Towards an Ethical Framework for AI Music: End-to-End Infrastructure

Whitepaper by Soundverse Inc.
Co-authored by Sourabh Pateriya, Riley Williams and Linnea Sundberg | August 2025

Abstract

The rapid advancement of generative AI in music has outpaced the frameworks of attribution, copyright protection, and ethical data sourcing, leaving rights holders vulnerable to unlicensed exploitation of their works. Soundverse Trace and DNA: A Trust Layer for AI Music proposes a structured framework to bridge the gap between innovation and artist integrity through a multi-stage ethical pipeline. At its foundation, Soundverse emphasizes ethical data sourcing, rejecting one-time royalty buyouts as insufficient and advocating for ongoing participation of artists in the AI lifecycle. Building on this, Trace integrates data attribution technologies based on influence functions, watermarking, license tagging, and advanced deep search capabilities designed to identify, trace, and attribute AI-generated outputs back to their influences with forensic precision.

The whitepaper further introduces the Soundverse Partner Program, piloted with 50 creators in 2024, which establishes tiered royalty models across different stages of AI music creation and usage. This program demonstrated the feasibility of scalable licensing frameworks while highlighting the complexity of balancing ease of AI model training with equitable artist compensation. Insights from high-profile copyright disputes underscore the urgency of such a system, as courts increasingly demand transparent methods for tracing influence, identifying derivative works, and adjudicating disputes in AI-generated compositions.

By embedding Trace and DNA as a trust infrastructure between rights holders, regulators, and AI innovators, Soundverse positions itself as the industry's ethical alternative to companies deploying opaque and exploitative models. This approach not only ensures compliance with emerging legal standards but also builds a foundation for a sustainable ecosystem where artists, technologists, and consumers all benefit. In doing so, Soundverse charts a path toward a win-win future: where AI music innovation accelerates without replicating the digital colonialism of the past, and where attribution, licensing, and royalties are transparent, fair, and enforceable at scale.

About Authors

1. **Sourabh Pateriya**: Sourabh Pateriya is the Founder & CEO of Soundverse Inc., previously leading product and design at Spotify, Tobii Tech, and Samsung Research. He is a music-tech innovator with 10+ patents and a member of the Forbes Tech Council. At Spotify, he built music creation experiences and co-invented Basic Pitch, an open-source ML model for audio-to-MIDI transcription. As an electronic music producer (BlueNucleus), he blends creativity with deep technical insight. Sourabh spearheads Soundverse's mission to build ethical and controllable AI music infrastructure, focusing on data attribution, artist-centric monetization, and tiered licensing models.
2. **Riley Williams**: Riley Williams is the Co-Founder and CTO of Soundverse, with prior experience as a research engineer at Meta Reality Labs and applied scientist at Amazon Photos. He has over a decade of expertise in computer vision and multimodal machine learning, paired with his background as an electronic music producer. Riley designs and scales Soundverse's AI generation and attribution systems, ensuring controllability and reliability. His work centers on embedding licensing and rights management directly into the platform's technical core.
3. **Linnea Sundberg**: Linnea Sundberg is a Soundverse Board Member and VP of Corporate Development at UnitedMasters, with previously senior leadership roles at Spotify and Splice. At Spotify, she helped shape global artist compensation and creator-facing business models, aligning incentives between the platform and rights holders. She brings 15 years of experience in music, tech, and creator-economy business models. Linnea specializes in rights management, IP strategy, and creator monetization. At Soundverse, she crafts fair, transparent licensing systems that balance innovation with artist and label interests.

Table of Contents

Towards an Ethical Framework for AI Music: End-to-End Infrastructure **1**

1. Introduction 4
    1.1 Background? 4
    1.2 Why Act Now? The Risk Clock is Ticking. 5
    1.3 Future Value Creation, Not Just Risk: Early-Mover Advantage 5
2. Problems without attribution (by stage) 6
    2.1 Model creation - Ethical Data Sourcing and Training Practices 7
    2.2 Application layer 8
    2.3 Inference 10
        2.3.1 Transparent and Fair AI Model Development - Challenges in the Model Training Stage 10
    2.3.2 Responsible Music Generation and User Interaction 11
    2.4 Export 12
    2.5 External audio 13
    2.6 Compensation 13
        2.6.1 The Problem: Broken Compensation Models in AI Music 14
        2.6.2 One Time Licensing Fee 14
    2.7 The Cost of Inaction 15
3. The Ethical AI Music Framework (Solutions by stage) 15
    3.1 Model creation 16
        3.1.1 Ethical Solutions: Licensed Data and Artist Partnerships 16
    3.2 Application layer 17
        3.2.1 Soundverse DNA: Artist-Trained AI Models 17
            3.2.1.1 What is DNA? 18
            3.2.1.2 DNA Marketplace 18
            3.2.1.3 Economic Framework: Royalties and Pricing 18
            3.2.1.4 DNA Creation Workflow 18
            3.2.1.5 Artist-Controlled Metadata 19
            3.2.1.6 Lifecycle of a DNA 19
        3.2.1.7 Why DNA Matters 19
    3.3 Inference 20
        3.3.1 Explainable Models and Attribution Technology 20
        3.3.2 Guardrails, User Guidelines, and Creative Collaboration 22
    3.4 Export 23
        3.4.1 Solutions: Clear Licensing, Attribution Credits 24
    3.5 External audio 26
        3.5.1 Trace by Soundverse 27
    3.6 Compensations 28
        3.6.1 The Solution: Soundverse Content Partner Program 28
            3.6.1.2 Here are the stages of a Content Partner Program 33
            3.6.1.3 Why This Model Matters 33
            3.6.1.4 Proposal to the Industry 33
        3.6.2 Learnings from the April 2024 Pilot of Content Partner Program (50 Creators) 33
            3.6.2.1 Onboarding & Data Ingestion 33

3.6.2.2 Metadata Quality Matters for Attribution Accuracy 33
3.6.2.3 Quantity Matters, Audio Quality Matters More 34
3.6.2.4 Attribution Thresholds Calibration 34
3.6.2.5 Royalty Distribution & Transparency 34
3.6.2.6 Tier System Validation 34
3.6.2.7 Creator Trust & Buy-In 34
3.6.2.8 Economic Insight 35
3.6.2.9 Regulatory & Industry Signaling 35
3.6.3 Summary of Pilot Learnings 35
4. Comparison of Attribution Approaches: Data Attribution vs. Deep Search 35
4.3 Risks, trade-offs & mitigations 39
4.4 Legal alignment: “substantial similarity” vs “originality” 40
5. Trace’s Deep Search - The Deep Dive 40
5.1 Use cases 41
5.2 Depth of search (Higher the accuracy, higher the cost) 41
5.3 Search types 43
5.4 Operating thresholds and outcomes 43
5.5 Data Attribution (influence) - status and limits 43
5.6 Watermarking & provenance 44
5.7 For labels: solving the Suno/Udio problem with Trace 44
5.9 Interoperability & standardization roadmap 45
5.9 Implementation guide for a rights-holder 45
6. Open Research Questions 45
7. Conclusion and Future Outlook 46
8. Glossary 47

1. Introduction

AI-generated music has moved from a curious novelty to a transformative force in the music industry, raising profound ethical and legal questions. Recent advances in **generative AI** allow algorithms to compose songs, mimic musical styles, and even clone human vocals with unprecedented fidelity. This innovation promises new creative opportunities, but it also blurs the lines of authorship, ownership, and often, if not always, eschews artistic compensation. The central challenge is clear: **How can we harness AI's creative potential without undermining the rights and livelihoods of human artists?**

This whitepaper proposes a comprehensive **Ethical AI Music Framework** spanning the entire lifecycle of AI music creation - from dataset curation and model training to music generation and distribution. At each stage, we examine the problems that have emerged (e.g. unlicensed training data, "black-box" models, AI mimicry of artists, and unclear rights for AI outputs) and present solutions and best practices (such as licensed data partnerships, explainable attribution technology, content generation guardrails, and transparent licensing models). Our focus is on **Soundverse's artist-centric approach** as a case study in ethical innovation, compared against other industry approaches, to illustrate how an end-to-end framework can balance technological progress with creative rights. We also analyze key trade-offs - for example, **similarity-based detection** vs **data attribution** methods - which highlight the practical considerations in building fair AI systems.

Audience: This document is written for music industry stakeholders - record labels, music publishers, artists' rights organizations, regulators, and AI technologists - who seek a deeper understanding of ethical AI practices. By detailing the framework stage by stage, with supporting research and real-world examples, we aim to inform policy and business decisions that will shape the future of music. The stakes are high: billions in revenue and the integrity of creative culture depend on getting this right. An ethical AI music ecosystem is not just a legal necessity but a cultural imperative to ensure technology uplifts creators rather than exploits them.

1.1 Background?

- **Litigation risk has materialized.** Major labels (Sony, UMG, Warner) sued Suno and Udio for training on copyrighted sound recordings without permission; complaints were filed in federal court (D. Mass.; S.D.N.Y.). (RIAA)
- **European collecting societies are moving.** GEMA sued Suno in Munich seeking compensation and a licensing model for training and outputs that are "confusingly similar" to world-famous works. (CISAC, Music Business Worldwide, mckoolsmith.com)
- **Regulatory expectations are shifting.** The U.S. Copyright Office's 2025 report questions broad "fair use" defenses for commercial AI training and encourages licensed markets for training data. (copyright.gov, Nixon Peabody LLP)

Implication for labels: Detection must cover *both* (a) **training-time influence** (attribution) and (b) **output-time similarity** (search). In this we look into the problems and solutions across the stages and propose solution like 'Trace', 'DNA' and 'Partner Programs' etc. 'Trace' is intended to address both, with near-term deployment centered on **Deep Search**, and R&D continuing on **influence-based attribution**.

Generative-AI music will be adopted safely only if *attribution, licensing, and enforcement* are embedded across the **entire lifecycle** of model building and usage-not bolted on at the end. This whitepaper proposes a practical, auditable framework spanning **six** stages:

1. **Model creation**: Trustworthy data intake, automated tagging, partner onboarding
2. **Application layer**: Permissioned AI models, license tags, upload checks
3. **Inference**: Data attribution (influence tracking) + deep search
4. **Export**: Watermarking + license metadata preservation
5. **External audio**: Industry-scale similarity scanning and copyright checks
6. **Compensation**: Paying compensation to rights-holders

We detail missing controls at each stage, the remedies, and trade-offs. We then deep-dive into Deep Search (use cases, depth, and search types), and how Soundverse Trace, DNA, Partner Program operationalizes it for music labels to detect influence from Suno/Udio-like models and monetize or block accordingly. We also map Soundverse's holistic approach to legal standards in the U.S. and EU, and suggest a path to standardization.

1.2 Why Act Now? The Risk Clock is Ticking.

Labels/publishers need dual capability immediately:

Training-time influence evidence (attribution) and 2) Output-time similarity detection (search). Your enforcement, licensing, and product teams must operate on both axes or you'll be negotiating from a weak position.

Clock: Decisions you make in the next 1-2 quarters determine whether you (a) accumulate a defensible data/ provenance moat or (b) enter 2026 with fragmented tools, stale evidence, and weaker negotiating leverage.

1.3 Future Value Creation, Not Just Risk: Early-Mover Advantage

Why move now (upside):

- **Deep search takes preparation and compute time:** if you take actions now, it'll take some time for onboarding, agreements, system setup, and the actual deep search. It's a minimum 3-6 month process to get significant outcomes - if delayed would put unethical players in the powerful position.
- **Attribution share compounds:** early partner catalogs contribute to early model states and earn outsized attribution over time (larger "market share" of influence).
- **Tiered economics:** higher-rights tiers (voice/likeness/sync) validated materially higher value per attribution point in the pilot, improving ROI for sophisticated partners.
- **Live transparency:** dashboards with live attribution/payout signals drive trust and faster commercial loops.

Takeaway: Acting now isn't only defense; it's how you lock in recurring, data-driven income streams that scale with actual usage of your catalog.Additionally, it's put labels/ publishers/ artists in a power position giving them the opportunity and time to safeguard their rights and earnings.

2. Problems without attribution (by stage)

This section breaks down the cascading problems that arise when AI music systems operate without attribution or transparent rights frameworks. At the model creation stage, unlicensed training data and opaque datasets create legal risk and strip artists of consent and future royalties. At the application layer, there are no secure, private models or provenance tracking, leaving creators unable to control how their style or catalogs are used. During inference, the “black box” nature of AI means no logging of which data influenced outputs, making disputes nearly impossible to resolve. At export, licensing metadata is lost and no proof of provenance is embedded, which complicates enforcement. The absence of industry-wide external audio scanning exposes labels to infringement, and finally, compensation systems are broken, relying on one-time fees or offering no payment at all, disconnecting creators from the long-term value their works generate. Together, these gaps create a rights vacuum where innovation races ahead, but trust, transparency, and fair economics lag behind.

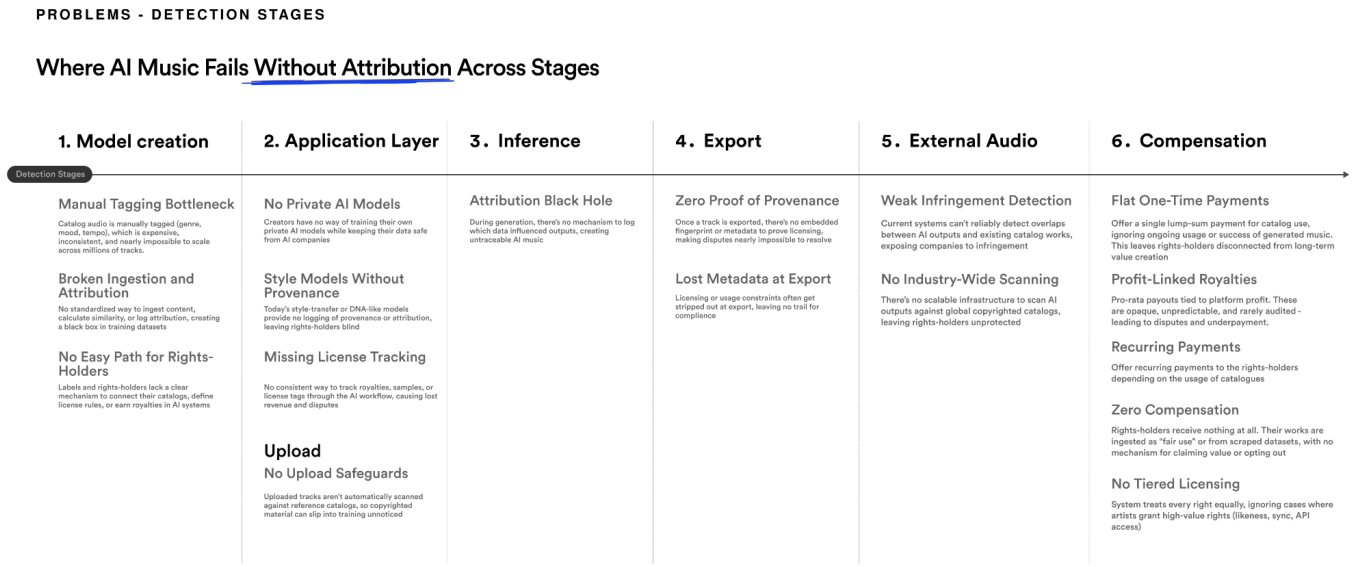
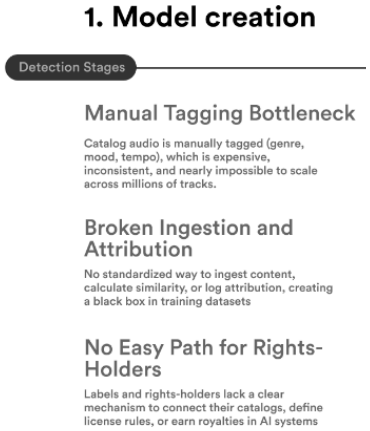


Diagram: Stages of an AI music system lifecycle - from training data ingestion, to model development, to music generation, and finally distribution/monetization - along with key ethical checkpoints at each stage (consent-based data sourcing, explainable models, responsible generation, and fair licensing).

