losses with design implications formed from the research, the resulting nine dimensions of loss formed a structured, cross-referenced framework rather than an expression of individual memory or nostalgia. The fieldwork therefore functions as a sensitising framework: it identifies experiential qualities that have been withdrawn from contemporary music exploration and directs attention towards how they might be actively designed into new contexts.

Designing Nourishment Inside Abundance

This approach imposes a central design criterion: every loss associated with physical media must be considered alongside the gains enabled by digital streaming. The aim is not to restore an earlier mode of listening wholesale, but to work within the tensions between what physical media made possible and what streaming makes possible today. These tensions form the basis of the design exploration; they must be made tangible in ways that produce a nourishing experience, rather than simply resolved in favour of either the past or the present.

Taken together, these claims position the design task not as recreating the record store, but as recovering the experiential functions it supported. A design proposal that counters streaming's advantages by reintroducing artificial scarcity or arbitrary friction would ask users to move backwards. Instead, this methodology is guided by the principle of reintroducing interpretive nourishment within abundance. The physical act of choosing an object functions as a pre-verbal "sign of the self" (Rochberg-Halton, 1984). Rather than requiring listeners to formulate and declare their intentions through a search box, the act of reaching allows an unnameable or not-yet-articulated state to enter the system through action. The effort is therefore valuable not because it is difficult, but because it creates space for interpretation, commitment, and self-expression within an otherwise frictionless environment.

The Limits of the Past as a Design Material

However, past practices can reveal which experiential qualities may be missing, but not how to translate them into a digital system. Decisions about interpretation, prompts, and software therefore remain matters of design judgement. Research and design exploration can expose tensions but cannot determine what a contemporary design should prioritise. Reintroducing interpretation, anticipation, or effort may support user agency, but can also conflict with the convenience and immediacy of streaming. The four outputs are thus provocations rather than finished solutions; they make these tensions tangible while treating the past as a starting point, not a design authority.

7.2 The Case for Personal Objects as Inputs

Requiring a listener to select a personal object and place it on a physical surface may initially appear to introduce unnecessary friction into an otherwise optimised digital domain. Yet in Curio, the object is not a decorative container for a digital query; it is the material prompt through which interpretation begins. Rather than asking listeners to declare what they want, Curio uses the act of selection to support embodied attunement. This section presents five arguments for using existing personal possessions as inputs for discovery in music streaming.

1. Material Pragmatism

From a practical design standpoint, using personal possessions draws on a rich, pre-existing material ecosystem within the domestic environment.

1. Avoiding Additional Material Consumption: Generating a tactile experience without requiring users to acquire another set of objects. By working with possessions already present in the home, it avoids adding the cost and waste associated with proprietary tokens. Materiality becomes part of the experience without turning the experience itself into a reason for increased consumption.

2. Authentic Material Variety: Ordinary household objects arrive with differences in shape, surface, colour, wear, glaze, scale, and material intact. These qualities provide a rich visual and material field for the multimodal AI to interpret.

3. Embedded in Domestic Life: Personal objects already inhabit the listener's environment, remaining present beyond moments of music listening. They can be encountered, used, and remembered across everyday activities, allowing their relationship to develop within the broader rhythms of domestic life rather than remaining confined to a single listening session.

2. Biographical Anchoring

The central challenge this project addresses in exploratory music discovery is: how to begin listening when one does not yet know what one wants. Personal objects help address this problem through their biographical histories.

1. Signs of the Self: Personal possessions can act as material anchors for identity and memory. They carry specific histories, emotional associations, and elements of the self that may be invisible to anyone outside the relationship. These same qualities shape how people relate to music, which is often experienced through personal memory, identity, and emotional attachment.

2. Pre-Verbal Choice as Interpretation: Reaching for this particular object rather than another is an informed, self-expressive gesture even when the listener cannot yet describe their state in language.

3. Externalising the Abstract: The listener's initial state is abstract, and not yet fully available in language. It is not a hidden preference waiting to be accurately classified. By selecting a personally meaningful object, the listener materialises that abstraction in a form the AI system can perceive.

4. A Starting Point Rather Than a Diagnosis: Curio does not need to identify exactly what the listener feels before the interaction begins. The object provides a concrete starting point from which the listener can explore an initially indeterminate state. The reach therefore functions as a query, but not because the object contains a readable answer. It transforms an abstract and pre-verbal condition into a material input that the system can respond to.

3. Private Meaning and Machine-Readable Form.

The object's personal significance remains primarily with the listener. The system may perceive the object's material form, but it cannot access the biography or meaning that makes the object important unless the listener explicitly provides that context. The interaction therefore does not eliminate ambiguity; it gives both listener and system a concrete point from which exploration can begin.

1. Meaning Is Held by the User: The object's personal significance does not determine the machine's output. That significance belongs to the listener and remains inaccessible to the system.

2. Nourishment Without Algorithmic Personalisation: Although the object's biography may have little or no bearing on the resulting recommendation, it can profoundly affect how the listener experiences the interaction. The personal meaning creates attention, anticipation, reflection, and emotional investment.

3. A Productive Decoupling: This separation is not a failure of recognition. It is what allows the object to serve two distinct roles: it is a machine-readable material input and a personally nourishing source of meaning for the listener.

4. Dual-History Accumulation

Using personal possessions introduces a temporal dimension to the physical interface.

1. A New Role for the Everyday: An ordinary domestic object, such as a ceramic cup, gains a new digital purpose without ceasing to be what it originally was. The cup remains a cup, while also becoming a conduit into a musical landscape.

2. Layering of Meaning: Through repeated use, the object can acquire a second, experiential history: the music it has prompted, the paths it has opened, and the moods or places it has come to signify.

5. Making Ambiguity Meaningful

Intentionally ambiguous AI experience requires more than a conventional software prompt.

1. Withholding Requires Commitment: An AI system that interprets an input while deliberately withholding a literal explanation can keep the experience felt but not fully legible. For this withholding to become meaningful rather than merely confusing, it must be attached to an act in which the listener has made some form of commitment.

2. Software Alone Can Appear Infallible: Because software is intangible, its processes and limitations remain hidden. An unexplained recommendation can therefore appear as an authoritative system output, either a correct answer that the user cannot access or a malfunction that has failed to explain itself. Physical objects, by contrast, make interpretation tangible and situated. Because the listener can see what is being presented to the system, they can also recognise that the input is partial and open to misreading. The system's response can therefore be understood not as an authoritative answer, but as a fallible interpretation of a material prompt.

3. The Physical Anchor: Locating an object, choosing one that carries personal significance, and placing it onto the device creates a concrete investment in the encounter.

7.3 The Case for Effective Ambiguity

Contemporary streaming interfaces make it easier for listeners to describe what they want. Natural-language features such as Spotify's Prompted Playlist expand the type of requests that can be expressed. However, they still depend on the listener being able to externalise a listening direction in language. This is a limitation when the listener's desire is felt before it can be named.

Following Gaver et al. (2003), this thesis treats ambiguity as a design resource, inviting the user to participate in interpreting a system rather than receive a fixed meaning. Curio develops this argument in a specific direction: ambiguity becomes productive when it is grounded in an uncertainty that the listener already feels but cannot yet articulate. In this thesis, Stimmung refers to a un-articulated orientation that may become clearer through the listener's interaction with objects, system responses and oneself.

This distinguishes ambiguity from opacity. Opacity concerns what the system withholds about its operation; ambiguity concerns what remains open for the listener to interpret. A system that is vague because it hides a deterministic recommendation process does not create productive ambiguity; it instead creates confusion or dissatisfaction. Simply presenting a conventional algorithm through natural language does not give the listener meaningful interpretive space.

The explorations identified two boundaries that effective ambiguity must avoid.

1. At one boundary is resolution. In Exploration 2, the AI's step-by-step reasoning appeared transparent, but its effect was persuasive rather than liberating. The explanation made the recommendations seem justified enough that I accepted tracks I might otherwise have questioned.

2. At the opposite boundary is arbitrariness. In Exploration 1, loose mismatches between an object and a track initially seemed interesting, but quickly became tedious. Without a recognisable connection, an intentionally strange recommendation was indistinguishable from a technical failure.

Effective ambiguity must therefore produce an encounter that is unexpected but retrospectively recognisable as connected to something real (André et al., 2009).

The role of embodiment follows from this distinction. In the Vibe Map, the listener searches the room, selects a personal object and places it before the camera. This physical commitment gives the system's ambiguity a material and experiential anchor. The object does not provide a fixed meaning that the system decodes. Instead, the listener's choice establishes a situation from which an associative response can develop. Moving the object across the ambiguous output display is not merely an adjustment after the recommendation; it is the way the listener navigates the uncertainty.

This makes the Vibe Map and its development Short-Pay the clearest manifestation of effective ambiguity in the project. It does not offer a final answer or ask the listener to specify a mood in advance. It gives them a way to move through an unresolved state, using action to test, reject and extend the system's associations.

Effective ambiguity in this context therefore has three conditions:

1. It is anchored in the listener's unresolved experience, rather than manufactured through system opacity.

2. It is bounded by a recognisable connection, avoiding both over-explanation and randomness.

3. It is navigable through embodied action.

Under these conditions, ambiguity is not a flaw to be eliminated but a designed space between what the listener cannot yet articulate and what the system offers for interpretation. Meaning emerges through the interaction between the two.

7.4 The Transferable Method

// Reintroducing Nourishment Beyond Music

Within this thesis, nourishment is defined as a design condition that systematically sustains the listener's own active interpretive work, rather than placating them by settling uncertainty on their behalf.

Beyond the physical provotypes built for music discovery, this research produces transferable design knowledge that can be abstracted and applied to other fields.

//The Five-Step Transferable Method

This repeatable framework is designed for any domain in which a superseded physical practice has been replaced by AI system that resolves choices on the user's behalf.

This step supersedes all others. Before anything else, the researcher must locate the output unit. Different output units can exist within the same medium; in this project, the song and the album are two different output units, each requiring different inputs and processes to achieve a nourishing interaction. Because the output unit is fixed, it should be selected first. The rest of the interaction should then unfold according to the affordances required by that particular format.

The Cost of a Wrong Turn: The physical and cognitive effort an interface can demand of a user scales directly with the commitment cost of the unit they are choosing, this materialised in my product in this way:

- An album (long-play) represents a high temporal commitment (roughly 40 minutes)

Because the cost of being wrong is high, the user is less experimental

- A song (short-play) costs only 1-3 minutes to skip

Because the commitment cost is so low, the user is highly willing to explore

Subsequently, the starting question for exploration of this type should be: *"What is the unit of delivery, and what does committing to one of them* cost the user?"

1. Locate

The designer must personally re-enter the superseded experience and document what it offers experientially and what affordances generate these experiences, avoiding reliance on retrospective or idealised memories.

2. Name

The designer must set the experience down as an explicit ledger of losses, gains, and tensions between the two.

Crucially, this ledger must treat the digital system's gains, such as immediate access and low cost, as vital evidence rather than ignoring them. Nostalgia alone is not design evidence. Acknowledging what streaming and digital platforms successfully delivered is essential to understanding what a new system must not attempt to dismantle.

3. Decide

The designer must systematically sort each identified item into three categories: what to recover (from what has been lost), what to retain (from what is gained), and what to hold unresolved (what is not yet clear and must be realised through making) .

4. Translate

The designer must decide whether to materialise the digital or digitalise the material. This means bringing digital interpretation into a physical form while allowing the user's physical actions to shape the digital response. This will then materialise where the designer has placed agency between the user, the physical and the AI.

Neither approach works alone: a physical object without meaningful digital interpretation becomes decorative, while putting a physical practice onto a screen simply recreates the passive digital systems it aims to replace. The goal is a two-way exchange between bodily action and machine interpretation.

5. Provoke

The designer must then materialise their conclusion to either test or make discussable what position they have arrived at.

Theoretical Application

There are many domains that are highly suited to this methodology, where a rich superseded physical practice is now fully dominated by automated convenience. Below, I have included some examples to function as potential future provocations .

- Film Discovery: The Video Rental Store vs Netflix-type streaming service.
- Wayfinding: Finding your way via a paper map vs using a Google Maps-type application.
- Dating: In-person meeting people at a bar vs Tinder-type application.

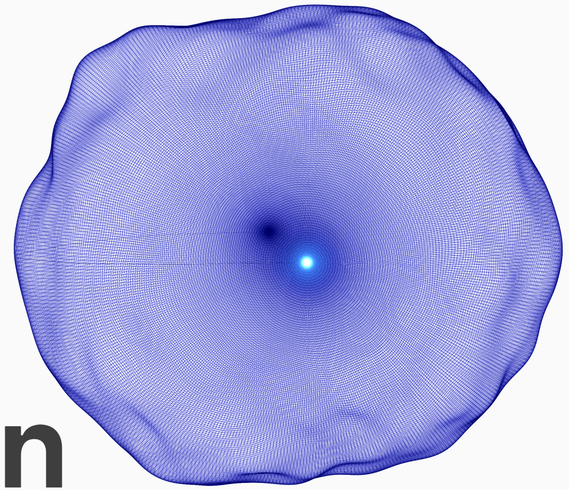
Methodological Boundaries

It is important to state that not every domain should choose nourishment over placation.

A task with a single, objectively correct, or highly efficient answer, such as looking up flight times or booking a medical appointment, thrives on frictionless prediction and should remain strictly optimised for placation.

The five-step method is specifically reserved for discovery experiences, where the value of the encounter lies in the journey of interpretation rather than the speed of the transaction.

Chapter 8



Conclusion

This chapter answers the research questions, states the project's contribution, sets out its limitations, and closes with a personal reflection on the work.

Answering the Research Questions
Reflection & Limitations
Personal Reflection

8.1 Answering the Research Questions

The central argument of this project is that music discovery can return interpretive work to the listener without rejecting the abundance and convenience of digital streaming. Doing so means rethinking what recommendation is meant to do / to create. Currently, mainstream recommendation systems work through prediction and prescription. They use past behaviour to anticipate what the listener might want and resolve uncertainty on the listener's behalf (Schedl et al., 2018). This thesis describes that mode as placation: the system soothes uncertainty by providing a finished answer, but risks leaving the listener passive.

Curio sets out an alternative: attunement. It begins with a present, physical input and deliberately keeps uncertainty open, allowing the listener to explore a mood or atmosphere that they may not yet be able to describe. This thesis calls that mode nourishment: a design structure that sustains the listener's interpretive work rather than replacing it. The three sub-questions listed below have aimed to address part of this question. The physical object provides an embodied, pre-declarative starting point; the AI provides an associative and qualitative reading rather than a definitive interpretation; and the listener remains responsible for connecting these elements and constructing meaning.