2.1 Model creation - Ethical Data Sourcing and Training Practices



AI music models are only as ethical as the data they are trained on. Early systems used massive amounts of copyrighted music without permission, leaving rights-holders uncompensated and sparking lawsuits. Training datasets are opaque, making it impossible for artists to know if their work was used, and there is no standard for tagging, attribution logging, or setting license rules. This results in a "black box" where artists' labor fuels models without consent, transparency, or recurring economic participation.

Problems in the stage of model creation are:

1. **Manual Tagging Bottleneck:** Catalog audio is manually tagged (genre, mood, tempo), which is expensive, inconsistent, and nearly impossible to scale across millions of tracks. Inconsistent genre/mood/tempo/identifiers at catalog scale.
2. **Broken Ingestion and Attribution:** No standardized way to ingest content, calculate similarity, or log attribution, creating a black box in training datasets.
3. **No Easy Path for Rights-Holders:** Labels and rights-holders lack a clear mechanism to connect their catalogs, define license rules, or earn royalties in AI systems. Once the upfront payment is made, creators are excluded from the continuous economic activity their works fuel.

2.1.1 Challenges in the Training Data Stage

The foundation of any AI music model is its **training dataset**. In early AI music efforts, this foundation was often built on sand: many models were trained on vast catalogs of songs **without permission or compensation** to rightsholders (soundverse.ai). This practice is considered the hallmark of "unethical" AI training - not merely for legal reasons, but because it disregards the creative labor of artists. Companies like **Suno** and **Udio**, which trained models on scraped copyrighted recordings which they've accepted publicly, have become emblematic examples of this controversy (musicbusinessworldwide.com). Their tools learned from thousands of popular songs by ingesting the actual recordings wholesale - far beyond the way a human artist might casually be "influenced" by listening - and could then regenerate music in those styles, sometimes *uncannily* close to the originals

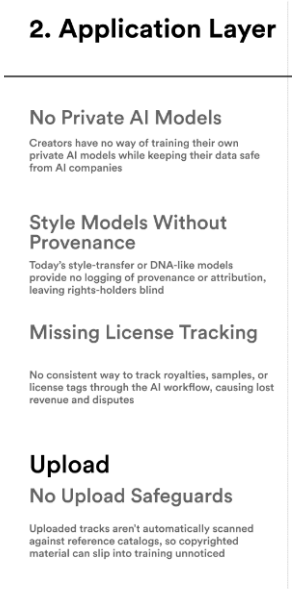
The ethical and legal pitfalls of such practices are now on full display. In 2024, major music labels (Universal Music Group, Sony Music, and Warner Music) filed landmark lawsuits against Suno and Udio, alleging **mass infringement** - essentially, that these AI platforms "*trained their models on copyrighted sound recordings at an almost unimaginable scale*," violating the rights of countless songs. Suno and Udio did not deny using the labels' music; instead, they attempted to justify it under the legal doctrine of *fair use*, claiming their outputs were sufficiently "transformative". This defense has drawn sharp rebukes from industry leaders, and regulators have begun to weigh in. In a recent U.S. Copyright Office report, authorities explicitly warned that the **knowing use of pirated datasets weighs against fair use**, affirming that copyright owners have the right to control how their works are accessed and

used in AI training. In Germany, songwriting society **GEMA** sued Suno after uncovering AI-generated songs that mirrored classics like Boney M.'s *"Daddy Cool"* and Alphaville's *"Forever Young"* nearly note-for-note. GEMA's evidence - side-by-side comparisons of Suno's AI outputs with famous tracks - underscored how **"confusingly similar"** the AI music was to specific copyrighted works. Such cases illustrate the risk: an AI model trained without ethical constraints can effectively produce *derivative works* or *soundalikes* that undermine the value of the originals. As GEMA's CEO put it, *"AI providers such as Suno Inc. use our members' works without their consent and profit from them"*, a practice that erodes the economic foundation of working musicians.

Another challenge at this stage is **lack of transparency**. Companies that trained on unlicensed data often kept their training sets secret, citing proprietary advantage. This opacity makes it impossible for artists or labels to even know if their songs were used, let alone to seek compensation. Without transparency, trust in AI tools plummets - as rightsholders fear an invisible theft of their catalog is occurring behind closed doors. In sum, the *status quo* in unethical data sourcing created a lose-lose scenario: AI developers faced litigation and mistrust, while artists faced uncompensated exploitation of their works and potential market harm from AI-generated copies.

Soundverse departs from this paradigm by embedding **recurring royalties** into its infrastructure. Instead of treating training data as a one-off commodity, Soundverse treats it as a **continuously accruing influence stream**, with payments tied directly to usage over time.

2.2 Application layer



Once models are built, creators still lack control over how their data is used. There is no way to train private, rights-aware models or to track provenance when style transfer or DNA-like models generate outputs. License tags and sample rights often disappear inside AI workflows, and uploads are rarely scanned for copyrighted material. Artists remain invisible participants in the system, and creators face technical and economic barriers that limit meaningful co-creation.

1. **No Private AI Models:** Creators have no way of training their own private AI models while keeping their data safe from AI companies. Creators can't train in secure silos with rights-aware logging.
2. **Style Models and Transfer Without Provenance**: Today's style-transfer or DNA-like models provide no logging of provenance or attribution, leaving rights-holders blind.
3. **Missing License Tracking:** No consistent way to track royalties, samples, or license tags through the AI workflow, causing lost revenue and disputes. No end-to-end royalty or sample/license flows.
4. **No Upload Safeguards:** Uploaded tracks aren't automatically scanned against reference catalogs, so copyrighted material can slip into training unnoticed. Copyrighted content can slip into training/reference sets.

For Artists:

a. Current AI platforms often scrape catalogs without consent, reducing artists to invisible training data with no compensation.
b. Musicians lack tools to monetize their *style* (not just their tracks) in the AI era.

c. Legacy royalty systems are fragmented, slow, and fail to reflect real-time usage.

For Creators:

d. Creative block and high barriers to entry limit productivity.
e. Access to professional studios, engineers, and years of production experience is costly.
f. Genre-blending is technically complex without expert knowledge.

For Consumers:

g. Generic AI outputs feel shallow, fun for a few minutes but lacking context or credibility.
h. There is no direct bridge between fandom and creation, no way to co-create "with" their favorite artists.

For rightsholders, the application layer is where revenue either becomes measurable—or leaks; controls installed here determine whether downstream attribution and payouts can work at all.

For Labels, Publishers & CMOs/PROs

1. **Catalog leakage & "style-by-proxy" risk.** Without permissioned models and app-layer controls, partners' catalogs can be used to generate "DNA-like" or style-transfer outputs that evade licensing while still trading on recognizable sonic identity. The current state lacks private, rights-aware models and robust provenance tracking at the application layer.
2. **Provenance blind spots = weak leverage.** Labels and publishers need both training-time influence evidence and output-time similarity detection; most application stacks offer neither, leaving enforcement teams with thin, post-hoc evidence packets and weaker negotiating positions.
3. **Royalty routing gaps from metadata loss.** License tags, sample rights, and usage constraints often disappear inside AI workflows before export; downstream, PRO/CMO distribution can't assign dynamic "AI shares" confidently when the app layer never captured or preserved them. (Result: unpayable overlaps, audit disputes, delayed settlements.)
4. **Deal asymmetry & settlement drag.** Absent app-layer logs (who used what, when, to what degree), labels are pushed toward broad, dataset-wide settlements rather than track/section-level claims—raising legal costs and shrinking recoveries. (This is a systemic issue called out across stages when provenance is missing.)
5. **Brand/moral-rights exposure from "confusingly similar" outputs.** Application-layer tools that enable soundalikes without guardrails amplify reputational and legal risk (e.g., GEMA's suit citing near note-for-note AI outputs). App-layer checks must prevent brand-diluting clones before they reach inference/export.
6. **Territory & repertoire governance shortfalls.** CMOs/PROs need territory-aware policy enforcement (consents, carve-outs, repertoire mandates). Today's app layers rarely honor granular license rules or region gating at upload/creation time, despite this being the natural place to enforce them. (The stage definition itself calls for permissioned models, license tags, and upload checks.)
7. **Net effect: measurable revenue leakage.** By the time inference/export tooling tries to catch issues, value has already leaked, usage went unlogged, tags were dropped, and enforcement starts late. This is why the **application layer** is the control point for prevention, not just after-the-fact detection.

2.3 Inference

Generation is where AI systems transform training data into outputs, but without attribution, it becomes a "black box." There is no logging of which songs or data influenced the music produced, so even licensed datasets cannot guarantee fair payout. True attribution requires understanding model influence, but deep learning models make this technically challenging and computationally expensive. Without solutions, infringement issues are discovered only after-the-fact, through audits or lawsuits.

- **Attribution Black Hole:** During generation, there's no mechanism to log which data influenced outputs, creating untraceable AI music. Nothing logs *which data* influenced generated tokens.

2.3.1 Transparent and Fair AI Model Development - Challenges in the Model Training Stage

Even with ethical data in hand, the process of training a music AI model presents its own ethical hurdles. Modern AI models - especially deep neural networks - are often **"black boxes"** where it's unclear how exactly the training data is being utilized internally to produce outputs. For music rights holders, this opaqueness becomes a critical issue: they want to know *which* parts of the model's output can be attributed to *which* training data. Without that, even a licensed dataset can lead to disputes if, for example, a particular artist's songs heavily influenced an AI-generated track but the compensation system doesn't reflect that. The core technical challenge is known as **AI attribution**: understanding and quantifying how specific training inputs affect the AI's generated output (waterandmusic.com).

Traditionally, in statistics and machine learning, one would use **influence functions** to tackle this problem. Influence functions measure the impact of a single data point on a model's predictions by conceptually retraining the model with and without that data point and seeing the difference. In principle, this is the gold standard for attribution - a *causal* measure of influence. For instance, imagine a generative model produced a new trap beat. To know if Drake's *"God's Plan"* or Kendrick Lamar's *"HUMBLE."* had a bigger influence on that beat, an influence-function approach would compare the AI's output with each of those songs removed from the training set. Whichever removal changes the output more indicates the song that had the greater impact. This method treats the model like a complex stew: remove one ingredient and see how the flavor changes. There are more methodologies such as gradient tracking and tracing, although slightly different - but worth exploring further.

However, applying pure influence functions to state-of-the-art music models is **computationally intractable**. Large generative models are trained on millions of songs; in theory one would have to retrain the entire model millions of times (each time leaving one song out) to exactly compute each song's influence. This is utterly impractical - it "*would take years even on advanced hardware*". Researchers have developed approximation techniques to estimate influence without full retraining, but these often **break down for neural networks**, yielding unreliable results. Neural nets are highly nonlinear and have redundant learned representations, so a simple approximation might say a song was influential when in fact the model would behave the same without it. In short, the straightforward path to attribution (leave-one-out testing) hits a dead end with deep learning.

Another challenge is that **influence in music models is diffuse and layered**. Unlike a simple regression where one data point might noticeably shift an output, a music generator learns abstract patterns (chord progressions, rhythmic grooves, timbral textures, etc.) spread across many training examples. A single generated song may incorporate dozens of micro-influences from the training data simultaneously. For example, an AI might internalize the *"four-on-the-floor"* kick drum pattern from hundreds of dance tracks, the use of a certain guitar tone from various rock songs, and the harmonic language of jazz standards - all culminating in one output. Tracing which specific training song contributed each element is extraordinarily complex. As the Water & Music research collective puts it,

“a single AI-generated track will draw from dozens of abstract layers of influence, which are challenging to map”. Moreover, models often include intermediate components (like a latent audio compression module trained on a broad audio corpus) that also affect the output in ways that are hidden from end-users. These *hidden dependencies* raise questions: Should data used in these intermediate stages also be credited, even if indirectly shaping every output?

Finally, **overfitting and memorization** are ethical concerns in model training. If a model memorizes parts of training songs and regurgitates them in outputs, it directly infringes copyright. Developers try to mitigate this via regularization and filtering, but it’s a fine line: we want the model to learn generalizable musical rules, not to store and copy exact clips. In early audio AI efforts like OpenAI’s Jukebox (2020), there were instances where recognizable segments of famous songs could be regenerated, showing the model had effectively memorized chunks of its training set. This risk is heightened if the training data is small or if certain songs are overrepresented. It’s an ethical duty in model development to ensure the AI isn’t just a “plagiarism machine” internally. Suno and Udio claimed their models did not reproduce any song *in full*, but as GEMA demonstrated with Suno’s outputs, even partial replication of melody and structure can occur. The lack of an attribution mechanism in those models meant such problems were only caught after the fact by human auditors (or lawsuits), rather than by the system itself.

2.3.2 Responsible Music Generation and User Interaction

Challenges in the Generation Stage

Once an AI music model is trained, it enters the **generation stage** - where end-users provide prompts or inputs and the AI produces music. This is where AI directly interfaces with creativity, and new ethical dilemmas arise regarding *how the tool is used*. One major concern is the ease with which AI can be used to **mimic specific artists or styles**, potentially infringing on an artist’s unique expression or deceiving listeners. For example, text-to-music generators might allow prompts like “create a song in the style of The Beatles” or even “generate a track that sounds exactly like [Artist X]’s voice and style.” Without constraints, this can lead to AI output that amounts to an unauthorized pastiche or impersonation of real artists. The Copyright Office has noted that even *“stylistic imitation may impact creators’ markets even without direct infringement” (*soundverse.ai). In other words, an AI doesn’t have to copy a melody verbatim to cause harm; simply flooding the market with sound-alike tracks can dilute the original artist’s economic opportunities. This is a form of **market harm** that regulators are watching closely. The notion of *“market dilution”* suggests that if AI-generated content saturates a genre or style, it could reduce the visibility and streaming of human creators, thereby shrinking their royalties. Ethically, we must consider: should an AI be allowed to generate infinite songs “in the style of Drake or Taylor Swift” if it means those artists (and songwriters) potentially lose revenue from substitute works?

Another issue is **voice cloning and deepfakes** in music. AI models now can replicate a singer’s vocal timbre with eerie accuracy. The Udio platform, for instance, was implicated in creating a viral deepfake track that mimicked Drake’s voice (often referred to as the “BBL Drizzy” incident). This track used Udio’s voice model to produce a song that many listeners initially thought was Drake singing. Such usage crosses into rights of publicity and artistic identity. Artists understandably worry that their voices could be used without consent to release new songs, potentially saying words or conveying messages they never agreed to - or even to flood streaming platforms with “fake” tracks that compete with their real work. From an ethical and legal standpoint, using an artist’s **likeness (voice)** without permission is problematic; it may not fall strictly under copyright (vocal style isn’t copyrighted, though the *recording* of a voice is), but it touches on issues of *misrepresentation and consent*. If left unchecked, fans could be misled by AI-generated performances, and artists’ reputations could be affected by creations they had no hand in.

User behavior with AI tools also needs guidance. Without proper **content policies**, users might attempt to generate outright infringing material - e.g. prompting the AI to recreate a specific copyrighted song ("Make a song that matches the melody of *Imagine* by John Lennon"). In fact, because these AI systems are stochastic, a user might repeatedly query the model to try to coax a close reproduction of a famous track. This kind of misuse can't be entirely prevented by the model training stage, so it must be addressed at the user interaction level.

Finally, there's the question of **human involvement** in generation. If an AI generates a song entirely by itself with no human creative input, current law (in the US and some other jurisdictions) would not recognize the output as a copyrighted work at all (soundverse.ai). While this might seem like a legal technicality, it has ethical implications: it means the user who generated the piece might not have exclusive rights, and also that the original training contributors aren't protected by traditional copyright on the output. Some argue that insisting on *"meaningful human contribution"* in the process (e.g. the user edits the AI output, writes lyrics, or mixes elements) can not only give the work copyright eligibility but also preserve the role of human creativity. The ethical framework thus may encourage a **collaborative AI usage** rather than fully automated generation that sidelines humans completely.

2.4 Export

At the point of export, the trail of attribution often disappears completely. Licensing metadata is stripped, there is no embedded fingerprint to prove compliance, and there is no standardized way to signal rights status downstream. This makes disputes hard to adjudicate and prevents automated royalty tracking. As a result, AI-generated music enters distribution with no clear proof of origin or license history.

1. **Zero Proof of Provenance:** Once a track is exported, there's no embedded fingerprint or metadata to prove licensing, making disputes nearly impossible to resolve
2. **Lost Metadata at Export:** Licensing or usage constraints often get stripped out at export, leaving no trail for compliance

Fair Distribution, Licensing, and Monetization of AI Music - Challenges in the Distribution Stage

The culmination of the AI music lifecycle is when the generated music leaves the lab and enters the world: being distributed to listeners, monetized on platforms, or used in media projects. This stage raises the pivotal question: **Who owns the AI-generated music, and who gets paid?** Without a framework, the default answers are problematic. As of now, purely AI-generated works (with no human author) are generally **not copyrightable** in jurisdictions like the US. That means the user who prompted an AI to make a song might not automatically have a legally protected right to exploit that song exclusively - a potential legal grey area for labels and content creators looking to release AI music. Moreover, even if the user claims ownership or a platform assigns ownership to the user, there's an underlying fairness issue: the AI's output was a product of a trained model that embodied the creative contributions of many artists in the training data. If those contributors are not acknowledged or compensated, we face a scenario where AI companies or users profit from music that, in spirit, carries pieces of countless human musicians' work.

From the perspective of record labels and publishers, one fear is a future where AI-generated tracks flood streaming services, **diluting the market** for human-made music. If AI songs (potentially built on their catalogs) take up listener attention and slice into the finite royalty pool, the original creators see reduced income - effectively being undercut by music that may owe its style to their own past creations. This is not a theoretical concern: streaming platforms like Spotify already had an incident in 2023 with an influx of AI-generated ambient tracks (via services like Boomy)

that led to the platform temporarily removing a large number of AI songs due to concerns of "algorithmic manipulation" of streams. The **"market saturation"** risk was explicitly highlighted by the Copyright Office, noting that AI's scale could cause *"unprecedented market damage"* if not managed.

Another challenge is the ambiguity around **royalties and rights for AI outputs**. Traditional music has well-defined rights (composition and recording rights) and royalty structures (mechanicals, performance, etc.) which assume human creators and identifiable works. AI outputs don't fit neatly. If an AI song inadvertently contains a melody very similar to an existing song, who is liable for that infringement? The user who generated it? The AI developer? Or consider if an AI-generated track trained on a specific artist's style becomes a hit - should that artist get a share, even if the song wasn't a direct copy? Without an agreed framework, either the original artists get nothing (leading to the ethical issue of uncompensated influence), or there will be after-the-fact litigation (which is inefficient for everyone). The lawsuits against Suno and Udio also highlighted this: the majors asked for statutory damages for each infringed work used in training, which could be astronomical. No label or AI company wants to be in a situation where every popular AI-generated song triggers a legal battle. What's needed is a **preemptive licensing and revenue model** so that by the time an AI track is distributed or commercialized, all parties' shares are already accounted for.

Finally, there's the issue of **public perception and authenticity**. If AI music is distributed without disclosure, audiences might be misled about who the artist is (some might assume a human band or producer, not an AI). There have been calls in some policy circles for labeling AI-generated content to maintain transparency. While not a legal requirement yet in most places, ethical practice might involve being clear that a piece is AI-generated, especially if using an artist's voice or a significant mimicry of style (to avoid false association). Balancing this, if an *human created AI augmented track*, the user might want to brand it under an artist alias/ project name as if it were any other music - that's a choice but ideally an informed one.

- Zero proof of provenance: no fingerprint/manifest at export.
- Lost metadata: usage constraints stripped in rendering/delivery.

2.5 External audio

Today, there is no global infrastructure to compare AI outputs against existing catalogs at scale. Detection tools are fragmented, low-recall, and low-precision, leaving rights-holders exposed to infringement. Without robust deep search and shared registries, AI platforms can unintentionally flood the market with infringing or confusingly similar works, creating legal and reputational risk for everyone in the value chain.

1. **Weak Infringement Detection:** Current systems can't reliably detect overlaps between AI outputs and existing catalog works, exposing companies to infringement. Low-recall/low-precision scanning.
2. **No Industry-Wide Scanning:** There's no scalable infrastructure to scan AI outputs against global copyrighted catalogs, leaving rights-holders unprotected. Fragmented coverage across DSPs/social/video.

2.6 Compensation

Even when attribution or licensing is acknowledged, payouts are often unfair or nonexistent. One-time buyouts disconnect artists from the long-term value their music generates, while opaque royalty pools and profit-linked payments are rarely transparent or audited. Many platforms offer zero compensation, citing "fair use," leaving

creators uncompensated as their catalogs fuel profitable AI products. This perpetuates historical inequities and sidelines artists from the economic upside of AI innovation.

Compensation answers *how those contributors are fairly compensated*. This stage is often overlooked by current AI music companies, yet it is the most consequential for creating an **ethical, sustainable economic system**.

1. **Flat One-Time Payments:** Offer a single lump-sum payment for catalog use, ignoring ongoing usage or success of generated music. This leaves rights-holders disconnected from long-term value creation
2. **Profit-Linked Royalties:** Pro-rata payouts tied to platform profit. These are opaque, unpredictable, and rarely audited – leading to disputes and underpayment.
3. **Recurring Payments:** Offer recurring payments to the rights-holders depending on the usage of catalogues
4. **Zero Compensation:** Rights-holders receive nothing at all. Their works are ingested as "fair use" or from scraped datasets, with no mechanism for claiming value or opting out
5. **No Tiered Licensing:** System treats every right equally, ignoring cases where artists grant high-value rights (likeness, sync, API access)

2.6.1 The Problem: Broken Compensation Models in AI Music

- **One-Time Buyouts**: As highlighted earlier, upfront licenses exclude creators from future participation, leaving them without recurring benefits.
- **Opaque Usage**: Even when royalties are promised, most AI platforms fail to provide transparency into *how* works are being used and *why* payments are calculated in a certain way.
- **Investor-First Economics**: Many platforms (e.g., Suno, Udio) are funded on rapid growth fueled by unlicensed training data. This approach prioritizes speed and valuation over fairness, with creators marginalized from the value chain.
- **No Tier Differentiation**: Current systems treat all rights equally, ignoring that some licensors grant broader rights (e.g., likeness, sync, API usage) that should command higher participation in payouts.

2.6.2 One Time Licensing Fee

A critical ethical challenge in AI music today lies in how training data is acquired and compensated. Many platforms (e.g., ElevenLabs Music, other one-time license marketplaces) adopt **single upfront payments** to rights-holders. While these deals may appear fair initially, they come with profound limitations:

- **Disconnection from Ongoing Value Creation**: Once the upfront payment is made, creators are excluded from the continuous economic activity their works fuel. If their music disproportionately shapes future AI generations, they see no additional benefit.
- **Asymmetry of Bargaining Power**: One-time deals are often struck when artists lack visibility into the long-term value of their data. Rights-holders effectively sell without knowledge of how much their catalog will power future AI creativity.
- **Perpetuation of Inequity**: Such models replicate historical patterns of undervaluing creators, where platforms and investors capture compounding returns while artists receive diminishing compensation.
- **Rights-holders adding an AI licensing clause**: Before signing a creator, right-holders often add an extra clause where creators unknowingly permit their data to be trained by AI. So artists are often not available,

2.7 The Cost of Inaction

If labels/ publishers/ rights-holders wait 6–12 months:

- **Revenue dilution:** AI look-alikes flood channels; even without verbatim copying, style-level substitution erodes your catalog's share of finite royalty pools.
- **Weaker enforcement:** absent provenance and search at export and external-audio stages, evidence packets are slower and costlier to assemble—shifting leverage to counterparties.
- **Negotiation penalty:** you'll be forced into generic, dataset-wide settlements because you can't show track- or section-level influence.

Counterfactual (act now): Bind rights + payouts at ingestion, deploy similarity scanning on external feeds, and accrue influence logs so you can **monetize or block** with policy-driven thresholds.

3. The Ethical AI Music Framework (Solutions by stage)

This section lays out the solution for an end-to-end trust layer for AI music - turning opaque, risk-heavy workflows into auditable, rights-aware pipelines. Across six stages, we replace scraping with licensed intake, black-box generation with attribution and guardrails, and ambiguous distribution with watermarking, license tags, and automated payouts. The result is a practical, enforceable framework that aligns creators, labels, platforms, and regulators, so innovation scales with consent, provenance, and compensation built in.

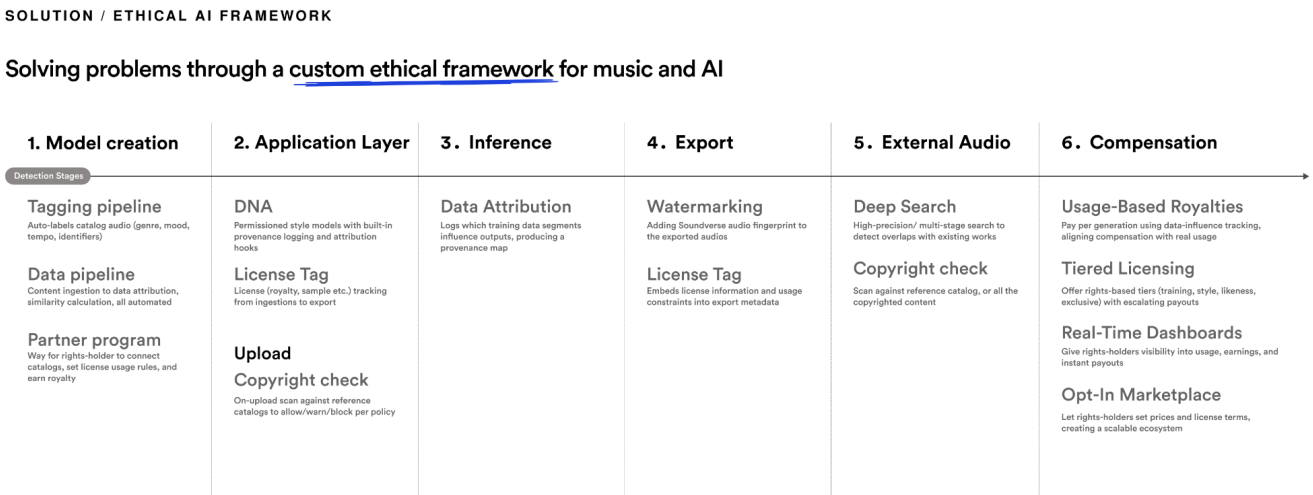


Diagram: Stages of an AI music system lifecycle and solution for each stage - from training data ingestion, to model development, to music generation, and finally distribution/monetization - along with key ethical checkpoints at each stage (consent-based data sourcing, explainable models, responsible generation, and fair licensing).

3.1 Model creation

This stage is the foundation of the entire ecosystem. It focuses on preparing, cleaning, labeling, and connecting source data in a way that is structured, rights-compliant, and scalable for training machine learning models. Rather than treating raw audio as an unstructured dump, this stage ensures that every file carries the right metadata, provenance information, and licensing rules from the start — making downstream attribution and compensation possible.

1. **Tagging Pipeline**: A fully automated system that analyzes incoming audio and assigns detailed metadata, including genre, mood, tempo, key, and unique identifiers. This improves dataset quality, enhances model training, and allows for granular filtering or search later.
2. **Data Pipeline**: Manages end-to-end ingestion, cleaning, normalization, and storage of catalog data. It also calculates similarity fingerprints and prepares attribution hooks so that every audio file's contribution to a model can be traced back.
3. **Partner Program**: Provides an interface for rights-holders to onboard their catalogs, define their preferred license usage rules (e.g., allow for training but not direct sampling), and start earning royalties through a fully automated, trust-building process.

3.1.1 Ethical Solutions: Licensed Data and Artist Partnerships

To address these issues at their root, the Ethical AI Music Framework begins with **consent-based, licensed data collection**. Rather than scraping the internet indiscriminately, ethical platforms curate training datasets through partnerships with artists, labels, and music rights holders. For example, **Soundverse** has pioneered a *Content Partner Program* (https://www.soundverse.ai/partner) that invites artists and rights holders to contribute their music to the training pool *voluntarily* and **in exchange for compensation**. In this model, every song in the training data is there with explicit permission. Soundverse proudly trains its AI models *"exclusively on licensed or copyright-free datasets,"* avoiding the "pirated dataset" problem altogether. This approach required forging relationships and trust with the creative community, but it pays off by building an AI that is ethically and legally grounded from day one.