//The Primary Research Question

How might a domestic, object-based music system use associative AI to support music discovery without requiring the listener to declare or resolve their listening state before the encounter?

This project shows that a domestic, object-based music system can support discovery by returning the active interpretive work of browsing and choice to the listener, without rejecting the convenience and abundance of modern streaming platforms. Curio does this by replacing the text-based search query with a personal household object used as a pre-declarative, embodied input. Rather than trying to model or classify the listener's internal state (which can be defined as a felt, pre-verbal atmosphere, or Stimmung, that resists explicit naming; Gumbrecht, 2012), the system treats the listener's inability to declare what they want as the starting condition of the interaction. Uncertainty is not treated as a deficiency to be algorithmically resolved, but as the condition within which discovery can begin.

This happens through the object itself, as it provides this uncertainty with a physical form. The listener must move through their home, notice something, select it, and commit to placing it on the system. This gesture creates an embodied starting point before any musical destination has been chosen. The object does not work as a stable query key, but provides an occasion for interpretation.

Exploration 1 also showed that discovery can begin before the music starts. Moving through the home, selecting an object, and deciding which object represented the current moment became part of the discovery process itself. The value of the experience therefore shifted partly from the recommendation received to the search undertaken. In this sense, Curio responds to what streaming often removes: the need to form a question before receiving an answer.

Curio then uses associative generative AI as an interpretive design material rather than as a prescriptive decision-maker. The AI is prompted to read only the expressive and qualitative features of the object while being prevented from naming the object or predicting the listener's preferences. This opens a generative gap between the private significance that the listener invests in the object and the associative reading returned by the machine. Because the machine's response stays ambiguous, it works as an evocative response rather than an authoritative choice, and the incompleteness is where the listener's own interpretation goes. The listener does the interpretive work of connecting the two and building their own meaning.

The project therefore reframes the purpose of music recommendation. Recommendation need not eliminate uncertainty before listening begins. It can instead provide a material and associative structure through which uncertainty becomes navigable. In Curio, the machine opens a field of possible meaning, but the listener decides whether and how that field becomes personally resonant.

//SQ1 · What was lost? - The Record Store

What qualities does record-based discovery, ownership, and use require from listeners and provide to them that streaming has removed? How might these qualities be translated into streaming without reproducing the limitations of physical media?

Answering this question means moving past a superficial craving for obsolete plastic. Through a duo-ethnographic record shop journey and a review of the literature, this project identified nine distinct dimensions of loss, LD1 to LD9. Crucially, the audit revealed that what has been lost is not physical vinyl itself, but the structural requirement for physical and cognitive effort on the part of the listener.

Physical record shops and vinyl playback required bodily movement: walking to a shop, scanning physical crates, handling authored album sleeves, paying a high temporal and financial cost, and committing to listening to an album whole and in order. These physical boundaries gave the listener's attention a physical anchor, transforming browsing from a transaction into a meaningful ritual. Modern streaming platforms treat all pre-playback action as friction to be engineered out. By automatically predicting what a listener wants on the basis of past behaviour, they resolve choices on the listener's behalf and maximise transactional efficiency at the cost of cognitive passivity, the state this research defines as placation. Translating these lost qualities cannot mean a nostalgic retreat into vinyl's scarcity. Listeners are unwilling to relinquish streaming's primary gains: abundance, instant access, portability, low cost and search. Translation must instead be understood as a deliberate redistribution of attention, choice and interpretation within abundance itself.

The decisive theoretical finding of this thesis is that translating these qualities is not a uniform, undifferentiated interaction-design task. It requires a systematic assessment of the temporal scale of the music medium, which dictates distinct configurations of physical and cognitive effort.

8.1 Answering the Research Questions

The album as a high-commitment unit (Provotype 3, Long Play)

An album represents roughly forty minutes of attention. Since the commitment cost is high, the interaction model can demand slower and higher-effort physical gestures to match those stakes. The finding from the explorations, that composing and withholding read as authorship, is carried here by the objects themselves. The listening persona is applied to and subsequently sits in the object, so a persona can be brought to the system or deliberately held back simply by choosing which object to carry over. The object's arrangement slowly steers a drifting background library. As an album sleeve already possesses an authored and atmospheric visual surface, Long Play deliberately forgoes screens and lets the record album art cover carry the pre-declarative framing. Lifting the stopper off its resting surface commits the listener to that album, played whole and in order, collapsing the AI's active agency to zero.

The song as a low-commitment unit (Provotype 4, Short Play)

A single song costs only three minutes to skip. Because leaving carries no risk, continuous low-cost exploration is strongly afforded. The interaction model must therefore be fluid, continuous and highly navigable, allowing the listener's hand to wander constantly across a responsive landscape without fear of a wrong turn. Short Play projects a label-free coordinate field of qualitative regions drawn from the object's facets. Position acts as the query, and the music and visual textures shift continuously with the hand. There is no final arrival and no playback resolution. The physical movement across the unlabelled field is the discovery encounter.

By scaling physical effort to the temporal granularity of the media unit, the physical constraints of the interface are experienced as nourishment rather than as frustrating, arbitrary barriers.

// SQ2 · How can this be brought back? - Physical Objects

How can AI-generated associations between a personal object and music remain meaningful and traceable without reducing the object to a single authoritative interpretation?

The explorations, particularly Exploration 2, produced an unexpected result: making the system's reasoning fully legible did not empower the listener, as it could actually placate them.

When the AI gave a clear, step-by-step explanation of why it had matched a personal object to a track, the explanation became an instrument of persuasion. Its tidy justification narrowed the listener's role to accepting or rejecting the machine's verdict. Rather than supporting interpretation, too much legibility encouraged passive agreement. This does not mean that maximum ambiguity is always desirable. Too much disclosure can make an association feel resolved before the music is heard, but too little disclosure creates the opposite problem. When the system provides no basis for interpreting the difference, an unexpected result and a poor result become indistinguishable. Therefore, the useful position sits between complete transparency and complete opacity. An association should be unexpected enough to preserve discovery, but connected enough that the listener can recognise the relationship retrospectively, which is the balance André et al. (2009) describe as unexpected but retrospectively recognisable as connected. The design problem is not to maximise either explanation or ambiguity, but to judge how much of the system's reading has to be exposed for the listener to stay engaged without being directed towards a predetermined interpretation.

The result is a design that is felt but not legible. The listener should be able to sense a connection between the object and the music without receiving a complete explanation of that connection. Withholding the account is not a shortfall in the interface; it is what obliges the listener to take part in making the meaning. To hold this balance, Curio relies on three structural decisions.

Withholding the name

After Exploration 1, the AI is never allowed to identify the object. By restricting its analysis to abstract material qualities, the system stays blind to the object's literal identity. This prevents direct puns, simplistic mappings and predictable associations.

Instability was therefore used as a design resource, but a conditional one; enough variation holds the encounter open and invites the listener to project their own mood into the reading. Too much severs the link between the object and the result, and the association stops being ambiguous and becomes arbitrary. Placation and arbitrariness are the two walls of the same corridor, and the object that is brought to the system largely decides which wall the encounter drifts towards.

Anchoring ambiguity through physical commitment

An ambiguous AI response might seem like poor usability when presented on a screen alone. However, the physical object gives that ambiguity a meaningful context. By searching for and placing an object with personal significance, the listener contributes their own history to the interaction. This makes the AI's associative response feel personally resonant, while leaving the listener to decide what it ultimately means

// SQ3 · Where should interpretation sit? - AI as a Collaborator

How can interpretive work be distributed among the listener, their object and the AI? And what does each arrangement mean for the listener's agency over where the discovery goes?

The project answers this question through the contrasts mapped across the annotated portfolio's four provotypes, two of which were built and two of which were designed and visualised in full. They are legible as a set precisely because they sit at different points on the same axis. Interpretive agency is not a fixed property belonging to either the human or the machine. It moves according to how the physical interaction and the AI pipeline are coupled.

In this project, agency means control over where the discovery goes next, rather than a general feeling of being in charge. Physical interaction alone does not create agency if the system still decides the interpretation and the musical result. Read across the portfolio, the arrangement resolves into two separate agencies rather than one shared quantity. The listener controls the supply: which object is brought to the system, what history it carries, and when the encounter is committed to. The system controls the ambiguity: how much of its reading it exposes, and whether it proposes a destination or withholds one. A design distributes interpretive work by deciding which of these two it hands over, and neither one on its own is sufficient.

Provotype 1: The Space Between

The Space Between occupies the placation pole, and it was built first so that the rest of the portfolio had something to be read against.

The object sits between the listener and the system. The machine reads it clearly and directly from the front, while the listener's own view of the same object is obscured. The gap between what the machine can see

8.1 Answering the Research Questions

and what the listener can see is the whole subject of the piece, and it is made physical rather than argued for. The machine then resolves its clear reading into a literal musical match and hands over the musical output as finished.

This provotype serves as an honest foil: it does not caricature mainstream recommendation but performs it effectively. It looks polished, responds quickly, and leaves little for the listener to do. It represents one end of the design space, where the system controls both what the result means and where discovery leads. Its smoothness reveals the cost: the listener has little ownership of the encounter.

Provotype 2: Finding Resonance

Finding Resonance addresses the limits of transparency by making the system's interpretation adjustable rather than simply revealing it. Its slider moves between step-by-step explanation and atmospheric reading, showing that explanation can either settle the listener or leave room to question the relationship. This demonstrates that agency depends not simply on how much the system reveals, but on where interpretive work is placed: the object provides an embodied starting point, the AI offers an associative reading, and the listener decides whether it resonates and how it should guide discovery. Stimmung is therefore not a problem for the system to resolve, but an uncertainty the interaction must hold open, enough to support recognition without dictating meaning. When the system closes that uncertainty, the result may be placation; when it holds it open and requires listener participation, it creates the possibility of nourishment.

Provotype 3: Long Play

Long Play materialises what nourishment looks like when paired with the demands of the album as an output medium.

Lifting the object commits to an album, played whole and in order. The listening persona is not held in the system or in any separate token; it sits in the objects themselves, associated onto them over time by the person who owns them. A persona can therefore be brought to the system or deliberately held back, simply by choosing which object to carry over. An album sleeve already carries an atmospheric representation of the music, so it is left to do the pre-declarative work.

Here, agency is expressed through physical authorship and ownership. Slow, deliberate gestures are matched to the format that asks most of the listener. The listener holds the supply, choosing which region of the library is in play and when the commitment is made, rather than accepting a finished recommendation.

Provotype 4: Short Play

Short Play materialises what nourishment pole looks like when paired with the demands of the single track.

It reads the object once in full and projects a responsive, unlabelled coordinate landscape drawn from the material qualities of the object the system is permitted to see. No region is named, and no destination is proposed. To explore the map, the listener has to keep moving the physical object across its surface, navigating the AI's qualitative reading through spatial movement and a two-way AI-to-listener call and response.

The listener therefore makes sense of the output from within their own uncertainty. The AI supplies an associative terrain, the object anchors the interaction, and the physical movement of the object across the map is the encounter itself. Short Play offers directional cues without ever prescribing an endpoint.

The commitment the medium can carry

Reading the four provotypes as a set produced the portfolio's most consequential finding, and it was not one the project set out to look for. Interpretive work cannot be distributed once and then applied uniformly. The distribution has to be set against the temporal size of the media unit being recommended, because that size determines how much physical and cognitive commitment an interface is entitled to ask for. There is no single interface for nourishment. The medium has to be assessed first, and the interaction model adjusted to it. An album asks for roughly forty minutes of attention. Because the stakes of committing are high, the interaction can afford to be slow, deliberate and discrete. Long Play therefore separates orientation from commitment: object arrangement orients the library over time, and a single lift commits to the whole record. The cost of the gesture is proportionate to the cost of the listening it initiates. A single track asks for roughly three minutes. Because leaving costs almost nothing, continuous exploration is not merely tolerable but desirable, and the interaction has to be fluid, continuous and cheap to abandon. Short Play therefore offers a landscape the hand can wander across without fear of a wrong turn, and never proposes a terminal destination at all.

Long Play and Short Play are not two attempts at the same problem. They are the same principle held against two different commitment costs, and the difference between them is not stylistic. It follows from the medium.

Across the portfolio, the object shifts from being primarily something the system reads to also being something the listener uses to steer discovery. The object does not exchange one role for another; it gains a second role while remaining an input to the system.

The four provocations therefore distribute agency differently:

1. The Space Between withholds control over both explanation and musical direction.
2. Finding Resonance provides control over the explanation but not over the selection.
3. Long Play provides control over the region of the library and over the moment of commitment.
4. Short Play provides direction without ever prescribing an endpoint.

The portfolio shows that interpretive work can be distributed in different ways:

- The object provides the embodied and personally situated starting point.
- The AI provides an associative, incomplete and revisable reading.
- The listener decides how to connect that reading to their own mood, memory and intentions.

The listener's agency is strongest when the system does not settle this relationship in advance, and when the demand the interface makes is proportionate to the commitment the medium actually requires.

8.1 Answering the Research Questions

// Concluding Synthesis

This thesis shows that music discovery can return interpretive work to the listener without giving up the main advantages of digital streaming.

By bringing personal domestic objects together with generative artificial intelligence, Curio challenges the prediction paradigm that dominates modern recommender systems. Rather than treating a listener's initial uncertainty as a deficiency to be resolved immediately by prediction, the project treats the state of not yet knowing what you want, a felt Stimmung that resists naming, as the condition within which design must operate.

The discovery process begins before the music starts: with the listener's movement through the home, selection of an object, and decision to commit it to the system. Curio therefore shifts some value away from the recommendation received and towards the interpretive search undertaken.

The primary contribution is therefore transferable design knowledge: a conceptual and physical demonstration that an automated recommender can be structured to support interpretive nourishment rather than passive placation.

The four provotypes show that interpretive agency does not belong automatically to any one point in a human-AI system. It shifts according to how a design holds and distributes ambiguity. The Space Between provides a baseline of placation by resolving the machine's reading into a destination. Finding Resonance hands the listener the account of the result but not the direction of the discovery. Long Play and Short Play then divide between them the finding that the portfolio arrived at last: that the commitment an interface can ask for is set by the temporal size of what it recommends, so the album is given slow, discrete, committing gestures and the single track is given a continuous landscape that costs nothing to leave.