How it works: Rights holders who join the Content Partner Program upload their tracks into the system (soundverse.ai). The platform's AI automatically checks and annotates the content (e.g. verifying the uploader's rights, tagging metadata). Once onboarded, these tracks become part of the training corpus for the AI music generator. Crucially, contributors are not donating their music in vain - they are entering a revenue-sharing arrangement. Soundverse's system tracks if and when a partner's music influences any AI-generated output, and the artist earns **royalties whenever their data is used** in the generation of new music. In effect, this creates a **data licensing marketplace**: artists provide training data under clear terms, and receive payment downstream, akin to licensing a sample or any other use of their work.

Other companies are following similar paths. For instance, the startup **LANDR** launched a "Fair Trade AI" program where artists can opt in their music for training, and in return 20% of all revenue from the AI product is paid back to contributors on a pro-rata basis. This simpler model avoids needing to measure each song's influence precisely (it just splits revenue by each artist's percentage of the dataset), which provides **predictability and transparency** to participants. Likewise, major industry players are striking **direct licensing deals**: in 2025, ElevenLabs (known for voice AI) introduced **Eleven Music** and announced partnerships with Merlin and Kobalt - organizations representing large catalogs of indie labels and publishers (musicbusinessworldwide). ElevenLabs made a point to **source data directly from rightsholders** and assure them their works are protected and compensated in the AI model. By securing these deals, Eleven Music can train on high-quality music libraries without legal risk, and rightsholders gain new revenue streams from AI.

A key advantage of consent-based training is **building trust and collaboration** rather than adversarial relations. When artists are invited to the AI table, they become stakeholders in the technology's success, not victims of it. As Soundverse's CEO explains, *"building our technology around ethical practices was essential… we envision a future where people and artists are at the center of the ecosystem, not AI" (*musically.com). By treating AI as a tool to **augment creativity** rather than replace it, Soundverse's ethos encourages artists to participate. Indeed, Soundverse's ethical stance has attracted over 200,000 creators to its platform by late 2024, collectively contributing to over half a million AI-generated tracks - a strong indication that many musicians will embrace AI if it's built on respect and shared benefit.

It's worth noting there is a trade-off in this stage: limiting training data to only licensed sources might constrain the size and diversity of the dataset compared to scraping the entire internet. However, initiatives like content partner programs are rapidly scaling up. By mid-2025, Soundverse had partnered with multiple independent labels and artists, and ElevenLabs' partnership with SourceAudio granted it access to **millions of pre-cleared songs** via a licensing program (musicbusinessworldwide.com). The industry is moving toward **collective datasets** that are both large and rights-cleared, as stakeholders realize the importance of ethical AI. Moreover, an emerging movement to certify "**fairly trained**" AI models (spearheaded by industry veterans) is putting positive pressure on developers to use only consented data ().

In summary, Stage 1 of the framework transforms data sourcing from a legal gray area into a collaborative, transparent foundation. By replacing the clandestine scraping with **cooperative licensing**, the AI model begins life on solid ethical ground, setting the stage for accountability in all subsequent phases of creation.

- **Auto-tagging pipeline.** ML tagging for genre, mood, tempo, key, timbre; quality gates + human QA for gold-standard subsets. (Accuracy ↑; ongoing bias monitoring required.). This automatic tagging can be extremely powerful when used for audio to description based attribution search. Additionally, this tagging pipeline makes search at an unimaginable scale extremely easy.
- **Data pipeline.** Normalized ingestion > similarity pre-compute > rights metadata binding.
- **Content Partner Program.** Onboard catalogs under agreed license terms; attach per-work rights metadata and usage rules; enable royalty pathways on detected influence.
- **Outcome:** A model-ready, rights-annotated corpus.

3.2 Application layer

This layer is where data meets user-facing functionality. It builds the mechanisms for controlled access, permissions, and tracking, ensuring that model usage stays within defined bounds and remains auditable.

- **DNA** – Represents permissioned, artist-style models with built-in provenance tracking. Each DNA carries a "loggable identity" so that every output can be linked back to the artist or dataset that inspired it, creating both accountability and monetization opportunities.
- **License Tagging** – A rules engine that attaches license types (training, sampling, commercial use, exclusive rights, etc.) to every dataset and propagates those rules through inference and export stages. This prevents unlicensed usage and ensures proper payout calculations.
- **Upload & Copyright Check** – On ingestion of any new dataset or user-uploaded content, the system performs a fast but robust scan against reference catalogs. It automatically flags, blocks, or warns based on rights-holder preferences, reducing liability and avoiding disputes.

3.2.1 Soundverse DNA: Artist-Trained AI Models

3.2.1.1 What is DNA?

A **DNA** is an AI model trained on a specific artist's catalog, capturing their sonic identity-the timbre, stylistic fingerprint, melodic tendencies, and production signatures that define their body of work. Unlike generic AI models, DNAs are ethically licensed, artist-controlled, and designed to monetize creative essence in a transparent, artist-first way.

Soundverse positions DNAs as the bridge between **artistic integrity** and **AI accessibility**: a system where creators, fans, and industry professionals can safely and legally generate music inspired by real artists.

Soundverse DNA is the solution layer where **artists become co-owners of their AI presence.**

- **Private Mode:** Artists build DNAs as a personal AI co-creator, enabling faster song ideation and remixing.
- **Public Mode:** Artists publish DNAs to the **DNA Marketplace**, where fans, collaborators, and prosumers can generate tracks in their style.

3.2.1.2 DNA Marketplace

The Marketplace allows:

- **Safe exploration:** Users can play with DNAs in-app, generating drafts powered by Soundverse tokens.
- **Paid exports:** Users purchase clearances to download or commercially use outputs.
- **Artist-set pricing:** Exports cost $0.99 to $9.99, with full price control left to the artist. (These numbers are tweakable)

3.2.1.3 Economic Framework: Royalties and Pricing

- **DNA Creation:** Free for artists. Sandbox creation incurs no costs.
- **DNA Usage:** Each generation costs Soundverse tokens (users buy tokens via credit card or in-app store).
- **DNA Export:** Paid event, triggering instant royalty split. Pricing between $0.99 to $9.99 to keep it economical for consumers.

Royalty Split:

- 50% to the artist/ label/ right-holder .
- 50% to AI Host (infrastructure, marketplace, and compliance operations).

While these numbers are tweakable, we believe as the demand and supply grows, the cost associated with AI model training and hosting at scale will significantly come down. We can't put a number to how low yet, but if AI compute costs and cloud storage costs drop significantly, and demand goes dramatically up - it may reach single digits for the AI host.

This system ensures artists earn from every monetized use of their DNA, unlike traditional one-time licensing models.

3.2.1.4 DNA Creation Workflow

Two Onboarding Paths:

1. **Upload Catalog:** Artist bulk-uploads tracks, stems, demos, and metadata (lyrics, production notes, annotations).
2. **Claim Profile:** Artist proves identity (via KYC and label verification) and links their catalog (Spotify, Apple Music, YouTube).

DNA Creation Times:

- On **CPU**: 2 minutes to 3 hours for ~100 audio files.
- On **A100 GPU**: 20 seconds to 15 minutes for ~100 audio files.

Sensitivity Selector:

- Automatically clusters catalog into multiple DNAs (e.g., “Acoustic Era DNA” vs. “Electronic Era DNA”).
- Label/ Artist can merge, split, or fine-tune clustering to capture stylistic periods or collaborative phases.

3.2.1.5 Artist-Controlled Metadata

Artists directly shape their DNAs by providing:

- **Audio Data**: Tracks, stems, unreleased demos.
- **Annotations (Powered by Soundverse Content Intelligence APIs)**: Lyrics, moods, genres and micro-genre extraction.
- **Profile Data**:
    - DNA name & cover art.
    - Price per export (from $0.99 to $9.99).
    - Permitted use cases (UGC only, sync licensing, commercial use).
    - Royalty split rules (default 50/50, but adjustable for collaborations).
- **Discovery Settings**: Whether DNA is public or private.

3.2.1.6 Lifecycle of a DNA

1. **Creation:** Artist uploads catalog; training completes in ~10 minutes on GPU.
2. **Marketplace Listing:** Artist sets DNA profile name, pricing, and allowed use cases.
3. **Usage:** Fans generate drafts in-app using tokens.
4. **Export & Royalties:** When a user pays $2.99 to export, Soundverse auto-splits royalties with the label/ artist.
5. **Transparency:** Artist dashboard tracks usage, influence, and real-time payouts (quarterly reconciliation for larger tiers).

3.2.1.7 Why DNA Matters

- **For Artists:** First ethical way to monetize *style* in the AI era, creating a new revenue stream beyond streaming.

- **For Fans:** The chance to *co-create* with favorite artists, generating tracks that feel authentic, not random.
- **For the Industry:** A transparent attribution system that shows AI can respect rights, create revenue, and enhance fandom rather than exploit it.

DNA transforms AI from an extractive threat into a participatory, artist-first ecosystem. By enabling ethical monetization of style, direct artist control, and real-time royalty distribution, it ensures AI music benefits creators, consumers, and rights-holders alike.

3.3 Inference

This stage ensures that model outputs are explainable, attributable, and defensible. Rather than treating generative AI as a "black box," this layer creates a transparent audit trail of how outputs are formed.

- **Data Attribution** – Every time the model generates music, it logs which training data segments influenced the output, creating a "provenance map." This is critical for trust, as it allows both rights-holders and end-users to see how data was used and for payouts to be aligned with actual influence rather than rough estimates.

3.3.1 Explainable Models and Attribution Technology

Addressing attribution at the model level requires blending *technical innovation* with *practical compromise*. An ideal solution is to build **explainability** into the model: techniques from the field of **Explainable AI (XAI)** can be adapted so that the model provides a rationale or trace for its outputs. Soundverse, for example, is investing in advanced attribution tools as part of its "Ethical AI" roadmap, aiming to go beyond simplistic dataset-wide royalty splits to **track influence at the model level**. While true influence functions are computationally expensive, Soundverse and others are exploring clever ways to approximate a similar result. One approach is to compute **gradients or influence scores** for each training example with respect to a given output without full retraining. These methods, inspired by influence functions, attempt to estimate "if song X were absent, how might this output change" in a more efficient way. Early research (e.g. TracIn, Influence Calculus) has shown some promise in ranking which training points most impact a particular output, but reliability is still a work in progress. Soundverse's team describes the ongoing development of a *"patent-pending in-house technology for music attribution,"* noting it was one of the toughest challenges they've faced. This underscores that **deep integration with the model's architecture** is required - a solution likely needs to be custom-built for each generative model type (whether it's a diffusion model, transformer, or hybrid system). In other words, attribution isn't a plug-and-play feature; it must be woven into the AI's design, which takes intensive R&D and may differ for each algorithmic approach. *For instance, integrating an attribution mechanism into Meta's MusicGen model might require different techniques than doing so for a diffusion-based audio model, given their internal representations.* This lack of one-size-fits-all solution means true data attribution currently **doesn't scale across different AI models** - each one demands a bespoke analysis framework.

While working towards model-based attribution, ethical AI developers also implement **pragmatic safeguards** to reduce egregious copying. One such safeguard is output **similarity checks** (discussed more in Stage 3 and the trade-offs section). After the model generates music, the system can automatically compare the output to the training database (or a broad music catalog) using audio fingerprinting or embeddings. If the output is too similar to a specific existing song (above some threshold), the system can flag or block it. This acts as a "plagiarism filter" to catch cases of memorization or reproduction. It doesn't explain *why* the model generated that result, but it can

prevent direct infringements from being released. For example, if an AI inadvertently spits out something that matches *"Happy Birthday"* or a Beatles melody, a deep search would catch it before it reaches the user. Many AI platforms quietly include such **guardrails**, using technology akin to YouTube's ContentID or Shazam to detect overlaps. However, relying solely on output similarity is reactive and **keeps the model a black box** - you're not peering inside to see influence, just inspecting the final product. Rights holders have pointed out that this is a limited comfort: *embeddings can flag what sounds alike, but cannot reveal why it sounds alike*. A song might resemble Taylor Swift's style in the output, but without model transparency one can't tell if it's because the AI trained on Taylor's songs or simply coincidentally arrived at a similar style.

Thus, the emerging best practice is a **hybrid approach**: develop as much internal attribution capability as feasible, and complement it with external similarity analysis for safety. Soundverse illustrates this hybrid strategy in its platform. On one hand, they are pioneering **"influence function" tools internally (cohort-level and eventually song-level tracking) to allocate credit**, and on the other hand, they promote **"cohort-based attribution"** as a scalable interim measure. Cohort-based attribution means grouping training data into meaningful categories (by genre, era, style) and attributing an output's influence to a cohort rather than an individual song. For example, if an AI output is determined to be in a synth-pop style, then all artists in the "synth-pop" cohort of the training data might share in the credit and royalties. **Lemonaide**, another startup, is pursuing this method: their system categorizes outputs by sub-genres and distributes revenue to all rights holders whose music falls in the dominant genre of the output. This sacrifices precision (it's more **correlation than causation**, essentially rewarding whoever is in the same style group as the output) but offers a practical path to scale. This approach leads to highly non-uniform distributions that could mean either one genre is dominantly generated, or consists of the largest 'training pool' of data. Additionally, it is more easily "gamed" in the sense that someone could target the most generated and upload low-quality audio to get a larger payout.

Soundverse's approach is to use such cohort or pro-rata models as a fallback to ensure everyone is paid something, while striving to refine the granular attribution over time. Second, it is more easily "gamed" in the sense that someone could target the most generated and upload low-quality audio to get a larger payout.

In simpler terms, ethical model development requires confronting a **trade-off**: *Black-box simplicity vs. explainable complexity.* Companies without attribution tech lean on simpler deals (fixed payments, pro-rata shares) because they **"cannot offer usage-based royalty models"** without knowing usage. This is easier to implement and explain (everyone understands a flat percentage), and it has advantages like predictability and minimal technical overhead. But it may not feel perfectly fair - some contributors might have heavily influenced an output yet get the same cut as others who had little influence. On the other hand, pushing for true data attribution offers fine-grained fairness but at the cost of **complexity and uncertainty** (influence algorithms might be hard to validate, and results might be so fragmented - e.g. 0.03% influence from one song - that it's hard to interpret). Soundverse's stance is that the long-term **innovation** in this stage is worth the effort: they describe their attribution system as a "roadmap for how to compensate artists fairly, even when the lines between inspiration and output get blurry". By building attribution into the *DNA* of the AI platform, Soundverse aims to eventually provide each artist with a clear answer to *"Did my music help shape this AI-generated song, and if so, how much?"* - and then pay them accordingly. Achieving this at scale will likely require continued research (possibly breakthroughs in XAI for generative models) and iterative improvement of algorithms like influence function approximations. Until then, a combination of reasonable heuristics (like similarity-based cohort crediting) and robust **transparency with rights holders** will maintain trust. As one AI music CEO noted, *"attribution is as much a business, ethics, and trust problem as it is a technical problem"*. Therefore, being forthright about what the model can and cannot yet explain is crucial in Stage 2. Platforms that manage expectations and invite stakeholder oversight (for example, allowing label partners to audit the training data or test the AI outputs for unwanted similarities) demonstrate the good faith needed to let the technology mature responsibly.

3.3.2 Guardrails, User Guidelines, and Creative Collaboration

An ethical AI music platform implements **responsible generation practices** to ensure that the technology augments creativity without facilitating impersonation or infringement. One concrete solution is to build **guardrails into the prompting system**. For example, ElevenLabs' Eleven Music platform explicitly bans users from inputting any real artist names, song titles, or other identifiers in their prompts. Their terms of service list *"prohibited inputs"* such as "any artist's real name or stage name" or "any song title". This prevents a straightforward way of asking the AI to copy a specific artist or song. Additionally, ElevenLabs prohibits users from generating output that they know is likely to infringe rights or **mimic a particular artist's voice or likeness in a misleading way**. These rules are backed by technical measures: presumably, the model or a pre-filter will reject prompts like "Give me a song that sounds exactly like *Bad Guy* by Billie Eilish" or "Use Freddie Mercury's voice." By disallowing direct references to protected material, the platform nudges users towards more original or generalized creative directions (e.g. "generate an upbeat pop song with female vocals" - which is style guidance, but not impersonation).

Soundverse follows a similar philosophy through design choices in its generation tools. Rather than offering a feature that says "clone this artist's style," Soundverse provides what it calls a **"Similar Song Generator"** or "vibe-twin" functionality. According to their documentation, this tool lets users input a reference track to capture a general mood or aesthetic, but it **avoids replicating specific melodies or vocals**. The idea is to create a new work that *feels* like the reference in vibe (tempo, energy, instrumentation) without crossing into plagiarism of composition. This aligns with an ethical stance: *inspiration instead of imitation*. Soundverse describes the results as "inspired by" rather than "identical to" the reference, consciously aiming to **evoke a style without copying** protected elements. In effect, they're encoding a principle of **"style guidance without mimicry"** - helping users achieve a certain genre or tone while steering them away from explicitly imitating an individual artist's signature elements. Such balancing acts may be reinforced by automatic checks; for instance, after using a reference track, the system could employ an audio similarity metric to ensure the generated piece isn't too close to the reference's melody or structure (similar to how some image AIs ensure outputs aren't too close to any single training image).

On the vocal front, ethical frameworks will typically **require consent for voice cloning**. Some platforms like Grimes' Elf.Tech have pioneered artist-sanctioned voice models - Grimes allowed anyone to generate music with her voice *provided they share royalties with her* - turning a potential infringement issue into an opt-in creative collaboration. Soundverse's approach to vocals (through its *AI Singing Generator* feature) emphasizes user-provided vocals or AI vocals that are not direct clones of real singers unless licensed. In practical terms, a platform could restrict the voice models available to either public domain voices or those it has legally licensed/created. Soundverse would never offer a famous-voice model, without artist's permission and a revenue share (given their overall ethos). If a user tries to misuse vocals (say, upload samples of a celebrity voice to train a custom voice model), a robust content policy and detection algorithm should intervene. We see this with some companies which have started tagging AI voices and requiring disclosure if a voice is synthetically generated, aiming to avoid deception.

Another safeguard in generation is **real-time plagiarism detection**. As mentioned earlier, deep search algorithms can be run on outputs. An ethical AI music system might implement a final check before delivering a track to the user: compare the generated audio against a database of known music (perhaps the training set plus commercially released songs via a fingerprint service). If a high similarity is detected (above a certain threshold in melody or audio fingerprint), the system could either: (a) alert the user that the output closely resembles an existing work and cannot be used as-is, (b) prevent the download and ask the user to retry with different parameters, or (c) automatically adjust the output (if feasible) to be more original. This kind of measure ensures that even if the user unknowingly or intentionally steers the AI toward a specific song, the platform won't produce a carbon copy that infringes copyright. It's an analogous concept to how Turnitin checks student essays for plagiarism against a corpus - here the AI checks itself against the music corpus. These **output guardrails** were something Suno and Udio apparently lacked initially (or were insufficient), as the lawsuits cited their outputs "closely mimicking" well-known songs. Learning from

that, newer platforms treat output filtering as essential. For instance, Udio later claimed to implement filters to prevent exact copies, and ElevenLabs' strict T&C suggests they have automated checks to enforce them.

Importantly, ethical AI generation also involves **user education and collaboration features**. Rather than pitching AI as an "instant hit song" button, Soundverse and similar platforms present it as a **co-creation tool**. Soundverse integrates a voice-controlled assistant (Soundverse Producer, previously known as SAAR) and real-time collaboration in its interface, encouraging musicians to use AI to **assist** their workflow (e.g. generate a beat, suggest a chord progression) while they continue to contribute creatively. By positioning the AI as a partner, the platform implicitly promotes *human oversight* and *creative input* at each step, mitigating the "AI replaces artist" narrative. This has a couple of ethical benefits: it keeps the human creator in charge (satisfying the likely requirement for human authorship for copyright and creative accountability), and it reduces the risk of blindly publishing AI output that might have hidden issues. A musician working *with* AI will review and edit the final output , meaning if something sounds too similar to a known song or if the AI-produced lyric is problematic, the human has the opportunity to catch it. Fostering this collaborative mindset is part of an ethical framework: **AI is a tool *in the artist's hands*, not an autonomous composer releasing music without oversight**.

By implementing these guardrails, platforms like Soundverse not only reduce legal risk but uphold an ethical standard: an AI-generated song should *not* cross the line into a derivative of someone else's work, and it should ideally involve a human creator's personal touch and approval. The output of the AI, in an ethical setting, becomes a jointly crafted piece - with clear provenance of influence and respect for boundaries, rather than an orphaned creation with murky origins. This helps to ensure the **creative process** remains an essential component of ethical creation and legal compliance, as it has always has, as opposed to only being examined by a model's output.

In Conclusion,

- **Data Attribution (R&D).** Soundverse explores model-integrated attribution that traces token/feature flows linking outputs to contributing training items (influence functions and activation tracing). This aligns with the Water & Music interview describing Soundverse's end-to-end token-flow logging within its own architecture. Water & Music
    - *Trade-off:* Architecture-specific; high scientific value (causation signal) but not yet portable across third-party models. Influence-function methods help but are brittle on deep non-convex nets and require approximations.
- **Prompt and input restrictions** to prevent explicit infringement requests (no using exact artist names or songs as prompts).
- **Impersonation safeguards** to protect artists' voices and styles (no misleading use of real artist vocals or branding).
- **Post-generation checks** against existing works
- **User guidelines and UX design** that encourage human creativity (presenting AI as assistive, requiring user edits, etc., aligning with a *human-AI collaboration* model).

3.4 Export

At the point of output, this stage ensures that generated files are both traceable and rights-compliant once they leave the platform. It "seals" every audio file with identity and license information so downstream ecosystems can read, verify, and respect usage constraints.

- **Watermarking** – Embeds a unique Soundverse audio fingerprint that survives transformations like compression or format change. This enables later identification and enforcement if an exported file is reused without proper attribution.
- **License Tag** – Inserts machine-readable license metadata (e.g., royalty-bearing, sample-cleared, personal use only) into the exported file. This means even if the file travels outside the Soundverse ecosystem, its rights information remains intact.

3.4.1 Solutions: Clear Licensing, Attribution Credits

In an ethical AI music framework, the distribution stage is governed by **built-in licensing agreements and royalty mechanisms** that ensure all stakeholders - the user, the platform, and the original rights contributors - are protected and rewarded. Soundverse's platform is an instructive model here. When a user creates a song with Soundverse AI, they are not left in a legal vacuum. Instead, Soundverse provides an **integrated licensing pipeline**: users can immediately license the AI-generated track for various uses (e.g. as background music for videos, commercial releases, etc.), with the terms made clear up front (soundverse.ai). For everyday creators, Soundverse offers a default **royalty-free license** for the generated music, meaning the user can use the track in their content without having to pay further royalties - *because the necessary royalties to training data contributors are already handled on the back-end*. In essence, Soundverse acts like a new kind of music publisher or PRO, automatically managing the royalty splits behind each AI song so that by the time it reaches the end-user, it is "pre-cleared" for use. This prevents scenarios like a streamer using an AI track and then getting a takedown because it sounded like someone's song - if it's from Soundverse, it's **"copyright-safe"** by design.

How are the original artists paid? Soundverse's *Trace* monitors the influence of each training song on generated outputs and allocates a share of revenue accordingly based on calculations explained later in the whitepaper. Soundverse mentions providing near **real-time royalty dashboards** to their content partners, where artists can see when their music helped train a popular track and how much they earned from it. This model-level tracking of influence is novel and ambitious. Even if the exact percentages are approximate (given the technical challenges described earlier), the principle is that artists receive a *direct stake* in the success of AI outputs that their work enabled. This stands in stark contrast to the opaque approach of other AI platforms. For instance, users of Suno/Udio were told they "owned" the AI-generated songs, but since those songs were potentially infringing, that ownership was on shaky ground (and no thought was given to compensating the original artists). Soundverse flips it: the *platform* ensures legality and fairness, so the user can confidently own or license the output, and original artists get their **"data dividends"** - a fair share of the AI's profits, akin to how an inspiration should be compensated.

Another facet is providing **clear tiers of rights for different users**. Soundverse offers tiered licensing options, from royalty-free personal use to more extensive rights for enterprise clients. An enterprise user can even secure full ownership rights of an AI track (for example, a brand commissioning a unique piece of music). In those cases, presumably the enterprise is paying a higher fee which covers buying out the necessary rights - again, with a portion flowing back to the training data creators. By having these options built in, the platform removes uncertainty. Compare this to other companies, whose marketplace had **vague revenue splits and fine-print licensing terms** that left users unsure of what they can or can't do with the music. Lack of clarity discourages serious creators from using the tool. Ethical design dictates that a user knows: "If I make a song with this AI, I have the right to release it on Spotify / use it in my video / sell it - and here are the conditions." Soundverse's approach of a **"built-in licensing marketplace"** with automated attribution tracking exemplifies the transparency artists and content creators want. Artists who contributed data also benefit from this transparency: they know exactly how they'll be paid (e.g. a certain split or rate when their contributions are used) rather than having to guess or audit the company.

It's interesting to note the variety of solutions other players have for monetization, as each reflects their attribution capabilities (or lack thereof). Some, like **SourceAudio's SongLab** or **Eleven Music**, avoid ongoing royalty complexities by negotiating **upfront license fees** for training data with strict usage constraints. They basically pre-clear the training data for specific uses and build the compensation into that licensing deal, rather than tracking every output. This provides stability - the rights holders were paid upfront and know the boundaries (e.g. the AI can use the data to create new music for certain clients, but not to outright replace commercial releases of the original tracks). The downside is that if an AI output becomes wildly successful, the original artists don't directly share in that upside (they only got the upfront fee). It's a safer, but more static model. On the other end, a platform like **Lemonaide** initially tried a very granular approach (one producer per AI model, sharing exactly 5% of revenue if a user's track blows up past 1M streams). That is extremely fair to that one producer but doesn't scale to multi-artist datasets, which is why they moved to cohort-based splits. **LANDR's pro-rata** 20% pool is another approach: it's simple - all contributors split 20% based on how many songs each added to the training data. That means if you contributed 100 songs out of a 1000-song dataset, you get 10% of whatever revenue the AI product makes, regardless of which songs the outputs resemble. It's straightforward but not granular; some might say it's "fair enough" and at least it's transparent. The key point is that **all these models aim to compensate creators** in some way, signaling a shift in the industry from the early "wild west" of AI to a more regulated environment. Ethical frameworks recognize that while the *methods differ*, the *objective* is to ensure human creators are not cut out of the value chain.

Soundverse's method leans toward the more **granular, usage-based side** (with automated tracking), which aligns with their long-term vision of fairness. However, they also understand the importance of reliability and **scalability**. If in any case the exact attribution is uncertain, Soundverse can default to cohort or pro-rata sharing so that no contributor is left with nothing. This hybrid ensures that even as the tech improves, artists are continuously benefiting. From the label perspective, this is very powerful: it means a label can allow its catalog to be used in AI and have confidence that whenever an AI-generated song owes something to their catalog, they will see a piece of the pie. Essentially, it could turn AI from a threat into a **new revenue stream** for rights holders. In an example scenario, a music label partners with an AI platform like Soundverse, contributes, say, 10,000 songs from their library into the training set, and then any user around the world using the AI might incidentally draw from that catalog's influence. Every time that happens and an AI song is monetized, the label (and its artists) automatically earn royalties, which could be tracked in real time. This beats the adversarial approach of lawsuits and takedowns; it's a system of **legal, opt-in monetization**.

Another element to consider is **licensing for end-users**. Soundverse and others allow users to **sell or license their AI-generated tracks** through the platform. This effectively creates an AI music marketplace. A YouTuber, for instance, could generate background music and then directly license it to other creators (similar to how stock music libraries operate), with all rights cleared. Because the original contributors are in the loop (getting their share from the initial creation), it prevents later conflicts like "Hey, that background track actually sounded like my composition." Everyone relevant has been accounted for financially from the start.

Finally, transparency is crucial at this stage too. Ethical frameworks favor **proper crediting** of AI music. This might mean that metadata for an AI-generated song could list "AI Model X" or the platform as a contributor, and possibly even credit the significant influences (if identified). While not common yet, one can imagine a future where AI songs come with an *Attribution Report* - e.g. "This track was AI-generated using Soundverse; it was influenced 10% by Song A, 5% by Song B, etc., and these artists have been compensated." This kind of transparency would go a long way to normalize AI content in the industry, showing it's not violating rights but rather paying homage in a measurable way. In line with that, **regulators may encourage labeling** AI content, so having an infrastructure of attribution and credit makes compliance easier if new rules come in.

To sum up Stage 4, the ethical framework ensures that when AI music leaves the studio and enters commerce:

- All **legal rights are addressed** (no hidden copyright bombs waiting to explode; the training data was licensed and the output is either copyrightable by a human or used under clear terms).
- **Revenue sharing is automatic** and fair, leveraging the attribution methods from Stage 2 to pay those who deserve credit (soundverse.ai).
- **Users have clarity and confidence** in what they can do with the AI music (license types, ownership details are provided upfront).
- **Creators (artists/labels) have dashboards and insight** into how AI is using their works and what they earn from it, fostering trust.

An ethically distributed AI song thus becomes not a point of contention, but a **collaborative product** of human artistry and AI creativity, with benefits shared across the ecosystem. In the words of Soundverse's team, *"royalties for AI-generated music that just work"* is the ultimate goal - meaning a system where if an indie musician's old track helps train an AI hit, they simply see the royalty in their account, without having to chase anyone down or litigate. That is a vision of an AI-powered music industry that augments incomes and diversity of content, rather than cannibalizing them.