Used this way, AI implementation does not diminish interpretation. It makes room for listeners to attend more closely to their homes, their objects, their music and themselves.

8.2 Reflection & Limitations

The limitations identified below clarify the kinds of claims this research can and cannot make. Several of these boundaries are deliberate scoping decisions rather than oversights discovered after the fact.

The Absence of Formal Participant Evaluation

The limitation: There is no formal participant study validating the provotypes in this report. Two encounters with other people did take place: the record-store journey and Test 2 (Exploration 2), which included a second participant because its testing goal could not be evaluated by myself alone.

The scope of the claims is also limited to the intended listener described in Section 1.2. This intended listener is a design construct rather than a validated demographic or behavioural category; the project does not make claims about listeners in general.

The record-store journey occupies a specific position within the evidence. It functioned both as a sensitising task and as a means of generating experiential material to guide later explorations. It should not, however, be treated as a participant evaluation. The journey was walked twice: first alone, to establish its structure, and then with a close friend. Although notes were taken during the mapping, most of the information was recorded afterwards. The mapping itself took place approximately two hours after its first stage, which introduced a retrospective element shaped by memory and interpretation. The journey therefore provided situated experiential material for design development, but it cannot support representative claims or be read as a set of directly comparable measurements.

The rationale: This project does not seek to establish universal laws of listener preference or make population-level claims about what listeners in general want. Within an RtD framework, its aim is to produce situated design knowledge. Testing highly conceptual provotypes with unfamiliar participants in a short session would primarily capture first impressions and novelty, rather than how the designs might support sustained experience.

Importantly, this rationale defines the scope of the contribution; it does not remove the need for later participant research.

The Auto-ethnographic Frame

The limitation: The researcher's musical taste, background, and cultural assumptions shaped the selection of objects, prompts, design decisions, and playlists throughout the project. Because the researcher also designed, tested, and interpreted the outcomes, the process is vulnerable to confirmation bias.

The rationale: In a project exploring subjective and partly pre-verbal experiential states such as Stimmung, first-person engagement is not simply a bias to be eliminated. It functions as a research instrument. Following the principles of autobiographical design, using oneself as the primary participant enabled sustained and reflective engagement with the material.

Where a judgement could not be made alone, a second person was brought in. Test 2 relied on a second participant precisely because its testing goal could not be validated from a single account. This does not remove the influence of the researcher's preferences or eliminate confirmation bias; it makes those influences explicit and establishes a boundary around the knowledge produced.

Uneven Materialisation and Technical Development

The limitation: The six design explorations and four provotypes do not carry equal evidential weight. Provotypes 3 and 4 were never engineered as working artefacts; they remain visual propositions. They therefore provide evidence of conceptual positions, but not of interactional performance, technical feasibility, or sustained use.

The rationale: An RtD portfolio does not require every outcome to have the same level of technical fidelity. The difference between the four provotypes is functional rather than simply technical. The Space

8.2 Reflection & Limitations

Between and Finding Resonance operate as materialised reflections for examining experienced realities, whereas Long Play and Short Play function as speculative propositions for modelling future possibilities.

Long Play and Short Play were never intended to become working artefacts. They are research culminations whose arguments are carried by visualisations rather than demonstrated through use. This decision was also shaped by the limited time, cost, and technical resources of the project. Developing these concepts further would have required additional hardware and software integration, AI testing, and music-service dependencies.

Consequently, these provocations cannot support claims about practical feasibility, reliability, or user behaviour. They support claims about the questions and positions that the designs make visible.

The Absence of Comparative Evaluation

The limitation: The project does not compare Curio directly with a conventional streaming interface or an existing recommendation system. It therefore cannot establish that the approach improves music discovery, listener agency, or meaningfulness in comparison with current alternatives.

The rationale: The project investigates what an object-based, associative system might make possible rather than measuring its performance against an established baseline. Its evaluation is therefore concerned with the qualities and questions opened by the design, rather than with demonstrating superior recommendation outcomes. The absence of comparative testing limits the claims to the design proposition itself; it does not show that Curio is more effective than conventional streaming.

Technical, Platform, and Data Dependency

The limitation: By using Spotify as the music source, Curio was limited to the user data, permissions, and API functions made available by the platform. The system could retrieve music through Spotify, but it could not access much of the users' metadata. This limited recommendation quality, as the system had less data to base these recommendations.

The rationale: Retaining Spotify allowed the project to work with a familiar and functioning music catalogue, keeping the focus on the relationship between object, AI interpretation, and music discovery rather than on building a better recommendation system. The findings are therefore situated within the data access and technical conditions Spotify provided during the project. They do not establish how Curio would behave with richer user data, another platform, or an independent catalogue.

Generalisability of the Five-Step Method

The five steps: Locate, Name, Decide, Translate, and Provoke were derived within the context of music discovery but are intentionally formulated as transferable design orientations. The central claim, that the physical granularity of a recommended unit should scale with its cognitive commitment cost, proposes a possible design guideline for systems seeking to reintroduce nourishment within digital abundance.

However, the method was articulated retrospectively from this project rather than specified in advance and independently tested. It has also not been applied outside music discovery. Its broader generalisability therefore remains an open proposition rather than an established finding. Future work would need to examine whether the five steps remain useful in other domains.

8.3 Personal Reflection

I began this journey believing I was designing towards a clear outcome; through the Research through Design process, however, that certainty was productively undone. Upon reflection, I am glad this was the case. Had I arrived at my starting point, there would have been little reason to undertake this research at all. Like many other designers working with AI, I had taken on the assumption that making a system's reasoning transparent was the clearest route to a good interaction. I believed that if a system could logically explain away the listener's Stimmung, this clarity would therefore increase the listener's feeling of agency over the system and therefore improve the interaction. However, the Research through Design process challenged that assumption. Through building and testing in my exploration phases, complete transparency and structured explanations hijacked the listener's authority. A tidy, logical machine justification "placated" the listener, settling them so efficiently that it flattened their own capacity to feel, question, and explore. Curio taught me to trust my instincts as a designer and embrace ambiguity as a profound resource. By looking at music streaming from a unique angle, I found that an AI which responds to abstract input with abstract outputs does not leave the user feeling isolated; it instead invites them back into the room. It transforms the passive act of choosing music into a slow, embodied ritual of self-attunement. Ultimately, this project has fundamentally shaped my philosophy as a designer: AI does not need to resolve our human uncertainty on our behalf; instead, design can build spaces where this uncertainty provides texture, texture that we can enjoy the feel of, instead of the smoothness that is increasingly served to us.

8.4 Refrences

Adomavicius, G., & Tuzhilin, A. (2005). Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. IEEE Transactions on Knowledge and Data Engineering, 17(6), 734–749. https://doi.org/10.1109/TKDE.2005.99

Anderson, A., Maystre, L., Anderson, I., Mehrotra, R., & Lalmas, M. (2020). Algorithmic effects on the diversity of consumption on Spotify. In Proceedings of The Web Conference 2020 (pp. 2155–2165). Association for Computing Machinery. https://doi.org/10.1145/3366423.3380281

Andjelkovic, I., Parra, D., & O'Donovan, J. (2019). Moodplay: Interactive music recommendation based on artists' mood similarity. International Journal of Human-Computer Studies, 121, 142–159. https://doi.org/10.1016/j.ijhcs.2018.04.004

André, P., schraefel, m. c., Teevan, J., & Dumais, S. T. (2009). Discovery is never by chance: Designing for (un)serendipity. In Proceedings of the Seventh ACM Conference on Creativity and Cognition (pp. 305–314). Association for Computing Machinery. https://doi.org/10.1145/1640233.1640279

Arevalo Escobar, S., Chaudhary, A., Ehrlich, D., & Smith, A. M. (2023). Creativity and consistency in musical perception of tangible objects (Report No. UCSC-SOE-23-08). University of California, Santa Cruz. https://tr.soe.ucsc.edu/research/technical-reports/UCSC-SOE-23-08

Bartmanski, D., & Woodward, I. (2015). The vinyl: The analogue medium in the age of digital reproduction. Journal of Consumer Culture, 15(1), 3–27. https://doi.org/10.1177/1469540513488403

Belk, R. W. (1988). Possessions and the extended self. Journal of Consumer Research, 15(2), 139–168. https://doi.org/10.1086/209154

Belkin, N. J. (1980). Anomalous states of knowledge as a basis for information retrieval. Canadian Journal of Information Science, 5, 133–143.

Bilgic, M., & Mooney, R. J. (2005). Explaining recommendations: Satisfaction vs. promotion. In M. van Setten, S. McNee, & J. Konstan (Eds.), Beyond Personalization 2005: A workshop on the next stage of recommender systems research (pp. 13–18). https://grouplens.org/beyond2005/bp2005.pdf

Blumer, H. (1954). What is wrong with social theory? American Sociological Review, 19(1), 3–10. https://doi.org/10.2307/2088165

Campolo, A., & Crawford, K. (2020). Enchanted determinism: Power without responsibility in artificial intelligence. Engaging Science, Technology, and Society, 6, 1–19. https://doi.org/10.17351/ests2020.277

Celma, Ò. (2010). Music recommendation and discovery: The long tail, long fail, and long play in the digital music space. Springer. https://doi.org/10.1007/978-3-642-13287-2

Cheng, Z., & Shen, J. (2016). On effective location-aware music recommendation. ACM Transactions on Information Systems, 34(2), Article 13. https://doi.org/10.1145/2846092

Chiquet, H., Evéquoz, F., & Lalanne, D. (2006). Elcano, a tangible personal multimedia browser. In Adjunct proceedings of the 19th annual ACM Symposium on User Interface Software and Technology (pp. 51–52). Association for Computing Machinery. https://uist.acm.org/archive/adjunct/2006/pdf/demos/p51-chiquet.pdf

Chung, J. (2022). Sound frame [Design prototype]. ECAL/University of Art and Design Lausanne, in collaboration with Yamaha Design Laboratory. https://jisanchung.com/sound-frame

Cirisano, T. (2026, June 15). Music subscriber market shares Q4 2025: The chess board is set. MIDiA Research. https://www.midiaresearch.com/blog/music-subscriber-market-shares-q4-2025-the-chess-board-is-set

Csikszentmihalyi, M., & Rochberg-Halton, E. (1981). The meaning of things: Domestic symbols and the self. Cambridge University Press. https://doi.org/10.1017/CBO9781139167611

DeNora, T. (2000). Music in everyday life. Cambridge University Press. https://doi.org/10.1017/CBO9780511489433

Desmet, P. M. A. (2018). Measuring emotion: Development and application of an instrument to measure emotional responses to products. In M. Blythe & A. Monk (Eds.), Funology 2: From usability to enjoyment (pp. 391–404). Springer. https://doi.org/10.1007/978-3-319-68213-6_25

Dong, M., & Chen, P. (2019). Incorporating facial expressions and social networks for musical contextual recommendations. Computer Science and Application, 9(5), 855–863. https://doi.org/10.12677/CSA.2019.95096

Dourish, P. (2001). Where the action is: The foundations of embodied interaction. MIT Press. https://doi.org/10.7551/mitpress/7221.001.0001

Eisenberg, E. (1987). The recording angel: Explorations in phonography. McGraw-Hill.

Epure, E. V., Deldjoo, Y., Sguerra, B., Schedl, M., & Moussallam, M. (2025). Music recommendation with large language models: Challenges, opportunities, and evaluation. arXiv. https://doi.org/10.48550/arXiv.2511.16478

Eriksson, M., Fleischer, R., Johansson, A., Snickars, P., & Vonderau, P. (2019). Spotify teardown: Inside the black box of streaming music. MIT Press. https://doi.org/10.7551/mitpress/10932.001.0001

Fernandez, K. V., & Beverland, M. B. (2019). As the record spins: Materialising connections. European Journal of Marketing, 53(6), 1152–1175. https://doi.org/10.1108/EJM-12-2016-0828

Frohlich, D. M., Clancy, T., Robinson, J., & Costanza, E. (2004). The audiophoto desk. In Proceedings of 2AD: Second International Conference on Appliance Design (pp. 139–140). Appliance Design Network.

Garcia-Gathright, J., St. Thomas, B., Hosey, C., Nazari, Z., & Diaz, F. (2018). Understanding and evaluating user satisfaction with music discovery. In Proceedings of the 41st International ACM SIGIR Conference on Research & Development in Information Retrieval (pp. 55–64). Association for Computing Machinery. https://doi.org/10.1145/3209978.3210049

Gaver, W. W. (2012). What should we expect from research through design? In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (pp. 937–946). Association for Computing Machinery. https://doi.org/10.1145/2207676.2208538

Gaver, W. W., Beaver, J., & Benford, S. (2003). Ambiguity as a resource for design. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (pp. 233–240). Association for Computing Machinery. https://doi.org/10.1145/642611.642653

Gaver, W. W., & Bowers, J. (2012). Annotated portfolios. Interactions, 19(4), 40–49. https://doi.org/10.1145/2212877.2212889

Gibson, J. J. (1979). The ecological approach to visual perception. Houghton Mifflin.

Goto, M., & Goto, T. (2009). Musicream: Integrated music-listening interface for active, flexible, and unexpected encounters with musical pieces. Journal of Information Processing, 17. 292–305. https://doi.org/10.2197/ipsjjip.17.292

Greb, F., Steffens, J., & Schlotz, W. (2019). Modeling music-selection behavior in everyday life: A multilevel statistical learning approach and mediation analysis of experience sampling data. Frontiers in Psychology, 10, Article 390. https://doi.org/10.3389/fpsyg.2019.00390

Guerra, P. (2022). House of golden records: Portugal's independent record stores (1998–2020). Societies, 12(6), Article 188. https://doi.org/10.3390/soc12060188

Gumbrecht, H. U. (2012). Atmosphere, mood, Stimmung: On a hidden potential of literature (E. Butler, Trans.). Stanford University Press. (Original work published 2011)

Hallnäs, L., & Redström, J. (2001). Slow technology – Designing for reflection. Personal and Ubiquitous Computing, 5(3), 201–212. https://doi.org/10.1007/PL00000019

Heidegger, M. (1962). Being and time (J. Macquarrie & E. Robinson, Trans.). Harper & Row. (Original work published 1927)

Hracs, B. J., & Jansson, J. (2020). Death by streaming or vinyl revival? Exploring the spatial dynamics and value-creating strategies of independent record shops in Stockholm. Journal of Consumer Culture, 20(4), 478–497. https://doi.org/10.1177/1469540517745703

8.4 Refrences

Hsia, C.-C., Lai, K.-H., Chen, Y., Wang, C.-J., & Tsai, M.-F. (2018). Representation learning for image-based music recommendation. arXiv. https://doi.org/10.48550/arXiv.1808.09198

Ishii, H., & Ullmer, B. (1997). Tangible bits: Towards seamless interfaces between people, bits and atoms. In Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (pp. 234–241). Association for Computing Machinery. https://doi.org/10.1145/258549.258715

Jang, S., Lee, W., Kim, B., Odom, W., & Park, Y.-W. (2022). Encountering cover versions of songs derived from personal music-listening history data: A design and field trial of Musée in homes. Interacting with Computers, 34(1), 24–42. https://doi.org/10.1093/iwc/iwac027

Janssen, J. H., van den Broek, E. L., & Westerink, J. H. D. M. (2012). Tune in to your emotions: A robust personalized affective music player. User Modeling and User-Adapted Interaction, 22(3), 255–279. https://doi.org/10.1007/s11257-011-9107-7

Jin, Y., Htun, N. N., Tintarev, N., & Verbert, K. (2019). ContextPlay: Evaluating user control for context-aware music recommendation. In Proceedings of the 27th ACM Conference on User Modeling, Adaptation and Personalization (pp. 294–302). Association for Computing Machinery. https://doi.org/10.1145/3320435.3320445

Jordà, S., Geiger, G., Alonso, M., & Kaltenbrunner, M. (2007). The reacTable: Exploring the synergy between live music performance and tabletop tangible interfaces. In Proceedings of the 1st International Conference on Tangible and Embedded Interaction (pp. 139–146). Association for Computing Machinery. https://doi.org/10.1145/1226969.1226998

Kaminskas, M., & Ricci, F. (2012). Contextual music information retrieval and recommendation: State of the art and challenges. Computer Science Review, 6(2–3), 89–119. https://doi.org/10.1016/j.cosrev.2012.04.002

Kaminskas, M., Ricci, F., & Schedl, M. (2013). Location-aware music recommendation using auto-tagging and hybrid matching. In Proceedings of the 7th ACM Conference on Recommender Systems (pp. 17–24). Association for Computing Machinery. https://doi.org/10.1145/2507157.2507180

Karana, E., Pedgley, O., & Rognoli, V. (2015). On materials experience. Design Issues, 31(3), 16–27. https://doi.org/10.1162/DESI_a_00335

Karmann, B. (2023). Paragraphica [Design project]. https://bjoernkarmann.dk/project/paragraphica

Kotkov, D., Wang, S., & Veijalainen, J. (2016). A survey of serendipity in recommender systems. Knowledge-Based Systems, 111, 180–192. https://doi.org/10.1016/j.knosys.2016.08.014

Kwon, D., & Lee, W. (2018). Artifact mixtape: Curating music in personal tangible artifacts. In Proceedings of the 2018 ACM Conference Companion Publication on Designing Interactive Systems (pp. 265–270). Association for Computing Machinery. https://doi.org/10.1145/3197391.3205447

Laurans, G., & Desmet, P. M. A. (2017). Developing 14 animated characters for non-verbal self-report of categorical emotions. Journal of Design Research, 15(3–4), 214–233. https://doi.org/10.1504/JDR.2017.089903

Lehtiniemi, A., & Holm, J. (2012). Using animated mood pictures in music recommendation. In 2012 16th International Conference on Information Visualisation (pp. 143–150). IEEE. https://doi.org/10.1109/IV.2012.34

Li, J., He, Z., Cui, Y., Wang, C., Chen, C., Yu, C., Zhang, M., Liu, Y., & Ma, S. (2022). Towards ubiquitous personalized music recommendation with smart bracelets. Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies, 6(3), Article 125. https://doi.org/10.1145/3550333

Lozano Murciego, Á., Jiménez-Bravo, D. M., Valera Román, A., De Paz Santana, J. F., & Moreno-Garcia, M. N. (2021). Context-aware recommender systems in the music domain: A systematic literature review. Electronics, 10(13), Article 1555. https://doi.org/10.3390/electronics10131555

Lupetti, M. L., & Murray-Rust, D. (2024). (Un)making AI magic: A design taxonomy. In Proceedings of the CHI Conference on Human Factors in Computing Systems (Article 1, pp. 1–21). Association for Computing Machinery. https://doi.org/10.1145/3613904.3641954

Magaudda, P. (2011). When materiality "bites back": Digital music consumption practices in the age of dematerialization. Journal of Consumer Culture, 11(1), 15–36. https://doi.org/10.1177/1469540510390499

McNee, S. M., Riedl, J., & Konstan, J. A. (2006). Being accurate is not enough: How accuracy metrics have hurt recommender systems. In CHI '06 extended abstracts on human factors in computing systems (pp. 1097–1101). Association for Computing Machinery. https://doi.org/10.1145/1125451.1125659

Melchiorre, A. B., Penz, D., Ganhör, C., Lesota, O., Fragoso, V., Fritzl, F., Parada-Cabaleiro, E., Schubert, F., & Schedl, M. (2023). Emotion-aware music tower blocks (EmoMTB): An intelligent audiovisual interface for music discovery and recommendation. International Journal of Multimedia Information Retrieval, 12(1), Article 13. https://doi.org/10.1007/s13735-023-00275-8

Neustaedter, C., & Sengers, P. (2012). Autobiographical design in HCI research: Designing and learning through use-it-yourself. In Proceedings of the Designing Interactive Systems Conference (pp. 514–523). Association for Computing Machinery. https://doi.org/10.1145/2317956.2318034

Nieuwpoort, S. A. (2017). Het kopen van vinyl [Buying vinyl] [Master's thesis, Erasmus University Rotterdam]. Erasmus University Thesis Repository. https://thesis.eur.nl/pub/40346

Norman, D. A. (1988). The psychology of everyday things. Basic Books.

North, A. C., & Hargreaves, D. J. (2008). The social and applied psychology of music. Oxford University Press.

Odom, W., & Duel, T. (2018). On the design of OLO Radio: Investigating metadata as a design material. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (Article 104, pp. 1–9). Association for Computing Machinery. https://doi.org/10.1145/3173574.3173678

Pariser, E. (2011). The filter bubble: What the Internet is hiding from you. Penguin Press.

Park, J., Kim, M., & Kim, H. Y. (2024). Image is all for music retrieval: Interactive music retrieval system using images with mood and theme attributes. International Journal of Human–Computer Interaction, 40(14), 3841–3855. https://doi.org/10.1080/10447318.2023.2201557

Park, J., Shin, H., Oh, C., & Kim, H. Y. (2024). "Is text-based music search enough to satisfy your needs?" A new way to discover music with images. In Proceedings of the CHI Conference on Human Factors in Computing Systems (Article 504, pp. 1–21). Association for Computing Machinery. https://doi.org/10.1145/3613904.3642126

Pichl, M., & Zangerle, E. (2021). User models for multi-context-aware music recommendation. Multimedia Tools and Applications, 80(15), 22509–22531. https://doi.org/10.1007/s11042-020-09890-7

Prey, R. (2018). Nothing personal: Algorithmic individuation on music streaming platforms. Media, Culture & Society, 40(7), 1086–1100. https://doi.org/10.1177/0163443717745147

Rayzhekov, A., & Murer, M. (2025). The well-tempered objects: Using everyday things for tangible composition. In Proceedings of the Nineteenth International Conference on Tangible, Embedded, and Embodied Interaction (Article 80, pp. 1–7). Association for Computing Machinery. https://doi.org/10.1145/3689050.3705994

Recording Industry Association of America. (2025). 2024 year-end music industry revenue report. https://www.riaa.com/reports/2024-year-end-music-industry-revenue-report-riaa/

Redström, J. (2006). Towards user design? On the shift from object to user as the subject of design. Design Studies, 27(2), 123–139. https://doi.org/10.1016/j.destud.2005.06.001

Rochberg-Halton, E. (1984). Object relations, role models, and cultivation of the self. Environment and Behavior, 16(3), 335–368. https://doi.org/10.1177/0013916584163003

8.4 Refrences

Schedl, M., Zamani, H., Chen, C.-W., Deldjoo, Y., & Elahi, M. (2018). Current challenges and visions in music recommender systems research. International Journal of Multimedia Information Retrieval, 7(2), 95–116. https://doi.org/10.1007/s13735-018-0154-2

Sloboda, J. A., O'Neill, S. A., & Ivaldi, A. (2001). Functions of music in everyday life: An exploratory study using the experience sampling method. Musicae Scientiae, 5(1), 9–32. https://doi.org/10.1177/102986490100500102

Spotify. (n.d.). About Spotify. Retrieved August 14, 2026, from https://newsroom.spotify.com/company-info/

Spotify. (2023, February 22). Spotify debuts a new AI DJ, right in your pocket. https://newsroom.spotify.com/2023-02-22/spotify-debuts-a-new-ai-dj-right-in-your-pocket/

Spotify. (2026a, January 22). Prompted Playlist in beta coming to Premium listeners in more markets. https://newsroom.spotify.com/2026-01-22/prompted-playlists-expansion/

Spotify. (2026b, February 10). Spotify reports fourth quarter 2025 earnings. https://newsroom.spotify.com/2026-02-10/spotify-q4-2025-earnings/

Spotify. (2026c, July 10). Spotify's discovery-driven playlists keep your weekly music rotation feeling fresh. https://newsroom.spotify.com/2026-07-10/discovery-playlists-release-radar-control-updates/

Spotify. (2026d, January 28). 3 easy ways to discover music that fits your moment on Spotify. https://newsroom.spotify.com/2026-01-28/music-discovery-features/

Sterne, J. (2003). The audible past: Cultural origins of sound reproduction. Duke University Press. https://doi.org/10.1215/9780822384250

Straw, W. (1997). Sizing up record collections: Gender and connoisseurship in rock music culture. In S. Whiteley (Ed.), Sexing the groove: Popular music and gender (pp. 3–16). Routledge.

Thonhauser, G. (2021). Beyond mood and atmosphere: A conceptual history of the term Stimmung. Philosophia, 49(3), 1247–1265. https://doi.org/10.1007/s11406-020-00290-7

Tsiakas, K., & Murray-Rust, D. (2024). Unpacking human-AI interactions: From interaction primitives to a design space. ACM Transactions on Interactive Intelligent Systems, 14(3), Article 18. https://doi.org/10.1145/3664522

Wang, X., Wang, Y., Hsu, D., & Wang, Y. (2014). Exploration in interactive personalized music recommendation: A reinforcement learning approach. ACM Transactions on Multimedia Computing, Communications, and Applications, 11(1), Article 7. https://doi.org/10.1145/2623372

Whitehouse, S. (2023). "Taking a chance on a record": Lost vinyl consumption practices in the age of music streaming. Consumption Markets & Culture, 26(1), 64–80. https://doi.org/10.1080/10253866.2022.2134124

Zhang, Y. C., Ó Séaghdha, D., Quercia, D., & Jambor, T. (2012). Auralist: Introducing serendipity into music recommendation. In Proceedings of the Fifth ACM International Conference on Web Search and Data Mining (pp. 13–22). Association for Computing Machinery. https://doi.org/10.1145/2124295.2124300

Zhao, S., Li, Y., Yao, X., Nie, W., Xu, P., Yang, J., & Keutzer, K. (2020). Emotion-based end-to-end matching between image and music in valence-arousal space. In Proceedings of the 28th ACM International Conference on Multimedia (pp. 2945–2954). Association for Computing Machinery. https://doi.org/10.1145/3394171.3413776

Zimmerman, J., & Forlizzi, J. (2008). The role of design artifacts in design theory construction. Artifact, 2(1), 41–45. https://doi.org/10.1080/17493460802276893

Zimmerman, J., Forlizzi, J., & Evenson, S. (2007). Research through design as a method for interaction design research in HCI. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (pp. 493–502). Association for Computing Machinery. https://doi.org/10.1145/1240624.1240704