Translating this vision into practice requires a set of robust, interoperable technical measures. These measures ensure that attribution, licensing, and provenance information remain inseparably tied to every AI-generated song as it moves through platforms and formats. Only with such infrastructure can the promise of fair, automatic royalty flows be realized at scale. Some solutions in the market tackling this problem are:

- **Watermarking.** Embed robust, inaudible marks and/or **C2PA** provenance manifests at render time to bind creator, model, and license data. Robust audio watermarks are emerging (e.g., SynthID for Lyria/YouTube) and claim resilience to compression and tempo changes, but ecosystem support is still maturing. The Verge, Google DeepMind
- **License Tag preservation.** Ensure metadata survives format conversions and distribution via C2PA Content Credentials. C2PA

3.5 External audio

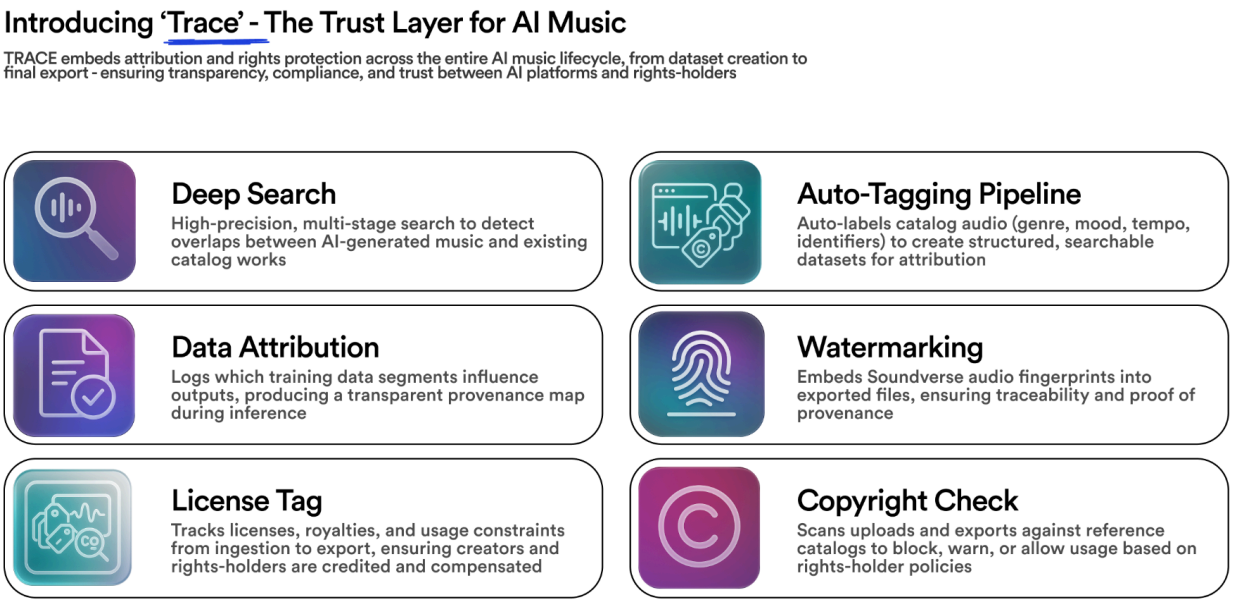
This stage ensures that Soundverse outputs remain safe and compliant in a broader musical landscape by comparing them to known catalogs and copyrighted works.

- **Deep Search** – Uses high-precision, multi-stage search to detect overlaps with existing works. This is essential for preventing unintentional copying, plagiarism claims, or legal exposure.
- **Copyright Check** – A final compliance layer that scans generated audio against reference databases of copyrighted content. Any flagged matches can be blocked, re-routed for rights clearance, or reported to the user before distribution.

Soundverse proposes following solutions:

- **Trace Deep Search.** High-precision, multi-depth search for overlaps with copyrighted catalogs, enabling **allow/warn/block** actions and royalty routing.
- **Copyright check at scale.** Industry-wide scanning endpoints for DSPs, UGC platforms, and distributors (analogous in spirit to YouTube Content ID's fingerprinting, but generalized to motifs, melodies, stems, and sections). (Columbia Electrical Engineering, Xposure Music)

3.5.1 Trace by Soundverse



'Trace' – The Trust Layer for AI Music

Trace embeds attribution and rights protection across the entire AI music lifecycle, from dataset creation to final export - ensuring transparency, compliance, and trust between AI platforms and rights-holders.

1. **Deep Search**
   High-precision, multi-stage search to detect overlaps between AI-generated music and existing catalog works.
2. **Auto-Tagging Pipeline**
   Auto-labels catalog audio (genre, mood, tempo, identifiers) to create structured, searchable datasets for attribution.
3. **Data Attribution**
   Logs which training data segments influence outputs, producing a transparent provenance map during inference.
4. **Watermarking**
   Embeds Soundverse audio fingerprints into exported files, ensuring traceability and proof of provenance.
5. **License Tag**
   Tracks licenses, royalties, and usage constraints from ingestion to export, ensuring creators and rights-holders are credited and compensated.
6. **Copyright Check**
   Scans uploads and exports against reference catalogs to block, warn, or allow usage based on rights-holder policies.

Trace gives rights-holders **control and visibility** over catalog. It continuously tracks appearances of your works **in the wild** via watermarking and multi-stage deep search. When potential infringement is detected, you receive **automated, evidence-ready reports** (timestamps, confidence, snippets, source) so you can act quickly and decisively.

3.6 Compensations

This final stage closes the feedback loop, rewarding rights-holders fairly and transparently. It transforms Soundverse from a “black-box AI generator” into a trusted ecosystem where creators, labels, and dataset owners are properly incentivized.

- **Usage-Based Royalties** – Uses data-attribution tracking to pay per generation based on actual influence. This makes payouts more precise and fair compared to blanket deals or profit-pools, aligning economic incentives with real creative contribution.
- **Tiered Licensing** – Allows different levels of participation: from training-only contributions, to style models, to likeness-based voice cloning, with escalating payouts for higher-value rights.
- **Real-Time Dashboards** – Gives rights-holders a live view of usage, earnings, and instant payouts, fostering trust and making Soundverse attractive as a licensing partner.
- **Opt-In Marketplace** – Lets rights-holders set their own price points, define license terms, and participate in a scalable marketplace that encourages collaboration rather than conflict.

While there can be many ways of solving compensation, we propose **The Partner Program** powered by data attribution that implements compensation by letting rights-holders onboard catalogs under configurable tiers (training-only, style models, likeness/voice, etc.), set permitted uses, and receive recurring payouts tied to actual usage across generations and exports. One-time fees misprice uncertain future value, sever creators from the ongoing activity their catalogs enable, and can’t adapt when demand shifts, often leading to disputes and underpayment. A recurring model aligns incentives: payouts scale with real use over time, higher-value rights earn more via tiers, and live dashboards provide settlement transparency. The result is a sustainable, opt-in system that rewards continued participation instead of a one-off sale.

3.6.1 The Solution: Soundverse Content Partner Program

Soundverse proposes the **influence-based recurring royalty model**. This program operationalizes attribution into a **transparent payout system**, ensuring that creators remain long-term beneficiaries of AI-generated music.

3.6.1.1 Why influence-based attribution (not Deep Search) fits the Partner Program

The Partner Program is designed for **causal credit**, not surface resemblance. Deep search is great for **correlation** (spotting sound-alikes, catching memorization, identifying candidates for infringement) but it over-includes common tropes and generic genre cues-bad inputs for payouts. Influence-based attribution aims to estimate **which specific catalog items materially shaped an output** and by how much, making thresholds, tiers, and settlements meaningful. Paying by similarity would be like **giving royalties to every pizza chef just because you made a pizza**; it doesn’t track who actually taught you to cook. Royalties should go to the chef you **learned from** (causal influence), not to a thousand others whose pizzas merely **look similar**.

3.6.1.2 Here are the stages of a Content Partner Program

Stage 1: Catalog Ingestion & Auto-Tagging

- **Partner Portal Access**: Once approved, partners get a secure portal for catalog ingestion.
- **Bulk Upload**: Supports thousands of tracks, with metadata, annotations, and lyrics.
- **Auto-Tagging Pipeline**:
    - Checks for copyright conflicts.
    - Tags each file with **genre, mood, tempo, timbre, instrumentation, motifs**.
    - Enriches discoverability and precision in attribution.
- **Optional Enhancements**: Partners can add lyrics, production notes, or annotations to improve attribution quality.

Stage 2: Attribution and Influence Calculation

Imagine a foundational AI model generates music:

- The proposed system decomposes outputs into their **influence components**.
- Example: A Lo-fi track with piano influences may map to 5 sources:
    - Track #896 = 51.31%
    - Track #34678 = 23.56%
    - Track #91782 = 17.80%
    - Track #127843 = 6.28%
    - Track #1987567 = 1.05%
- Only influences ≥ **x% threshold** and **max y songs** qualify for royalty to bring simplicity to the system and avoid unnecessary costs due to the computation of long-tail influences with no material impact on the generated track. System designers are free to change threshold and maximum tracks depending on their budgeting and algorithm tuning. x% and y can be decided by the system designers.
- Each company/system must calibrate thresholds to block or review, and that it might vary by location, type of content, sensitivity of model/ catalogue etc.

This ensures fair distribution without fractional noise, while still capturing meaningful contributions.

Stage 3: Attribution Points System

- Each influence percentage is converted into **attribution points**.
- **Partner Tier** (1-6) multiplies the points earned. This rewards rights-holders who grant broader or deeper rights with **higher long-term value**.

Why are Attribution Points needed?

- **Normalize noisy influence into a stable unit.** Raw per-generation influence %s are volatile and model-version–dependent. Points convert those momentary signals into an **additive, comparable credit unit** across outputs, time periods, and model updates.
- **Decouple science from finance.** Points let you evolve attribution algorithms, thresholds, and models without rewriting past payouts. Finance maps **points to pool share** each period; R&D can keep improving attribution.
- **Enable clear policy levers.** Tiers, caps/floors rules are simple to apply to points. This keeps budgets predictable and preserves fairness (e.g., avoiding long-tail micro-noise).

- **Reduce operational overhead.** Instead of millions of nano-payments per output, partners accrue points and settle quarterly - **auditable, lightweight, and reproducible**.
- **Harden against gaming.** Thresholds + max 5 tracks + points aggregation blunt spam strategies (e.g., flooding low-quality data to harvest tiny %s).
- **Protect model IP while staying transparent.** Reporting points and totals gives evidence without exposing sensitive model internals.

What happens if we remove Attribution Points?

- **Payout fragmentation & volatility.** Directly turning every influence % into cash creates millions of tiny, version-sensitive entries that are hard to reconcile, easy to dispute.
- **Budget unpredictability.** No stable unit to govern caps/tiers means finance can't set reliable pools; settlements swing with model tweaks.
- **Audit and legal friction.** Regulators and partners get sprawling, inconsistent ledgers tied to specific model states - **high dispute risk**.
- **Incentives misalign.** Without thresholds + point aggregation, long-tail micro-influences siphon value; spammers are rewarded, quality catalogs are diluted.
- **Innovation slowdown.** Any change to attribution math can retroactively perturb payouts, discouraging model improvements.

Attribution Multiplier by Tier:

- **Tier 1:** ×2.5
- **Tier 2:** ×1.8
- **Tier 3:** ×1.5
- **Tier 4:** ×1.0
- **Tier 5:** ×0.8
- **Tier 6:** ×0.5

Disclaimer: the attribution multiplier numbers need more investigation and refining before finalising. Soundverse has internally debated these numbers a lot depending on the value certain licenses bring to the table, and chose these numbers.

Why tier multipliers are needed

- **Price the scope of rights.** Different tiers grant very different permissions (training-only vs. likeness, sync, API, external commercial use). Multipliers ensure payouts reflect the **option value and legal exposure** of those broader rights.
- **Incentivize participation where it matters.** Higher multipliers reward partners who unlock higher-impact, higher-risk use cases, preventing "free-riding" by low-rights catalogs.
- **Keep finance predictable.** A fixed multiplier ladder lets you control pool spend while still differentiating value; payouts = *(points × multiplier)*—simple, auditable.
- **Avoid adverse selection.** Without multipliers, high-value rights gravitate elsewhere; with them, top-tier catalogs have a reason to opt in.

Formula (conceptual):

$$\text{Partner share} = \frac{\sum(\text{points}\times\text{multiplier})_{partner}}{\sum(\text{points}\times\text{multiplier})_{all}} \times \text{Royalty Pool}$$

Where do ×2.5 / ×1.8 come from?

They're **calibrated weights**, not arbitrary numbers - set by:

1. **Pilot/backtests** on historical usage to keep total multiplier-weighted points within the pool budget,
2. **Rights value modeling** (uplift per added permission: likeness, sync, external commercial, API), and
3. **Market signaling** (clear separation between tiers without blowing the pool).

The current ladder (2.5 / 1.8 / 1.5 / 1.0 / 0.8 / 0.5) creates:

- ~**2.5×** uplift for Full Spectrum vs. baseline (captures the big jump in rights & risk),
- a **meaningful but moderate** uplift for Tier 2 (**×1.8**) where voice+music likeness are included but with slightly narrower scope,
- **smooth steps** below that to avoid cliffs and gaming.

These coefficients are **governance knobs**: reviewed periodically against real usage mix, dispute rates, and pool stability, and adjusted if needed to keep incentives aligned and payouts predictable.

Stage 4: Royalty Pool Distribution

- **Profit Allocation**: Each quarter, AI companies can allocate **30% of profits** generated from its consumer subscription business back into a global royalty pool. Additionally, 50% revenue coming from DNA export payouts, is also sent directly to rights-holders.
- **Distribution Mechanism**: Attribution points across all partners are tallied. Payouts are distributed **proportionally**.
- **Example**:
    - Royalty Pool = $3M.
    - Total attribution points = 3M ($ = 1 attribution point).
    - Jack's Track #1 (Tier 1) earns 500 points > Jack earns 500 x 2.5 i.e. **$1250** for that quarter.
    - Jack's additional tracks may further increase his share.

This creates a **recurring, compounding income stream**: the more a track influences, the more it earns.

Stage 5: Partner Dashboard & Transparency

- **Live Attribution Reports**: See exactly when and where your catalog influenced a generation.
- **Royalty Calculator**: Breaks down percentages, points, and payments.
- **Quarterly Payouts**: Minimum threshold = $10. Frequency scales up with ecosystem growth.
- **Historical Insights**: Artists can track influence over time, enabling them to identify trends and future opportunities.

Stage 6: Partner Tiers (1-6)

To balance flexibility and fairness, we propose **six partner tiers**. Each tier reflects the **scope of rights granted by rights-holder** and **level of participation** in AI ecosystems.

Tier 1 - Full Spectrum Partnership (Highest Value)

- **Scope of Rights**:
    - AI generations allowed (download & transformation).
    - Commercial & external use permitted.
    - Royalty-free licensing for Soundverse ecosystem.
    - Monetization in **general use, sync licensing, samples**.
    - **Voice likeness + Music likeness licensing**.
    - API access for enterprise integrations.
- **Benefits**:
    - Priority support.
    - Exclusive projects & collaborations.
- **Attribution Multiplier**: × 2.5 (highest).