8.5 Image Refrences

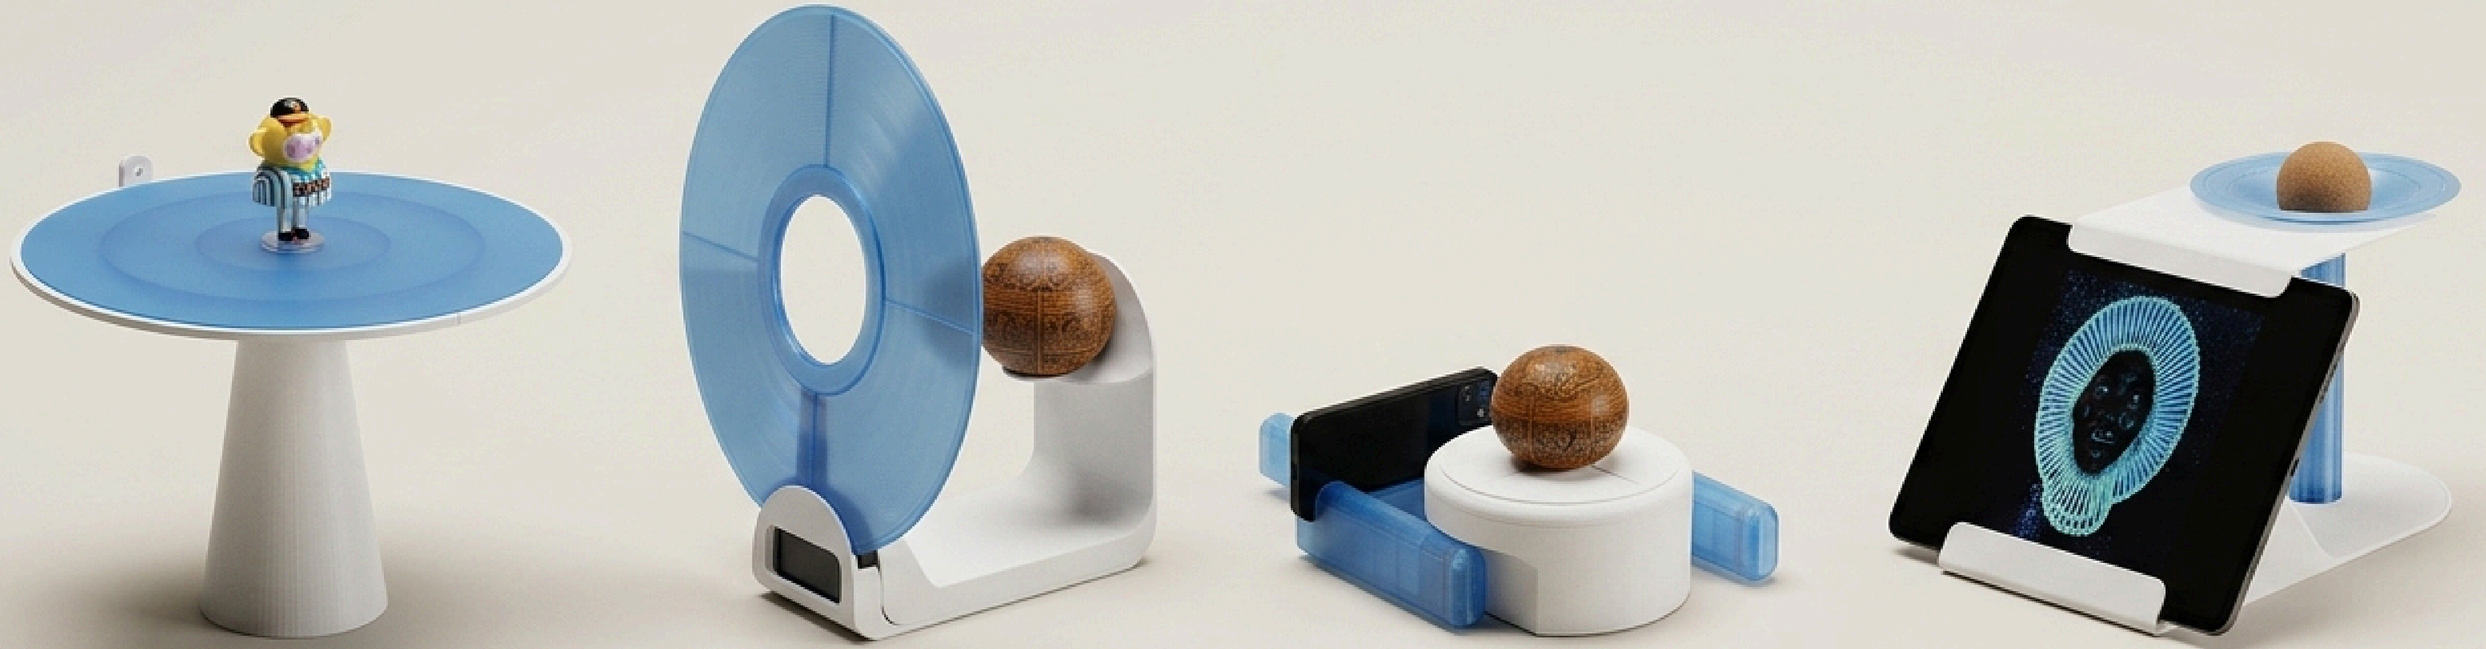
Geiger, A. (2015). SUR-FAKE [Photograph series]. https://web.archive.org/web/20190413215317/https://antoinegeiger.com/filter/photo/SUR-FAKE

Miclet, B. (2023, December 22). Rap, soul, chants lettons : notre sélection d'albums de Noël [Rap, soul, Latvian songs: Our selection of Christmas albums] [Photograph]. Harper's Bazaar France. https://www.harpersbazaar.fr/culture/rap-soul-chants-lettons-notre-selection-dalbums-de-noel_1688

Pierce, D. (2023, March 8). Spotify's new design is part TikTok, part Instagram, and part YouTube [Screenshots]. The Verge. https://www.theverge.com/2023/3/8/23630821/spotify-design-home-music-podcasts-audiobooks-app

Record Mania. (n.d.). Location, opening hours [Photograph]. Retrieved August 20, 2026, from http://www.recordmania.nl/storelocation-opening-hours/

Thompson, F. M. (2015, October 1). Amsterdam travel: Concerto Music Store - Amsterdam's best music shop [Photograph]. As the Bird Flies. https://www.asthebirdfliesblog.com/posts/concerto-record-store-amsterdam



Appendix A - Refrences for Spotify Features Graph

The sources below are the ones behind Figure 2.1 and the number against each entry matches the number that event carries on the figure. The timeline was compiled with the assistance of Claude. Its purpose is illustrative rather than evidential; it is there to show the pace and direction of a shift, not to carry an argument. Its sources are therefore listed here rather than in the main reference list, which holds only the works cited in the body text.

1. Spotify launches (October 2008)

Spotify. (2018a, October 18). Celebrating a decade of discovery on Spotify. Spotify Newsroom. https://newsroom.spotify.com/2018-10-18/celebrating-a-decade-of-discovery-on-spotify/

2. Spotify Radio (16 December 2011)

Empson, R. (2011, December 16). The Echo Nest to power new Spotify Radio (which begins rolling out today). TechCrunch. https://techcrunch.com/2011/12/16/the-echo-nest-to-power-new-spotify-radio-which-begins-rolling-out-today/

O'Brien, T. (2011, December 16). Echo Nest is the man behind the Spotify Radio curtain. Engadget. https://www.engadget.com/2011-12-16-echo-nest-is-the-man-behind-the-spotify-radio-curtain.html

3. Annoy open-sourced (2013)

Bernhardsson, E. (2023). Annoy (Version 1.17.3) [Computer software]. GitHub. https://github.com/spotify/annoy

Sobot, P. (2023, October 25). Introducing Voyager: Spotify's new nearest-neighbor search library. Spotify Engineering. https://engineering.atspotify.com/2023/10/introducing-voyager-spotifys-new-nearest-neighbor-search-library

Note. The repository states no public release date. The year rests on Spotify's own retrospective wording in the Voyager post, "back in 2013", with the repository's creation in April 2013 as corroboration.

4. Echo Nest acquisition (6 March 2014)

Spotify. (2014, March 6). Spotify acquires The Echo Nest [Press release]. Cision News. https://news.cision.com/spotify/r/spotify-acquires-the-echo-nest,c9547435

Etherington, D. (2014, March 6). Spotify acquires The Echo Nest, gaining control of the music DNA company that powers its rivals. TechCrunch. https://techcrunch.com/2014/03/06/spotify-acquires-the-echo-nest/

5. Seed Scientific acquisition (24 June 2015)

Constine, J. (2015a, June 24). Spotify buys Seats' analytics provider Seed Scientific. TechCrunch. https://techcrunch.com/2015/06/24/pulling-the-data-rug-out-from-under-apple/

Spotify acquires data provider to the UN - and Beats Music. (2015, June 25). Music Business Worldwide. https://www.musicbusinessworldwide.com/spotify-acquires-data-provider-to-the-un-and-beats-music/

6. Discover Weekly (20 July 2015)

Spotify. (2015, November 18). What made Discover Weekly one of our most successful feature launches to date? Spotify Engineering. https://engineering.atspotify.com/2015/11/what-made-discover-weekly-one-of-our-most-successful-feature-launches-to-date

Spotify. (2020, July 9). Spotify users have spent over 2.3 billion hours streaming Discover Weekly playlists since 2015. Spotify Newsroom. https://newsroom.spotify.com/2020-07-09/spotify-users-have-spent-over-2-3-billion-hours-streaming-discover-weekly-playlists-since-2015/

Constine, J. (2015b, July 29). Spotify's latest trick is a personalized news feed-style "Discover Weekly" playlist. TechCrunch. https://techcrunch.com/2015/07/29/why-browse-when-you-can-play/

Kokalitcheva, K. (2015, July 29). Spotify's new feature will help you fight the Monday blues. Time. https://time.com/3985678/spotify-playlists/

7. Fresh Finds (early March 2016)

Spotify. (2021d, May 28). Spotify's new Fresh Finds program helps independent artists build their careers. Spotify Newsroom. https://newsroom.spotify.com/2021-05-28/spotifys-new-fresh-finds-program-helps-independent-artists-build-their-careers/

Luckerson, V. (2016, March 7). Spotify's new tool can help find your next favorite song. Time. https://time.com/4249597/spotifys-new-tool-can-help-find-your-next-favorite-song/

Wilson, S. (2016, March 2). Spotify launches playlists designed to predict the next big thing. Fact Magazine. https://www.factmag.com/2016/03/02/spotify-fresh-finds-playlists-launched/

8. Release Radar (4 August 2016)

Spotify. (2016, August 4). Say hello to Release Radar. Spotify for Artists. https://artists.spotify.com/blog/say-hello-to-release-radar

Spotify. (2021e, September 8). Spotify's Release Radar personalized playlist celebrates five years and 18 billion streams. Spotify Newsroom. https://newsroom.spotify.com/2021-09-08/spotifys-release-radar-personalized-playlist-celebrates-five-years-and-18-billion-streams/

Constine, J. (2016, August 8). Spotify follows Discover Weekly with personalized new releases playlist. TechCrunch. https://techcrunch.com/2016/08/08/spotify-release-radar/

9. Daily Mix (27 September 2016)

Lunden, I. (2016, September 27). With Daily Mix, Spotify taps algorithms to make infinite playlists from your favorite tunes. TechCrunch. https://techcrunch.com/2016/09/27/spotify-daily-mix/

Wilkie, E. (2016, September 27). Spotify launches new Daily Mix feature. Music Week. https://www.musicweek.com/digital/read/spotify-launches-new-daily-mix-feature/068671

10. Wrapped (December 2016)

Spotify. (2024c, December 4). We're commemorating a decade of Spotify Wrapped with our best and boldest Wrapped yet. Spotify Newsroom. https://newsroom.spotify.com/2024-12-04/10-years-spotify-wrapped/

Spotify. (n.d.-d). Spotify: A visual history. Spotify Newsroom. Retrieved July 28, 2026, from https://newsroom.spotify.com/spotify-timeline/

11. Sonalytic acquisition (7 March 2017)

Perez, S. (2017a, March 7). Spotify acquires audio detection startup Sonalytic. TechCrunch. https://techcrunch.com/2017/03/07/spotify-acquires-audio-detection-startup-sonalytic/

Evanson, D. (2017, March 15). Spotify buys Imperial audio recognition start-up. Imperial College London. https://www.imperial.ac.uk/news/178176/spotify-buys-imperial-audio-recognition-start-up/

12. Mediachain acquisition (28 April 2017)

Perez, S. (2017b, April 28). Spotify acquires blockchain startup Mediachain to solve music's attribution problem. TechCrunch. https://techcrunch.com/2017/04/28/spotify-acquires-blockchain-startup-mediachain-to-solve-musics-attribution-problem/

Higgins, S. (2017, April 28). Spotify acquires blockchain startup Mediachain. CoinDesk. https://www.coindesk.com/markets/2017/04/28/spotify-acquires-blockchain-startup-mediachain

13. Niland acquisition (17 May 2017)

Keane, J. (2017, May 17). Spotify acquires Paris-based AI search startup Niland.io. Tech.eu. https://tech.eu/2017/05/17/spotify-niland-io/

Russell, J. (2017, May 18). Spotify buys AI startup Niland to develop its music personalisation and recommendations. TechCrunch. https://techcrunch.com/2017/05/18/spotify-buys-ai-startup-niland/

McIntyre, H. (2017, May 18). Spotify has acquired machine-learning startup Niland. Forbes. https://www.forbes.com/sites/hughmcintyre/2017/05/18/spotify-has-acquired-machine-learning-startup-niland/

14. Bart ranking research (October 2018)

McInerney, J., Lacker, B., Hansen, S., Higley, K., Bouchard, H., Gruson, A., & Mehrotra, R. (2018). Explore, exploit, and explain: Personalizing explainable recommendations with bandits. In Proceedings of the 12th ACM Conference on Recommender Systems (RecSys '18) (pp. 31–39). Association for Computing Machinery. https://doi.org/10.1145/3240323.3240354

Spotify. (2018b, October 2). Explore, exploit, explain: Personalizing explainable recommendations with bandits. Spotify Research. https://research.atspotify.com/publications/explore-exploit-explain-personalizing-explainable-recommendations-with-bandits

15. Discovery Mode (2 November 2020)

Spotify. (n.d.-b). Discovery Mode. Spotify for Artists. Retrieved August 5, 2026, from https://artists.spotify.com/discovery-mode

Dredge, S. (2020, November 2). Spotify's new artist tool could boost streams (but with a lower royalty rate). Music Ally. https://musically.com/2020/11/02/spotify-artist-tool-boost-streams-discounted-royalty-rate/

Murray, C. (2025, November 5). Spotify hit with class action lawsuit alleging Discovery Mode is a "pay-for-play scheme." Forbes. https://www.forbes.com/sites/conormurray/2025/11/05/spotify-hit-with-class-action-lawsuit-alleging-discovery-mode-is-a-pay-for-play-scheme/

Note. The date comes from the Music Ally report. The Spotify for Artists product page carries the mechanism but is undated.

16. Podz acquisition (17 June 2021)

Spotify. (2021c, June 17). Spotify accelerates podcast discovery with acquisition of Podz. Spotify Newsroom. https://newsroom.spotify.com/2021-06-17/spotify-accelerates-podcast-discovery-with-acquisition-of-podz/

Silberling, A. (2021, June 17). Spotify acquires Podz, a podcast discovery platform. TechCrunch. https://techcrunch.com/2021/06/17/spotify-acquires-podz-a-podcast-discovery-app/

Spangler, T. (2021, June 18). Spotify buys Podz, a podcast discovery tech startup whose backers included Katie Couric and Paris Hilton. Variety. https://variety.com/2021/digital/news/spotify-podz-acquisition-1235000150/

17. Blend (31 August 2021)

Spotify. (2021b, August 31). How Spotify's newest personalized experience, Blend, creates a playlist for you and your bestie. Spotify Newsroom. https://newsroom.spotify.com/2021-08-31/how-spotifys-newest-personalized-experience-blend-creates-a-playlist-for-you-and-your-bestie/

Perez, S. (2021, August 31). Spotify officially launches Blend, allowing friends to match their musical tastes and make playlists together. TechCrunch. https://techcrunch.com/2021/08/31/spotify-officially-launches-blend-allowing-friends-to-match-their-musical-tastes-and-make-playlists-together/

18. Enhance (9 September 2021)

Spotify. (2021a, September 9). Get perfect song recommendations in the playlists you create with Enhance. Spotify Newsroom. https://newsroom.spotify.com/2021-09-09/get-perfect-song-recommendations-in-the-playlists-you-create-with-enhance/

Adorno, J. (2021, September 9). Spotify wants to improve your playlists with new Enhance feature. 9to5Mac. https://9to5mac.com/2021/09/09/spotify-wants-to-improve-your-playlists-with-new-enhance-feature/

19. Basic Pitch (June 2022)

Bittner, R. M., Bosch, J. J., Rubinstein, D., Meseguer-Brocal, G., & Ewert, S. (2022). A lightweight instrument-agnostic model for polyphonic note transcription and multipitch estimation [Preprint]. arXiv. https://arxiv.org/abs/2203.09893

Bittner, R. (2022, June 1). Meet Basic Pitch: Spotify's open source audio-to-MIDI converter. Spotify Engineering. https://engineering.atspotify.com/2022/06/meet-basic-pitch

Spotify. (2022a). Basic Pitch [Computer software]. GitHub. https://github.com/spotify/basic-pitch

Spotify. (2022b, September 1). Rachel Bittner on Basic Pitch: An open source tool for musicians. Spotify Newsroom. https://newsroom.spotify.com/2022-09-01/rachel-bittner-on-basic-pitch-an-open-source-tool-for-musicians/

20. Sonantio acquisition (13 June 2022)

Spotify. (2022d, June 13). Spotify to aquire Sonantio, an AI voice platform. Spotify Newsroom. https://newsroom.spotify.com/2022-06-13/spotify-to-acquire-sonantio-an-ai-voice-platform/

Spotify. (2023a, March 8). Behind the scenes of Spotify's new AI DJ. Spotify Newsroom. https://newsroom.spotify.com/2023-03-08/spotify-new-personalized-ai-dj-how-it-works/
Also cited at node 22.

Note. Announced 13 June 2022 and completed 11 July 2022.

21. Kinzen acquisition (5 October 2022)

Spotify. (2022e, October 5). Spotify continues to ramp up platform safety efforts with acquisition of Kinzen. Spotify Newsroom. https://newsroom.spotify.com/2022-10-05/spotify-continues-to-ramp-up-platform-safety-efforts-with-acquisition-of-kinzen/

Perez, S. (2022, October 5). Spotify acquires content moderation tech company Kinzen to address platform safety issues. TechCrunch. https://techcrunch.com/2022/10/05/spotify-acquires-content-moderation-tech-company-kinzen-to-address-platform-safety-issues/

22. AI DJ (22 February 2023)

Spotify. (2023d, February 22). Spotify debuts a new AI DJ, right in your pocket. Spotify Newsroom. https://newsroom.spotify.com/2023-02-22/spotify-debuts-a-new-ai-dj-right-in-your-pocket/

Spotify. (2023a, March 8). Behind the scenes of Spotify's new AI DJ. Spotify Newsroom. https://newsroom.spotify.com/2023-03-08/spotify-new-personalized-ai-dj-how-it-works/
Also cited at node 20.

Spotify. (2023e, August 8). Spotify expands DJ to now be available in 60 markets around the world. Here's how to find it. Spotify Newsroom. https://newsroom.spotify.com/2023-08-08/ai-dj-expanded-new-markets-how-to-use-feature/

Perez, S. (2023a, February 22). Spotify launches 'DJ,' a new feature offering personalized music with AI-powered commentary. TechCrunch. https://techcrunch.com/2023/02/22/spotify-launches-dj-a-new-personalized-music-lineup-that-includes-ai-powered-commentary/

23. Smart Shuffle (8 March 2023)

Spotify. (2023a, March 8). Smart Shuffle breathes new life into your Spotify playlists. Spotify Newsroom. https://newsroom.spotify.com/2023-03-08/smart-shuffle-new-life-spotify-playlists/

Perez, S., & Malik, A. (2023, March 8). Spotify revamps its app with TikTok-style discovery

Appendix A - Refrences for Spotify Features Graph

feeds, Smart Shuffle for playlists and more. TechCrunch. https://techcrunch.com/2023/03/08/spotify-revamps-its-app-with-tiktok-style-discovery-feeds-smart-shuffle-for-playlists-and-more/

24. daylist (12 September 2023)

Spotify. (2023b, September 12). Get fresh music sunup to sundown with daylist, your ever-changing Spotify playlist. Spotify Newsroom. https://newsroom.spotify.com/2023-09-12/ever-changing-playlist-daylist-music-for-all-day/

Perez, S. (2023b, September 12). Spotify's latest personalized playlist, daylist, evolves with you throughout the day. TechCrunch. https://techcrunch.com/2023/09/12/spotifys-latest-personalized-playlist-daylist-evolves-with-you-throughout-the-day/

25. Voice Translation (25 September 2023)

Spotify. (2023f, September 25). Spotify's AI voice translation pilot means your favorite podcasters might be heard in your native language. Spotify Newsroom. https://newsroom.spotify.com/2023-09-25/ai-voice-translation-pilot-lex-fridman-dax-shepard-steven-bartlett/

Spangler, T. (2023, September 25). Spotify is testing AI-powered podcast language translation, which mimics the podcaster's own voice. Variety. https://variety.com/2023/digital/news/spotify-ai-podcast-language-translation-1235732548/

Capoot, A. (2023, September 25). Spotify will use AI to replicate podcasters' voices and translate them to other languages. CNBC. https://www.cnbc.com/2023/09/25/spotify-will-use-ai-to-replicate-and-translate-podcasters-voices.html

28. AI Playlist beta (7 April 2024)

Spotify. (2024b, April 7). Spotify Premium users can now turn any idea into a personalized playlist with AI Playlist in beta. Spotify Newsroom. https://newsroom.spotify.com/2024-04-07/spotify-premium-users-can-now-turn-any-idea-into-a-personalized-playlist-with-ai-playlist-in-beta/

Spotify. (2024a, September 24). AI Playlist is rolling out in beta in the US, Canada, Ireland, and New Zealand: Get started with these pro tips. Spotify Newsroom. https://newsroom.spotify.com/2024-09-24/ai-playlist-expanding-usa-canada-ireland-new-zealand/

Spotify. (2025b, April 24). Spotify expands AI Playlist in beta to Premium listeners in 40+ new markets. Spotify Newsroom. https://newsroom.spotify.com/2025-04-24/spotify-expands-ai-playlist-in-beta-to-premium-listeners-in-40-new-markets/

27. AI Wrapped podcast (4 December 2024)

Spotify. (2024d, December 4). Your Spotify Wrapped AI podcast is here to help you reconnect with the music that defined your year. Spotify Newsroom. https://newsroom.spotify.com/2024-12-04/your-spotify-wrapped-ai-podcast-is-here-to-help-you-reconnect-with-the-music-that-defined-your-year/

Google. (2024, December 4). Listen to your first-ever 2024 Spotify Wrapped AI podcast, built with NotebookLM. Google Blog. https://blog.google/innovation-and-ai/models-and-research/google-labs/notebooklm-spotify-wrapped/

Perez, S. (2024, December 4). Spotify Wrapped 2024 adds an AI podcast powered by Google's NotebookLM. TechCrunch. https://techcrunch.com/2024/12/04/spotify-wrapped-2024-adds-an-ai-podcast-powered-by-googles-notebooklm

28. Generative-AI ads (2 April 2025)

Spotify. (2025d, April 2). Spotify tunes up platform to make it easy for advertisers to buy, create, and measure. Spotify Newsroom. https://newsroom.spotify.com/2025-04-02/spotify-tunes-up-platform-to-make-it-easy-for-advertisers-to-buy-create-and-measure/

Spotify. (n.d.-c). Gen AI Audio Ad Creator. Spotify Ads. Retrieved July 28, 2025, from https://ads.spotify.com/en-US/creative-lab/gen-ai/

Perez, S. (2025b, April 3). Spotify debuts Gen AI ads, programmatic ad buying. TechCrunch. https://techcrunch.com/2025/04/03/spotify-debuts-gen-ai-ads-programmatic-ad-buying

29. AI DJ requests (13 May 2025)

Spotify. (2025e, May 13). Spotify's DJ now takes requests, enhancing real-time music discovery. Spotify Newsroom. https://newsroom.spotify.com/2025-05-13/dj-voice-requests/

Spotify. (2025c, October 15). DJ levels up with Spanish-language requests, text requests, and personalized prompts. Spotify Newsroom. https://newsroom.spotify.com/2025-10-15/dj-spanish-text-requests-update/

Mehta, I. (2025, May 13). Spotify's AI DJ now lets you use voice commands to personalize your tunes. TechCrunch. https://techcrunch.com/2025/05/13/spotifys-ai-dj-now-lets-you-use-voice-commands-to-personalize-your-tunes/

30. AI-music policy (25 September 2025)

Spotify. (2025a, September 25). Spotify strengthens AI protections for artists, songwriters, and producers. Spotify Newsroom. https://newsroom.spotify.com/2025-09-25/spotify-strengthens-ai-protections/

Perez, S. (2025c, September 25). Spotify to label AI music, filter spam and more in AI policy change. TechCrunch. https://techcrunch.com/2025/09/25/spotify-updates-ai-policy-to-label-tracks-cut-down-on-spam/

Stassen, M. (2025, September 25). Spotify has deleted 75m+ "spammy tracks" – as it unveils new AI music policies. Music Business Worldwide. https://www.musicbusinessworldwide.com/spotify-has-deleted-75m-spammy-tracks-as-it-unveils-new-ai-music-policies/

31. ChatGPT connector (8 October 2025)

Spotify. (2025g, October 8). Your prompts, Spotify's personalized picks: Introducing Spotify in ChatGPT. Spotify Newsroom. https://newsroom.spotify.com/2025-10-08/spotify-personalized-prompts-chatgpt/

Forristal, L. (2025, October 10). You can now connect your Spotify account to ChatGPT: Here's how to do it. TechCrunch. https://techcrunch.com/2025/10/10/you-can-now-connect-your-spotify-account-to-chatgpt-heres-how-to-do-it/

32. WhoSampled acquisition (10 November 2025)

Spotify. (2025f, November 10). Spotlighting the people, connections, and stories behind your music. Spotify for Artists. https://artists.spotify.com/en/blog/spotlighting-the-people-connections-and-stories-behind-your-music

WhoSampled. (2025, November 10). WhoSampled is joining Spotify! https://www.whosampled.com/news/2025/11/10/whosampled-is-joining-spotify/

Perez, S. (2025a, November 10). Spotify acquires music database WhoSampled. TechCrunch. https://techcrunch.com/2025/11/10/spotify-acquires-music-database-whosampled/

Tenser, D. (2025, November 10). Spotify acquires song sample and remix database WhoSampled. Music Business Worldwide. https://www.musicbusinessworldwide.com/spotify-acquires-song-sample-and-remix-database-whosampled/

Note. The date is taken from Spotify's own post. The Music Business Worldwide report is internally inconsistent on the day.

33. Prompted Playlist (22 January 2026)

Spotify. (2026h, January 22). Prompted Playlist in beta coming to Premium listeners in more markets. Spotify Newsroom. https://newsroom.spotify.com/2026-01-22/prompted-playlists-expansion/

Perez, S. (2026a, January 22). Spotify brings AI-powered Prompted Playlists to the US and Canada. TechCrunch. https://techcrunch.com/2026/01/22/spotify-brings-ai-powered-prompted-playlists-to-the-u-s-and-canada/

34. About the Song beta (6 February 2026)

Spotify. (2026d, February 6). Introducing About the Song, a new way to explore the stories behind the music. Spotify Newsroom. https://newsroom.spotify.com/2026-02-06/about-the-song-beta/

Spotify. (2025a, April 28). 3 new ways to dive deeper into the music you love. Spotify Newsroom. https://newsroom.spotify.com/2025-04-29/deeper-music-features-about-the-song-dj-songdna/

Malik, A. (2026, February 6). Spotify's new feature lets you explore the story behind the song you're listening to. TechCrunch. https://techcrunch.com/2026/02/06/spotifys-new-feature-lets-you-explore-the-story-behind-the-song-youre-listening-to/

Davies, R. (2026, July 21). Pop singer Lorde slammed Spotify's AI 'About the Song' feature, claiming it's full of nothing but false details, and now subscribers are starting to catch on. TechRadar. https://www.techradar.com/audio/spotify/pop-singer-lorde-slammed-spotifys-ai-about-the-song-feature-claiming-its-full-of-nothing-but-false-details-and-now-subscribers-are-starting-to-catch-on

35. Taste Profile beta (13 March 2026)

Spotify. (2026b, March 13). A new era of personalization: Shape your Taste Profile on Spotify. Spotify Newsroom. https://newsroom.spotify.com/2026-03-13/taste-profile-beta-announcement/

36. SongDNA beta (24 March 2026)

Spotify. (2026g, March 24). Meet SongDNA, a new beta feature that reveals the creative connections behind your favorite tracks. Spotify Newsroom. https://newsroom.spotify.com/2026-03-24/songdna-announcement-beta/

Perez, S. (2026e, March 24). Spotify's new SongDNA feature maps how your favorite songs are connected. TechCrunch. https://techcrunch.com/2026/03/24/spotifys-songdna-feature-lets-you-explore-the-connections-behind-your-favorite-songs/

Dalugdug, M. (2026, March 24). Spotify launches SongDNA in beta for Premium subscribers, tracing how songs connect through shared producers, samples and covers. Music Business Worldwide. https://www.musicbusinessworldwide.com/spotify-launches-songdna-in-beta-for-premium-subscribers-tracing-how-songs-connect-through-shared-producers-samples-and-covers/

37. AI Credits beta (17 April 2026)

Spotify. (n.d.-a). AI credits on Spotify. Spotify for Artists Support. Retrieved August 5, 2026, from https://support.spotify.com/us/artists/article/ai-credits/

Dredge, S. (2026, April 17). Spotify rolls out its own AI-transparency feature in beta. Music Ally. https://musically.com/2026/04/17/spotify-roll-out-its-own-ai-transparency-feature-in-beta/

Note. Spotify's support article is undated. The date comes from the Music Ally report.

38. Verified by Spotify (30 April 2026)

Spotify. (2026e, April 30). Introducing Verified by Spotify, a signal of authenticity and trust for the artists behind the music. Spotify Newsroom. https://newsroom.spotify.com/2026-04-30/verified-by-spotify-badge-artist-details/

39. Podcast Q&A (21 May 2026)

Spotify. (2026j, May 21). Spotify's 2026 Investor Day recap: Raising ambition for the next era of media. Spotify Newsroom. https://newsroom.spotify.com/2026-05-21/investor-day-recap/
Also cited at node 41.