Tier 2 - Comprehensive Commercial Use

- AI generations allowed (download & transformation).
- Commercial external use permitted.
- Royalty-free within the ecosystem.
- Monetization in general use, sync licensing, samples.
- **Voice likeness + Music likeness included**.
- Exclusive projects & priority support.
- **Attribution Multiplier**: ×1.8.

Tier 3 - Mid-Level Commercial Licensing

- AI generations allowed (download & transformation).
- Commercial external use permitted.
- Royalty-free within the ecosystem.
- Monetization in general use & samples.
- **Music likeness allowed** (voice likeness excluded).
- Exclusive collaborations possible.
- **Attribution Multiplier**: × 1.5.

Tier 4 - Standard Commercial Participation

- AI generations allowed (download & transformation).
- Commercial external use permitted.
- Royalty-free within the ecosystem.
- Sample licensing permitted.
- **No likeness licensing**.
- **Attribution Multiplier**: × 1.0.

Tier 5 - Limited Commercial Rights

- AI generations allowed (download & transformation).
- Royalty-free within the ecosystem.

- **Non-commercial sample use only**.
- **Attribution Multiplier**: × 0.8.

Tier 6 - Non-Commercial Partnership (Entry Level)

- AI generations allowed (download & transformation).
- Strictly **non-commercial use**.
- **Attribution Multiplier**: × 0.5 (lowest).

3.6.1.3 Why This Model Matters

- **For Artists & Songwriters**: Provides **recurring income**, ensuring that every AI generation referencing their content or their audio DNA yields compensation.
- **For Labels & Publishers**: Offers a **tiered framework** to monetize catalogs strategically - aligning licensing rights with revenue outcomes.
- **For Regulators**: Demonstrates a **transparent, enforceable framework** where payments are traceable, tied to documented attribution.
- **For the Industry**: Moves beyond exploitative one-time buyouts to a **living royalty economy**, positioning AI as a tool for *shared prosperity*.

3.6.1.4 Proposal to the Industry

The Soundverse Partner Program is not just an internal system - it is a **blueprint for recurring royalties in the AI era**. It transforms attribution into structured, auditable payouts, ensuring creators remain **perpetual stakeholders** rather than **disposable inputs**.

By combining **tiered licensing flexibility** with **influence-based attribution**, Soundverse sets the foundation for a **trustworthy, recurring royalty economy** in generative AI music.

3.6.2 Learnings from the April 2024 Pilot of Content Partner Program (50 Creators)

3.6.2.1 Onboarding & Data Ingestion

- **Learning:** Bulk upload pipelines need to handle diverse file formats (WAV, MP3, stems). Majority of the creators had little to no metadata.
- **Outcome:** Improved **auto-tagging pipeline** for genre, timbre, and mood classification; simplified partner portal UX.

3.6.2.2 Metadata Quality Matters for Attribution Accuracy

- **Learning:** Tracks with detailed annotations (lyrics, production notes, stems) were more **precisely attributable** in AI generations.

- **Outcome:** Soundverse encouraged richer metadata submissions and added incentives (bonus attribution points for well-tagged datasets).

3.6.2.3 Quantity Matters, Audio Quality Matters More

- **Learning:** A brand-new bedroom producer may submit a large volume of tracks through the partner program, but if the audio quality is poor, much of that data is not valuable for model training. More data does not always equal better data. Just as not every track on Spotify gets plays, not every submitted track will be used or compensated - in AI training.
- **Outcome:** Soundverse recommends stricter data quality filters in its ingestion pipeline, making it clear that low-quality or unusable submissions receive little to no compensation. At the same time, Soundverse is considering a tiered evaluation system that explains to creators why certain tracks are excluded, reinforcing transparency while encouraging quality over sheer volume.

3.6.2.4 Attribution Thresholds Calibration

- **Learning 1:** Setting the **x% influence threshold** was validated. Without it, too many micro-contributions cluttered payout calculations. With it, the system stayed fair and computationally efficient.
- **Learning 2:** We set a max number of payout to top y tracks that had above x% influence.
- **Learning 3:** A straight up simple number such as x% influence, didn't give us insight into why they are similar. This opened up further areas of research for us. This is why Trace's multi-stage search was introduced.
- **Outcome:** Reinforced threshold policy; allowed creators to clearly see which tracks *meaningfully* influenced outputs.

3.6.2.5 Royalty Distribution & Transparency

- **Learning:** Creators wanted near **real-time visibility**, not just quarterly payouts. Many asked for a "live earnings estimate" to understand usage patterns.
- **Outcome:** Partner dashboard evolved to include **live attribution reports**, showing ongoing influence even before payout cycles.

3.6.2.6 Tier System Validation

- **Learning:** Higher-tier partners (granting broader rights like likeness or sync use) saw significantly more value per attribution point. This validated the **tiered royalty multiplier model**.
- **Outcome:** Expanded communication about the value of moving "up tiers" - labels and publishers showed interest in Tier 1-2, while indie artists preferred Tier 4-6 for simplicity.

3.6.2.7 Creator Trust & Buy-In

- **Learning:** Creators expressed **strong preference for recurring royalties** vs. one-time buyouts. Many compared it to traditional PROs (ASCAP, BMI, PRS) but appreciated the **granularity and transparency** unique to Soundverse.

- **Outcome:** Confirmed positioning: Soundverse should brand itself as the **“Spotify for AI royalties”** - a living, growing payout system instead of static licensing.

3.6.2.8 Economic Insight

- **Learning:** A small subset of tracks disproportionately influenced generations (popular genres like lo-fi, trap, cinematic scores).
- **Outcome:** Helped shape **catalog acquisition strategy** - focus on high-demand genres first to maximize partner payouts and ecosystem fairness.

3.6.2.9 Regulatory & Industry Signaling

- **Learning:** Labels, publishers, and regulators saw the pilot as a **proof-of-concept** for ethical AI monetization. Early discussions revealed that **data transparency + recurring payout** could become an industry standard.
- **Outcome:** Strengthened Soundverse’s regulatory narrative: unlike Suno/Udio, Soundverse can prove *exactly* how AI models compensate original creators.

3.6.3 Summary of Pilot Learnings

The pilot demonstrated:

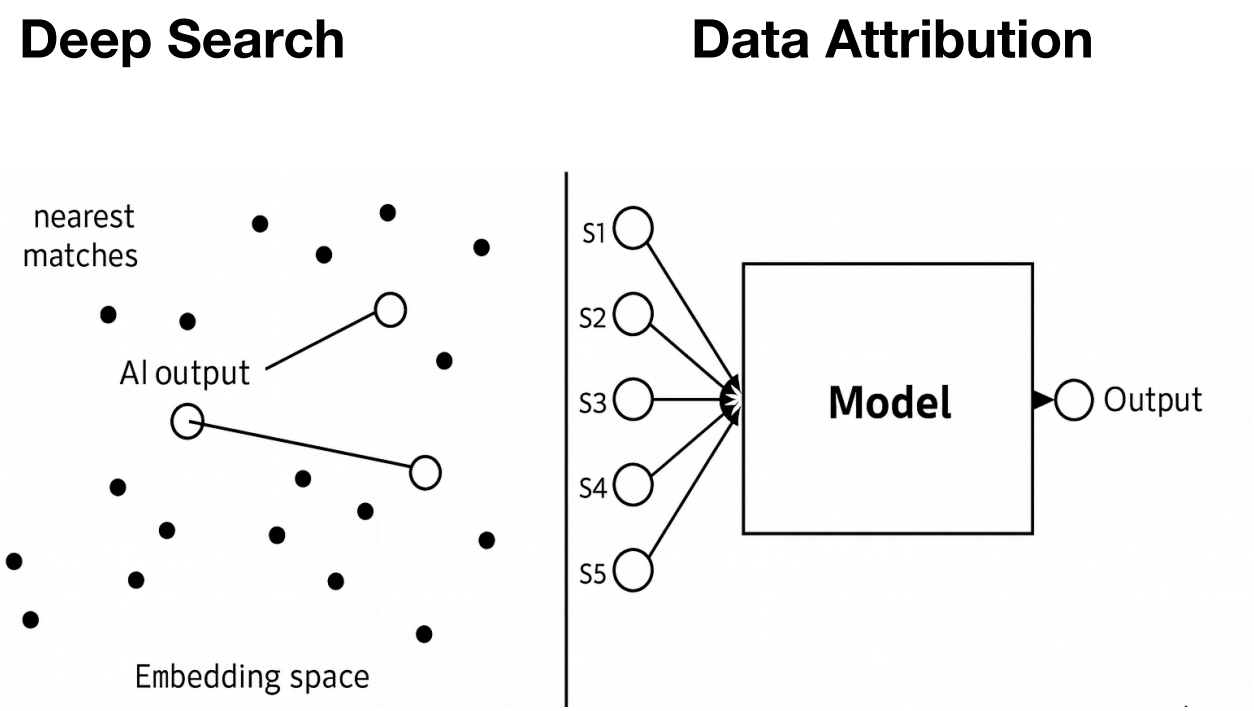
- **Operational Learnings**: Need for metadata enrichment, attribution thresholds, near real-time dashboards.
- **Economic Learnings**: Validation of tier multipliers, uneven influence distribution across genres.
- **Strategic Learnings**: Strong creator preference for recurring royalties; regulatory interest in transparency.

4. Comparison of Attribution Approaches: Data Attribution vs. Deep Search

A recurring theme throughout these stages is the challenge of **attribution** - linking AI outputs back to inputs - and the different methods to achieve it. It’s worth pausing to explicitly compare the two primary approaches we’ve discussed: **data attribution (influence-based)** versus **deep search (output-based)**, outlining their pros, cons, and appropriate uses. These are not mutually exclusive; they can complement each other, but they represent different philosophies and technical strategies.

4.1 Data Attribution (Influence Function Approach)

- This method attempts to assess causation by examining the AI model's inner workings or training process. Ideally, it answers: *"Which specific training data caused this output, and to what extent?"*



Pros:

- **Causal Insight:** It aims to provide a true explanation of influence. If implemented successfully, this is more than just saying "X and Y sound alike" - it can reveal that "Track X in the training data directly contributed 7% of the model's decision-making for this output". This kind of insight is the most **fair** basis for rewarding artists, as it reflects actual usage of their work in generating the new song.
- **Direct Compensation Link:** Because it identifies contributions at the model level, it enables a precise revenue-sharing model (e.g. pay out based on those influence percentages) which feels just and can scale to many contributors, each getting their rightful micro-share. It's akin to crediting writers for their exact contributions in a songwriting camp.
- **Not Limited to Similar Sound:** Data attribution might catch influences that are not obvious to the ear. For example, an AI might have internalized a rhythm pattern from a song even if the output doesn't sound superficially similar. Influence metrics could detect that under-the-hood impact, whereas an external similarity measure might miss it if the output doesn't outwardly resemble the source. This speaks to **invisible influence** - reward the inspiration even if it's subtle.

Cons:

- **High Complexity & Integration:** Implementing influence-based attribution requires deep integration with the AI model's architecture and training process. As discussed, it can demand retraining models or computing complex gradients for each output. Each AI model (MusicGen vs. a Diffusion based vs. an LLM-based music model) might need a custom approach, so it doesn't generalize easily. This makes it **time-intensive and expensive** to develop for every new model or update. We can't just bolt it on; it has to be designed into the system from the ground up, which is why no company/ research lab has fully achieved it so far with high accuracy.
- **Scaling Issues:** Because true influence calculation is impractical at large scales, current implementations use approximations and heuristics, which may not be perfectly reliable. There's a risk of either over- or under-estimating a song's influence. Also, if a model is retrained or new data added, past attribution might change, making it complex to maintain a consistent system. In contrast, deep search can be rerun on demand for any output without reengineering the model.
- **Limited Portability:** As hinted, an attribution method developed for one AI model may not work for another. If the industry has many AI music models, a label would prefer a universal tool to monitor them, rather than integrating into each model's proprietary system. Data attribution is largely **model-specific** - for example, Soundverse's patent-pending tech is unique to Soundverse's platform. It can't yet be applied to, say, an open-source model running elsewhere. This limits broader ecosystem adoption unless standards emerge.
- **Overwhelm of Data:** If done at a very granular level, attribution could yield extremely fragmented results - e.g. dozens or hundreds of songs with tiny influence percentages. Distributing meaningful royalties in such a scenario can be complex, and many recipients might get only pennies (which can be demotivating or practically pointless to process). There's an argument that beyond a certain granularity, it's better to pool (hence cohort approaches).
- **Easily Outdated**: If a custom attribution solution is built for a model, and if the model architecture itself goes through an evolution - the existing attribution system may not work with the new model architecture at all, and sometimes may even require additional research to support that component. Since this is an evolving field, our internal belief is that it will take some more time for this to fully mature with high accuracy. Additionally, depending on the complexity of the model's architecture this is tougher to solve- however not impossible.

4.2 Deep Search (Embedding/Output-Based Approach)

- This method treats attribution as a problem of finding resemblance between the AI-generated output and existing songs. It answers: *"Which existing track(s) does this AI output most closely resemble?"* - essentially using **similarity as a proxy for influence** (waterandmusic.com).

Pros:

- **Easier Integration & Compatibility:** Deep search operates on the *output*, so it doesn't matter how the AI generated it. You can use standard audio analysis techniques (audio fingerprinting, embeddings, spectral comparisons) across any system. This makes it easy to plug into existing workflows. For instance, a music label could run a similarity scanner on all new AI tracks released to see if any are too close to their catalog, without needing the AI developer's cooperation. In an ecosystem sense, it's more universally applicable.

- **Leverages Mature Technology:** Audio similarity detection is not new - services like Shazam, Audible Magic, and YouTube Content ID have years of development. Modern embedding techniques (turning audio into vectors in high-dimensional space) allow precise similarity calculations. These can handle large databases efficiently. Thus, an AI company can implement a similarity-based attribution with relatively

off-the-shelf components or APIs. It's far less R&D effort than inventing influence functions for deep models.

- **Transparent to Stakeholders:** The concept of "this AI song has a 90% similarity to song X" is straightforward for rights holders to grasp, much like a plagiarism report. It doesn't require understanding neural networks. This can make negotiations easier - labels might be satisfied with a system that at least flags high resemblances, as it aligns with how infringement cases are typically argued (by showing substantial similarity).

- **Scalable and Real-Time:** Deep search can be run in real-time or batch on many outputs without retraining models. This makes it feasible to screen every output on-the-fly. It provides a **quick safety net** for plagiarism without deep introspection of the model.