Mehta, I. (2026, May 21). Spotify adds AI-powered Q&A and briefing generation features to podcasts. TechCrunch. https://techcrunch.com/2026/05/21/spotify-adds-ai-powered-qa-and-briefing-generation-features-to-podcasts/
Also cited at node 40.

40. Personal podcasts and briefings (21 May 2026)

Spotify. (2026o, May 21). Create, control, and personalize your listening across every moment. Spotify Newsroom. https://newsroom.spotify.com/2026-05-21/studio-by-spotify-labs-launch/
Also cited at node 41.

Mehta, I. (2026, May 21). Spotify adds AI-powered Q&A and briefing generation features to podcasts. TechCrunch. https://techcrunch.com/2026/05/21/spotify-adds-ai-powered-qa-and-briefing-generation-features-to-podcasts/
Also cited at node 39.

41. Studio by Spotify Labs (21 May 2026)

Spotify. (2026o, May 21). Create, control, and personalize your listening across every moment. Spotify Newsroom. https://newsroom.spotify.com/2026-05-21/studio-by-spotify-labs-launch/
Also cited at node 40.

Spotify. (2026j, May 21). Spotify's 2026 Investor Day recap: Raising ambition for the next era of media. Spotify Newsroom. https://newsroom.spotify.com/2026-05-21/investor-day-recap/
Also cited at node 39.

42. AI covers and remixes, licensed with UMG (21 May 2026)

Spotify. (2026i, May 21). Spotify and Universal Music Group announce landmark licensing agreements for fan-made covers and remixes. Spotify Newsroom. https://newsroom.spotify.com/2026-05-21/universal-music-group-spotify-licensing-agreements-fan-made-covers-remixes/

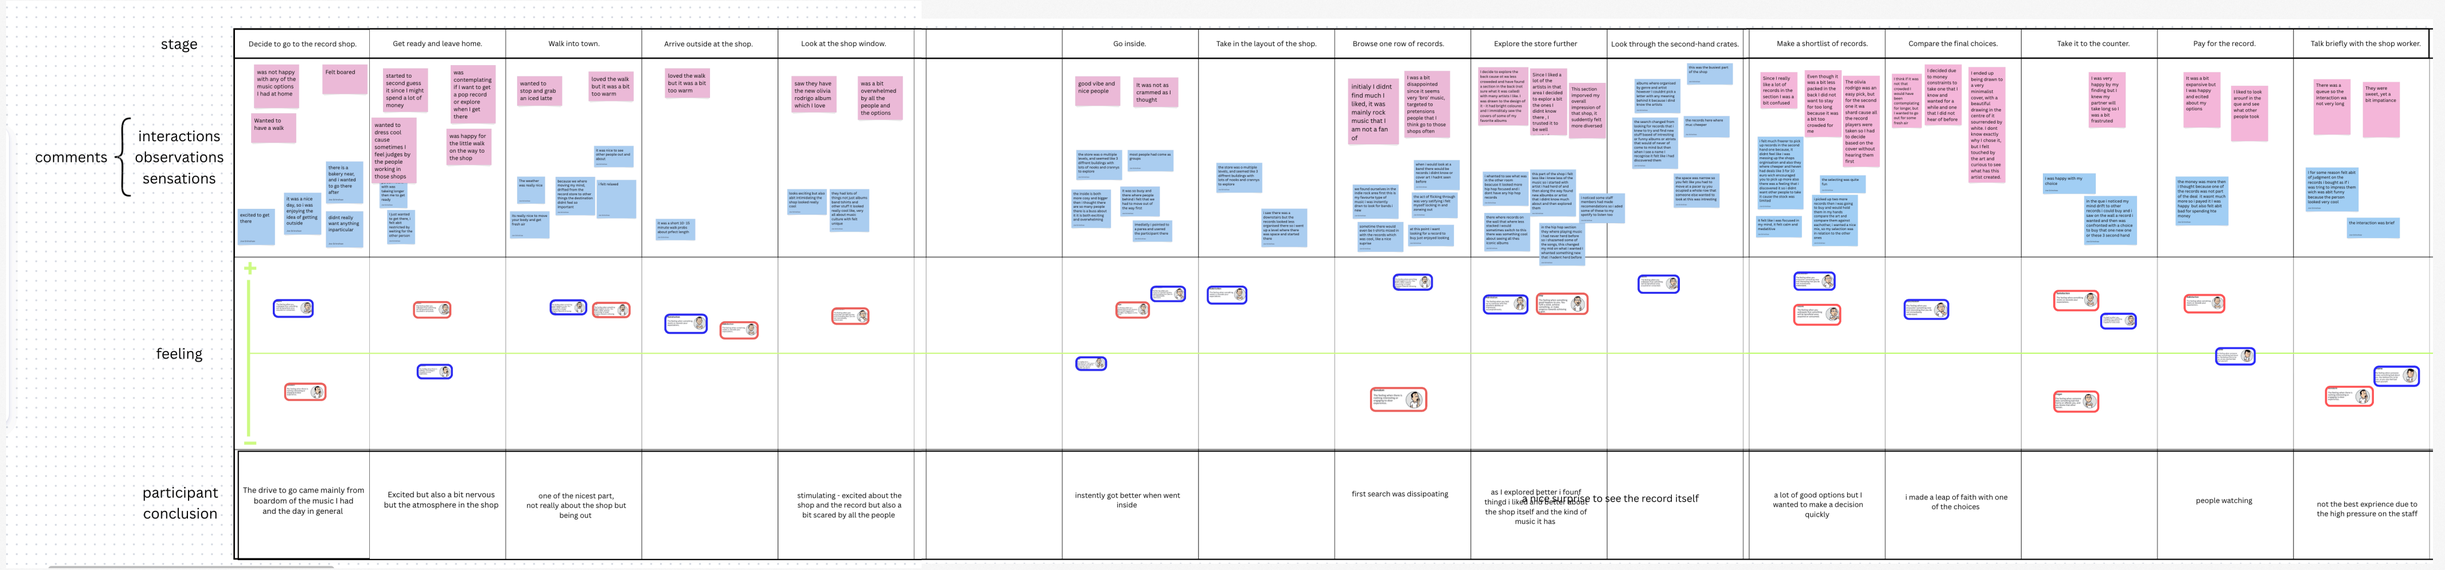
Schmelzer, R. (2026, May 28). Spotify's AI Remix looks to turn fans into revenue. Forbes. https://www.forbes.com/sites/ronschmelzer/2026/05/28/spotifys-ai-remix-looks-to-turn-fans-into-revenue/

43. Talk to Spotify beta (14 July 2026)

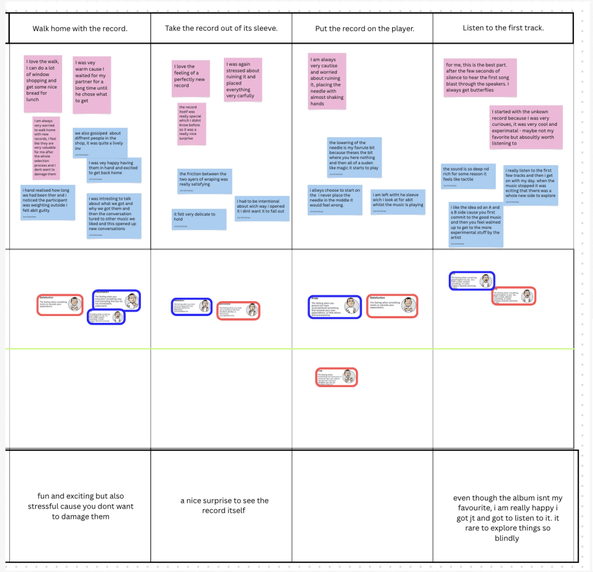
Spotify. (2026f, July 14). Just say the word: A more personal way to ask, discover, and listen. Spotify Newsroom. https://newsroom.spotify.com/2026-07-14/talk-to-spotify-announcement-beta/

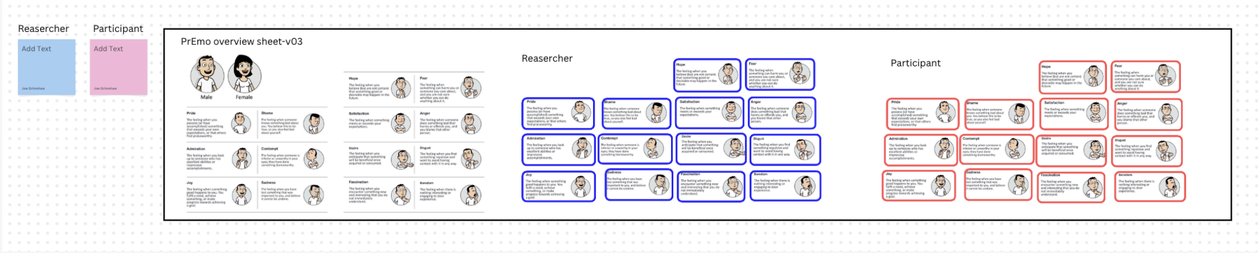
Perez, S. (2026b, July 14). Spotify expands its AI push with a ChatGPT-like music assistant. TechCrunch. https://techcrunch.com/2026/07/14/spotify-expands-its-ai-push-with-a-chatgpt-like-music-assistant/

Appendix B - Jorney Map from Participant Study (section 1)



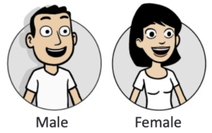
Appendix B - Jorney Map from Participant Study (section 2)





Appendix C - PrEmo Tool used from Participant Study

PrEmo overview sheet-v03



Pride
The feeling when you possess (or have accomplished) something that exceeds your own expectations, or that others find praiseworthy.



Shame
The feeling when someone knows something bad about you. You believe this to be true, so you also feel bad about yourself.



Admiration
The feeling when you look up to someone who has excellent abilities or impressive accomplishments.



Contempt
The feeling when someone is inferior or unworthy in your eyes; they have done something blameworthy.



Joy
The feeling when something good happens to you. You fulfil a need, achieve something, or make progress towards achieving a goal.



Sadness
The feeling when you have lost something that was important to you, and believe it cannot be undone.



Hope
The feeling when you believe (but are not certain) that something good or desirable may happen in the future.



Fear
The feeling when something can harm you or someone you care about, and you are not sure whether you can do anything about it.



Satisfaction
The feeling when something meets or exceeds your expectations.



Anger
The feeling when someone does something bad that harms or offends you, and you blame that other person.



Desire
The feeling when you anticipate that something will be beneficial once acquired or consumed.



Disgust
The feeling when you find something repulsive and want to avoid having contact with it in any way.



Fascination
The feeling when you encounter something new and interesting that you do not immediately understand.



Boredom
The feeling when there is nothing interesting or engaging to door expereince.



Appendix D - Test 1 Results

Test 1 -- Results

Object	Day 1	Day 2	Day 3	Mean	Spread
1	2.4	1.8	2.1	2.1	0.6
2	3.2	2.5	3.0	2.9	0.7
3	4.2	3.5	4.0	3.9	0.7
4	3.8	3.0	3.7	3.5	0.8
5	4.5	3.8	4.0	4.1	0.7
6	1.5	1.2	1.2	1.3	0.3
7	3.5	2.8	3.3	3.2	0.7
8	3.0	3.8	3.4	3.4	0.8
9	2.6	2.0	2.3	2.3	0.6
10	3.6	4.3	3.8	3.9	0.7
11	4.1	3.4	3.9	3.8	0.7
12	3.4	4.0	3.7	3.7	0.6
13	3.5	2.7	3.1	3.1	0.8
14	3.7	3.0	3.5	3.4	0.7
15	1.2	1.0	1.1	1.1	0.2
16	2.2	1.6	1.9	1.9	0.6
17	2.3	1.7	2.0	2.0	0.6
18	1.8	2.4	1.8	2.0	0.6
Set mean	**3.0**	**2.7**	**2.9**	**2.9**	**0.6**

Appendix E - Jorney Map - Participant Consent Forms

PARTICIPANT INFORMATION SHEET

Template 1 — Informed Consent Materials

Study Title	Exploring Tangible Interfaces for AI-Mediated Music Discovery
Researcher	Joe Grimshaw | jgrimshaw@tudelft.nl
Supervisors	Dr. Dave Murray-Rust | D.S.Murray-Rust@tudelft.nl Prof. Gerd Kortuem | G.W.Kortuem@tudelft.nl
Date	12/04/2026
Version	v1.0

1 Introduction

You are being invited to participate in a research study titled **Exploring Tangible Interfaces for AI-Mediated Music Discovery**. This study is being done by **Joe Grimshaw**, a Master's student in the Department of Design for Interaction, Faculty of Industrial Design Engineering (IDE), at **Delft University of Technology (TU Delft)**. The project is supervised by **Dr. Dave Murray-Rust** (TU Delft) & **Prof. Gerd Kortuem.** This research has no external funding and is conducted solely as part of a Master's graduation project at TU Delft.

The purpose of this research study is to investigate how tangible (physical) interfaces may influence the way people explore and discover music using AI-mediated recommendation systems. Current music streaming platforms primarily rely on screen-based interaction and algorithmic recommendations. This study explores whether physical interaction with a prototype interface changes how users experience exploration and discovery, with potential implications for the design of future music interfaces and human–AI interaction more broadly. The study forms part of a Master's thesis that will be submitted to TU Delft and may be made publicly available through the TU Delft repository. Your session will take approximately **15–30 minutes** to complete. In this time, we will be asking you to interact with the prototype

2 What you will need to do

We will be asking you to interact with a series of experimental tangible interface prototypes, each combining physical controls with a digital user interface for exploring AI-generated music recommendations. Specifically, you will be asked to: perform simple music exploration tasks with each prototype; complete a short questionnaire about your experience after each prototype; and take part in a brief discussion with the researcher at the end of the session to share your reflections.

3 Data & Privacy

As with any online activity, the risk of a breach is always possible. To the best of our ability, your answers in this study will remain confidential. We will minimise any risks by doing the following:
Your name and email address will be collected to arrange the testing session. These details will be stored separately from your research data on the researcher's password-protected computer and TU Delft's secure systems, and will be deleted once your participation has finished.

During the session, the following personally identifiable research data may be collected:

- photographs or video recordings of your interaction with the prototype
- audio recordings of the discussion after the task
- questionnaire responses
- interaction logs (such as music selections and interface events)

Audio recordings will be deleted after transcription, and all transcripts will be anonymised before analysis. Any images or video used in publications will ensure that you cannot be identified (for example, by blurring or cropping faces or focusing on the prototype). Recordings will not be shared outside the research team.

Data Storage and Confidentiality

Anonymised research data (such as interaction logs and questionnaire responses) will be stored securely on the researcher's TU Delft OneDrive and kept for up to 10 years after the thesis is completed, in accordance with TU Delft data management guidelines. Only the researcher and supervisor will have access.

Publication and Open Data

The results of this study will be published as a TU Delft Master's thesis, which may appear in the TU Delft public repository and may be presented at conferences or submitted to academic journals. Any published material will use anonymised data only, and you will not be identifiable. Anonymous quotes from discussions may be used in the thesis or related publications.

Your participation in this study is entirely voluntary and **you can withdraw at any time**, without having to give a reason. You are free to omit any questions or tasks. **If you wish to have your data removed, please request this before within 3 working days of the session;** once data has been anonymised, it may no longer be possible to locate and delete your specific records. You also have the right to request access to, rectification of, or erasure of any identifiable personal data we hold about you (such as your contact details) at any time before that data is deleted. To exercise these rights, please contact the researcher directly using the details below.

4 Contact details and procedure for making complaints

Ifyou have any questions about this study before, during, or after your participation, please contact:

Corresponding Researcher	Joe Grimshaw jgrimshaw@tudelft.nl
Responsible Researcher(Supervisor)	Prof. Dave Murray-Rust D.S.Murray-Rust@tudelft.nl Department of Design for Interaction, Faculty IDE, TU Delft

AGREEMENT

By signing anddating the consent form on the following page, you confirm that you have read and understood this information sheet, that your questions have been answered to your satisfaction, andthat you agree to participate in this study on the terms described above.

Name (printed): _____Yael Kind _____

Signature: _____ Date: ___12/04/2026_____

Appendix E - Jorney Map - Participant Consent Forms

EXPLICIT CONSENT FORM

Template 2 — Informed Consent Materials

Study Title	Exploring Tangible Interfaces for AI-Mediated Music Discovery
Researcher	Joe Grimshaw | jgrimshaw@tudelft.nl
Supervisor	Dr. Dave Murray-Rust | D.S.Murray-Rust@tudelft.nl Prof. Gerd Kortuem | G.W.Kortuem@tudelft.nl
Date	12/04/2026

PLEASE TICK THE APPROPRIATE BOX FOR EACH STATEMENT	YES	NO

A: GENERAL AGREEMENT — RESEARCH GOALS, PARTICIPANT TASKS AND VOLUNTARY PARTICIPATION

1. I have read and understood the participant information sheet dated [12/04/2026], or it has been read to me. I have been able to ask questions about the study and my questions have been answered to my satisfaction.	✓	□
2. I consent voluntarily to be a participant in this study and understand that I can refuse to answer questions and withdraw at any time, without giving a reason, and can request data deletion within 3 working days of the session.	✓	□
3. I understand that taking part in the study involves: • *Interacting witha series of experimentaltangible interface prototypes* • *Performing simple music exploration tasks* • *Completing a short self-administered questionnaire after each prototype* • *Taking part in a brief post-task discussion with the researcher.* *The post-task discussion will be audio-recorded only if I give separate consent below (see point 6).* *My prototype interactions may be photographed or video-recorded only if I give separate consent below (see point 6).*	✓	□

B: POTENTIAL RISKS OF PARTICIPATING

4. I understand that taking part in the study involves the following risks, and that these will be mitigated as described: *Emotional: Music can occasionally evoke responses such as nostalgia or mild sadness. The study does not intentionally aim to elicit strong emotions. I may stop the session at any time without explanation, and if I indicate any discomfort the session will be stopped immediately.*	✓	□

B2: DATA PROTECTION AND PRIVACY

5. I understand that to schedule my session, my name and email address (personally identifiable information, PII) has been collected. *I understand that this information will be stored securely on the researcher's password-protected computer and TU Delft systems only, will not be linked to my research data, and will be deleted immediately once my participation has concluded.*	✓	□
6. I understand and consent to the following personally identifiable research data to be collected during the session: • Photographs and/or video of my interaction with the prototypes. • Audio recording of the post-task discussion • My questionnaire responses and interaction logs (music selections and interface events). • I understand there is a small risk of re-identification from audio/visual recordings, which will be mitigated as described in point 8. *The following has been taken to minimise the threat of a data breach:* • *Audio recordings will be deleted promptly after transcription.* • *Transcripts will be anonymised before analysis and storage.* • *Any images or videos used in publications will ensure I am not identifiable (e.g., faces blurred or cropped, recordings focused on the prototype).* • *No recordings will be shared beyond the research team, and anonymised data will be stored securely on the researcher's personal TU Delft OneDrive with access restricted to the researcher and supervisor only.*	✓	□

C: RESEARCH PUBLICATION AND DISSEMINATION

7. I understand that after the research study, the de-identified information I provide will be used for the following purposes: a TU Delft Master's thesis (which may be made publicly available through the TU Delft repository); and potentially academic conference presentations or journal publications. All published material will use anonymised data and I will not be identifiable in any research output.	✓	□
8. I agree that my anonymised responses, views, or other input can be quoted in research outputs such as the thesis, conference papers, or publications. I understand that my real name will not be used in any quoted material.	✓	□

D: DATA STORAGE ACCESS AND REUSE

9. I understand that anonymised research data (questionnaire responses, interaction logs) will be retained for up to 10 years following completion of the thesis, stored securely on TU Delft systems in line with TU Delft data management guidelines.	✓	□

SIGNATURES

PARTICIPANT	Name (printed): _____Yael Kind _____ Signature: _____ Date: __12/04/2026_____
RESEARCHER	Joe Grimshaw (printed): _____ Signature: _____ Date: _____ 12/04/2026 _____ *I confirm that I have accurately explained the study to the participant and answered all questions to the best of my ability.*

Study contact: Joe Grimshaw| jgrimshaw@tudelft.nl |Supervisor: D.S.Murray-Rust@tudelft.nl

Appendix F - Test 2 - Participant Consent Forms

PARTICIPANT INFORMATION SHEET

Template 1 — Informed Consent Materials

Study Title	Exploring Tangible Interfaces for AI-Mediated Music Discovery
Researcher	Joe Grimshaw | jgrimshaw@tudelft.nl
Supervisors	Dr. Dave Murray-Rust | D.S.Murray-Rust@tudelft.nl Prof. Gerd Kortuem | G.W.Kortuem@tudelft.nl
Date	24/04/2026
Version	v1.0

1 Introduction

You are being invited to participate in a research study titled **Exploring Tangible Interfaces for AI-Mediated Music Discovery**. This study is being done by **Joe Grimshaw**, a Master's student in the Department of Design for Interaction, Faculty of Industrial Design Engineering (IDE), at **Delft University of Technology (TU Delft)**. The project is supervised by **Dr. Dave Murray-Rust** (TU Delft) & **Prof. Gerd Kortuem**. This research has no external funding and is conducted solely as part of a Master's graduation project at TU Delft.

The purpose of this research study is to investigate how tangible (physical) interfaces may influence the way people explore and discover music using AI-mediated recommendation systems. Current music streaming platforms primarily rely on screen-based interaction and algorithmic recommendations. This study explores whether physical interaction with a prototype interface changes how users experience exploration and discovery, with potential implications for the design of future music interfaces and human–AI interaction more broadly. The study forms part of a Master's thesis that will be submitted to TU Delft and may be made publicly available through the TU Delft repository. Your session will take approximately **15–30 minutes** to complete. In this time, we will be asking you to interact with the prototype

2 What you will need to do

We will be asking you to interact with a series of experimental tangible interface prototypes, each combining physical controls with a digital user interface for exploring AI-generated music recommendations. Specifically, you will be asked to: perform simple music exploration tasks with each prototype; complete a short questionnaire about your experience after each prototype; and take part in a brief discussion with the researcher at the end of the session to share your reflections.

3 Data & Privacy

As with any online activity, the risk of a breach is always possible. To the best of our ability, your answers in this study will remain confidential. We will minimise any risks by doing the following:
Your name and email address will be collected to arrange the testing session. These details will be stored separately from your research data on the researcher's password-protected computer and TU Delft's secure systems, and will be deleted once your participation has finished.

During the session, the following personally identifiable research data may be collected:

- photographs or video recordings of your interaction with the prototype
- audio recordings of the discussion after the task
- questionnaire responses
- interaction logs (such as music selections and interface events)

Audio recordings will be deleted after transcription, and all transcripts will be anonymised before analysis. Any images or video used in publications will ensure that you cannot be identified (for example, by blurring or cropping faces or focusing on the prototype). Recordings will not be shared outside the research team.

Data Storage and Confidentiality

Anonymised research data (such as interaction logs and questionnaire responses) will be stored securely on the researcher's TU Delft OneDrive and kept for up to 10 years after the thesis is completed, in accordance with TU Delft data management guidelines. Only the researcher and supervisor will have access.

Publication and Open Data

The results of this study will be published as a TU Delft Master's thesis, which may appear in the TU Delft public repository and may be presented at conferences or submitted to academic journals. Any published material will use anonymised data only, and you will not be identifiable. Anonymous quotes from discussions may be used in the thesis or related publications.

Your participation in this study is entirely voluntary and **you can withdraw at any time**, without having to give a reason. You are free to omit any questions or tasks. **If you wish to have your data removed, please request this before within 3 working days of the session**; once data has been anonymised, it may no longer be possible to locate and delete your specific records. You also have the right to request access to, rectification of, or erasure of any identifiable personal data we hold about you (such as your contact details) at any time before that data is deleted. To exercise these rights, please contact the researcher directly using the details below.

4 Contact details and procedure for making complaints

Ifyou have any questions about this study before, during, or after your participation, please contact:

Corresponding Researcher	Joe Grimshaw jgrimshaw@tudelft.nl
Responsible Researcher(Supervisor)	Prof. Dave Murray-Rust D.S.Murray-Rust@tudelft.nl Department of Design for Interaction, Faculty IDE, TU Delft

AGREEMENT

By signing anddating the consent form on the following page, you confirm that you have read and understood this information sheet, that your questions have been answered to your satisfaction, andthat you agree to participate in this study on the terms described above.

Name (printed): _____Yael Kind _____

Signature: _____ Date: ___24/04/2026_____

Appendix F - Test 2 - Participant Consent Forms

EXPLICIT CONSENT FORM

Template 2 — Informed Consent Materials

Study Title	Exploring Tangible Interfaces for AI-Mediated Music Discovery
Researcher	Joe Grimshaw | jgrimshaw@tudelft.nl
Supervisor	Dr. Dave Murray-Rust | D.S.Murray-Rust@tudelft.nl Prof. Gerd Kortuem | G.W.Kortuem@tudelft.nl
Date	24/04/2026

PLEASE TICK THE APPROPRIATE BOX FOR EACH STATEMENT	YES	NO

A: GENERAL AGREEMENT — RESEARCH GOALS, PARTICIPANT TASKS AND VOLUNTARY PARTICIPATION

1. I have read and understood the participant information sheet dated [24/04/2026], or it has been read to me. I have been able to ask questions about the study and my questions have been answered to my satisfaction.	✓	□
2. I consent voluntarily to be a participant in this study and understand that I can refuse to answer questions and withdraw at any time, without giving a reason, and can request data deletion within 3 working days of the session.	✓	□
3. I understand that taking part in the study involves: • *Interacting witha series of experimentaltangible interface prototypes* • *Performing simple music exploration tasks* • *Completing a short self-administered questionnaire after each prototype* • *Taking part in a brief post-task discussion with the researcher.* *The post-task discussion will be audio-recorded only if I give separate consent below (see point 6).* *My prototype interactions may be photographed or video-recorded only if I give separate consent below (see point 6).*	✓	□

B: POTENTIAL RISKS OF PARTICIPATING

4. I understand that taking part in the study involves the following risks, and that these will be mitigated as described: *Emotional: Music can occasionally evoke responses such as nostalgia or mild sadness. The study does not intentionally aim to elicit strong emotions. I may stop the session at any time without explanation, and if I indicate any discomfort the session will be stopped immediately.*	✓	□

B2: DATA PROTECTION AND PRIVACY

5. I understand that to schedule my session, my name and email address (personally identifiable information, PII) has been collected. *I understand that this information will be stored securely on the researcher's password-protected computer and TU Delft systems only, will not be linked to my research data, and will be deleted immediately once my participation has concluded.*	✓	□
6. I understand and consent to the following personally identifiable research data to be collected during the session: • Photographs and/or video of my interaction with the prototypes. • Audio recording of the post-task discussion • My questionnaire responses and interaction logs (music selections and interface events). • I understand there is a small risk of re-identification from audio/visual recordings, which will be mitigated as described in point 8. *The following has been taken to minimise the threat of a data breach:* • *Audio recordings will be deleted promptly after transcription.* • *Transcripts will be anonymised before analysis and storage.* • *Any images or videos used in publications will ensure I am not identifiable (e.g., faces blurred or cropped, recordings focused on the prototype).* • *No recordings will be shared beyond the research team, and anonymised data will be stored securely on the researcher's personal TU Delft OneDrive with access restricted to the researcher and supervisor only.*	✓	□

C: RESEARCH PUBLICATION AND DISSEMINATION

7. I understand that after the research study, the de-identified information I provide will be used for the following purposes: a TU Delft Master's thesis (which may be made publicly available through the TU Delft repository); and potentially academic conference presentations or journal publications. All published material will use anonymised data and I will not be identifiable in any research output.	✓	□
8. I agree that my anonymised responses, views, or other input can be quoted in research outputs such as the thesis, conference papers, or publications. I understand that my real name will not be used in any quoted material.	✓	□

D: DATA STORAGE ACCESS AND REUSE

9. I understand that anonymised research data (questionnaire responses, interaction logs) will be retained for up to 10 years following completion of the thesis, stored securely on TU Delft systems in line with TU Delft data management guidelines.	✓	□

SIGNATURES

PARTICIPANT	Name (printed): _____Yael Kind _____ Signature: _____ Date: __24/04/2026_____
RESEARCHER	Joe Grimshaw (printed): Signature: _____ Date: _____ 24/04/2026_____ *I confirm that I have accurately explained the study to the participant and answered all questions to the best of my ability.*

Study contact: Joe Grimshaw| jgrimshaw@tudelft.nl |Supervisor: D.S.Murray-Rust@tudelft.nl