Cons:

- **Correlation, Not Causation:** The biggest drawback is that similarity does not prove that a particular training item *caused* the output. An AI track might coincidentally sound similar to a Beatles song even if the Beatles song wasn't in training (or had minimal impact). Conversely, an AI might have been heavily influenced by a track but generated something that sounds quite different. Deep search would misattribute in both cases - flagging some innocent outputs and missing some real influences. It's essentially a **black-box approach**: you treat the model as a music black box and only inspect its output, thus you gain no insight into the internal creative process. The model's learning remains opaque.

- **Potential False Positives/Negatives:** Relying on similarity thresholds can lead to contentious cases. Two songs might share a genre or common elements and thus seem similar above a threshold, but that doesn't mean one was derived from the other (it could be independent creation). If an AI makes a blues song that by nature sounds like many 12-bar blues classics, a naive similarity metric could flag many "influences." Also, creative transformations might evade similarity: an AI could take a melody but invert it or change key and tempo; a basic similarity check might not catch the nuanced relation whereas a human might.

- **One-Time Check:** Deep search usually treats attribution as a one-time mapping of an output to nearest neighbors. It might say "this output is 70% similar to Song A, 65% to Song B" etc., and you might credit those. But it doesn't account for the cumulative *influence* of many tracks or the combinatorial creativity of the AI. It tends to credit the closest match or matches, which may oversimplify the true creative mixture. It could overvalue one prominent similarity (like a chord progression) and undervalue a bunch of smaller influences. This might lead to overcompensating a few rightsholders while undercompensating others (for example, a model might have drawn rhythm from dozens of funk songs but melody from one pop song; similarity might pick up the melody match and ignore the rhythmic influence).

- **No Model Improvement**: Using similarity does nothing to improve the model's accountability or reduce infringement internally; it's a policing mechanism rather than a learning one. Influence-focused approaches, by contrast, could possibly be used to *improve* models (e.g. by identifying and removing overly influential data, or balancing training). Similarity detection just watches after the fact. It doesn't close the black box or help artists understand how the AI operates; it only provides a safety filter.

In practice, many modern ethical AI systems use a **combination**: deep search as a pragmatic immediate tool, and working towards deeper attribution for the long run. For example, Soundverse's framework could use deep search to ensure no blatant copy leaves the platform (addressing urgent legal risks in output), while simultaneously developing influence-function-based attribution to divide royalties in the fairest way behind the scenes. This combined use acknowledges that **technical limitations shape deal structures** - if fine-grained attribution isn't fully there yet, you use similarity as a proxy to craft a workable compensation scheme. Indeed, many companies have chosen "good enough" similarity or cohort methods to get started, prioritizing *trust and simplicity* (predictable outcomes) over precision. As one founder noted, *"a technically sound attribution system is worthless if industry rightsholders won't adopt it"* - implying that sometimes a less perfect but easy-to-understand solution (like splitting revenue by genre or contribution count) can build the necessary trust and buy-in first.

Pros and Cons Summary:

- *Data Attribution (Influence-Based)* - **Pro:** More causal, artistically fair, directly links model usage to reward. **Con:** Difficult to implement, model-specific, computationally heavy, may yield overly granular payouts.

- *Deep Search (Output-Based)* - **Pro:** Easier to deploy across systems, uses proven tech, intuitive to stakeholders, scalable for filtering. **Con:** Only finds correlations, doesn't truly explain model decisions, can be fooled by coincidences or miss hidden influences, doesn't integrate into model training.

4.3 Risks, trade-offs & mitigations

Topic	Deep Search	Data Attribution (influence)
Strength	Mature primitives; interpretable hits; ecosystem-portable	Causation-leaning links from outputs to training items
Limit	Correlation not causation; can miss diffuse, style-only influence	Architecture-specific; heavy R&D; limited portability today
Bias	Feature choices can over/under-weight genres/instruments : mitigate with balanced evaluation sets, label QA	Influence metrics can over-credit frequent items : normalize with exposure priors and per-cohort baselines
Cost	Depth ↑ ⇒ compute ↑; stems/DTW expensive	Training/inference hooks add overhead and require proprietary access

In an ethical AI music framework, we view deep search as a **"safety and interoperability tool,"** and data attribution as a **"fairness and insight tool."** Both have roles: similarity measures reassure labels that no direct plagiarism is occurring (and can plug into their existing infringement detection systems easily), while data attribution promises a future of precise, accountable AI creativity where every contributor's role is acknowledged. The trade-off is essentially between **short-term practicality** and **long-term idealism**. Soundverse's strategy, and that of a few pioneering startups, leans toward developing the ideal (attribution) without ignoring the practical (using similarity and

cohorts in the interim). This dual approach is likely to become the industry norm: use whatever works now to ensure no one gets a raw deal, and keep pushing the envelope so that eventually AI music can be as transparent and fair.

4.4 Legal alignment: “substantial similarity” vs “originality”

- **U.S.: Substantial similarity** focuses on whether the “ordinary observer” would regard the aesthetic appeal as the same; tests vary by circuit (Second Circuit’s *ordinary observer*; Ninth Circuit’s extrinsic/intrinsic). Access + substantial similarity establish infringement. Wikipedia, jipel.law.nyu.edu, law.lclark.edu

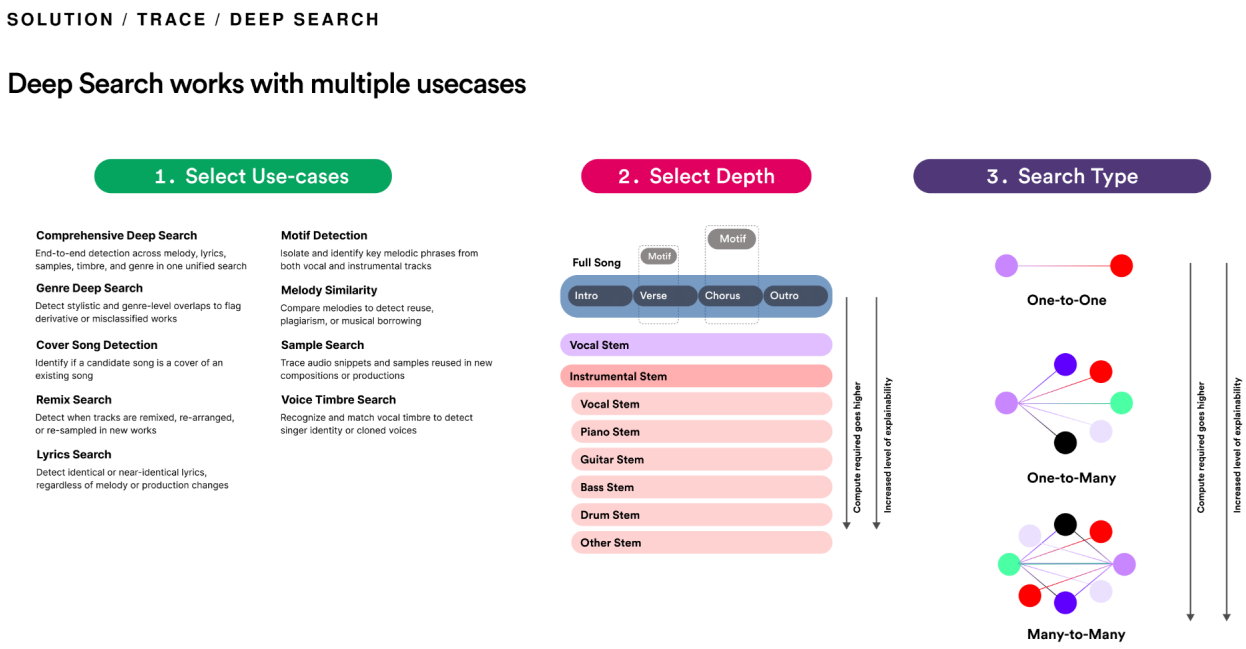
- **EU/UK: Originality** requires a work to be the *author’s own intellectual creation*, i.e., reflecting free and creative choices (CJEU *Infopaq* line). Curia, Fieldfishe. rivir.nl

Working hypothesis for Soundverse:

- **Blocking threshold (operational):** similarity score ≥ **0.xx** (x%) triggers *block and review* (to mirror or exceed “ordinary observer” risks in U.S. practice while remaining conservative in EU contexts).

- **Payout path:** where similarity is <0.xx but motif/section/stem matches demonstrate partial reuse or strong influence, route to **Trace payout** tied to contribution share. *(Exact commercial thresholds are policy, not law-Trace exposes tunable knobs so labels can align to local risk appetite.)*

5. Trace’s Deep Search - The Deep Dive



5.1 Use cases

1. **Comprehensive deep search** - unified search across melody, lyrics, timbre, stems, and structure
2. **Genre/style deep search** - stylistic overlap flags for derivative/ misattributed works
3. **Cover-song detection** - chroma/beat-sync matching, tempo/key-invariant
4. **Remix/derivative detection** - structural and motif-level reuse across edits
5. **Lyrics similarity** - near-duplicate and paraphrase detection (Deep Search + Semantic text embeddings)
6. **Motif detection** - isolate repeated phrases
7. **Melody similarity** - predominant melody extraction + contour/interval sequences; DTW alignment for time warps
8. **Sample search** - audio fingerprinting robust to noise/compression
9. **Voice timbre search** - singer/clone detection using speaker embeddings

5.2 Depth of search (Higher the accuracy, higher the cost)

Trace exposes depth controls so labels can “pay for precision where it matters.”

Depth	What it inspects	Accuracy gain	Cost/latency impact
Track-level	Whole-song embeddings	Good recall; lower precision on short borrowings	**Low**
Section-level	Intro/Verse/Chorus /Bridge	Catches localized reuse	**Medium** (4-8× per track)
Stem-level	Vocal, drums, bass, guitar, keys	Big lift on mix-masked reuse; vital for vocal motifs	**High** (2-6× depending on stems)
Motif-level	Short phrases (1-8 bars)	Highest precision on melodic borrowing	**Very high** (windowed search, many candidates)
Fingerprint-level	Micro-snippets & samples	Near-duplicate/snippet match; robust under noise	**Low-Medium**

Why cost rises with depth: moving from a single global vector to many local windows multiplies comparisons; stems require separation; DTW is quadratic without pruning (though exact/streaming variants and lower bounds help). csdl.ics.hawaii.edu, program.ismir2020.net

5.3 Search types

- **One-to-One (1 : 1).** Compare a candidate track to a suspected reference. Fastest; ideal for takedown triage.
- **One-to-Many (1 : N).** Candidate against a label catalog or global registry
- **Many-to-Many (N↔N).** Continuous fleet-wide dedup/remix/cover discovery

We operate three complementary *search modes* (1:1, 1:N, N↔N) across three *reference corpora* (A: internal training set, B: curated/partner set, C: external web-scale). This lets labels, publishers, and CMOs/PROs choose speed vs. depth per workflow, and decide *where* to look: inside the model's own diet, inside licensed partner catalogs

Reference corpora

A. Internal training dataset (closed)
Used to detect memorization. Highest value for compliance audits and model-governance reviews; not suitable for public takedowns (private corpus).

B. Curated/partner dataset (consented)
Rights-holders onboard catalogs via the Partner Program or simply share the catalog via online storage URL/

C. External web-scale (open-world)
Search across platforms (YouTube, TikTok/IG Reels, Spotify, etc.) via audio fingerprints/embeddings and public/partner registries. Note: web-scale audio indexing is largely a solved infrastructure problem; the *hard* problem is accurate, durable attribution and thresholding over time.

Mode 1. One-to-One (1:1)

Compare a **candidate** to a **suspected reference**. Fastest path for **triage** and internal legal reviews.

- **Typical use:** "Does this new AI output collide with Track X?"
- **Depth options:** global-embedding cosine; section-DTW around melody/chorus; motif-level n-gram on predominant melody; stem-level (vocal/bass/drums) if needed.
- **Corpora:** A (training) for leak checks; B for partner disputes; C for public claims.
- **Why it matters:** single-pair analysis supports quick block/allow calls and red-flag escalation.
- **Notes:** start at coarse global vector; escalate to section/motif/stem only if the global hit is near threshold to control cost.

Mode 2. One-to-Many (1:N)

Compare one **candidate** against a **catalog/registry**. This is the workhorse for **pre-release** screening and **partner safety**.

- **Typical use:** "Screen every export against our label catalog + industry registry."
- **Corpora:** B first (partner), then A (training) for leakage, then C (external) if distribution is planned.
- **Tuning:** recall-optimized at the ANN layer; final policy threshold set with counsel ("substantial similarity" risk tolerance).

Mode 3. Many-to-Many (N↔N)

Continuous, fleet-wide **dedup/remix/cover discovery** across rolling batches. Suited for **market monitoring** and **catalog hygiene**.

- **Typical use:** scan new outputs vs. recent releases and user uploads to catch cover-style or motif reuse at scale.
- **Corpora:** B↔B for partner ecosystems (e.g., multi-label exchanges), C↔C for platform-wide monitoring; A is used periodically to regress test memorization.
- **Outcome:** automatically surfaced clusters (near-dupes, remixes, chorus-matches) for human review and escalation.

Where each corpus fits operationally

- **A. Internal training set** : *Compliance & science.* Proves/denies memorization; not used for public claims.
- **B. Curated/partner set** : *Business & payouts.* Clean metadata, consented use, influence hooks, attribution points/tiers downstream.
- **C. External web-scale** : *Enforcement.* Takedown/claim flows on UGC/DSPs; requires robust fingerprint registries and ingestion pipelines.

A note on web-scale indexing

Building and operating external (C) indexes - fingerprints, embeddings, ingestion from YouTube/TikTok/Spotify - is a *separate* engineering problem that many vendors have already solved. Replicating it for AI attribution is feasible; the *hard* problem is sustained accuracy, cost control with deeper searches (section/stem/DTW), and stable thresholds that align with legal risk and royalty policy over time.

5.4 Operating thresholds and outcomes

- **Block/Review:** overall similarity ≥ **0.xx** (policy default).
- **Flag for payout:** strong *localized* matches (motif/section/stem) even if global similarity <0.xx; routed to rights workflows
- **Explainability:** Trace returns *what matched* (timecodes, sections, stems, or lyric snippets) + *how strong* (score, alignment type, features used).

5.5 Data Attribution (influence) - status and limits

- **Model-integrated logging.** Soundverse's approach-tracking token/feature flows during training/generation within its own architecture-offers a *causation-leaning* view of which training items influenced an output (not just correlation). Water & Music
- **Influence-function lineage.** Classical influence functions map predictions back to training points; they illuminate behavior but become approximate on deep non-convex nets, requiring careful validation.
- **Reality check.** Influence is an *active research area*; robust, model-agnostic attribution remains unsolved. Meanwhile, Trace deploys **Deep Search** (already production-proven primitives like fingerprinting, chroma

covers, ANN vector search) to deliver immediate value.

5.6 Watermarking & provenance

- **Constraints today.** Watermarks (e.g., SynthID) are promising for audio and video and claim resilience to compression/speed/noise, but they require ecosystem adoption and aren't yet a standalone proof of rights for adversarial cases. The Verge

- **Provenance manifests.** Adopt **C2PA Content Credentials** to bind license tags and model metadata to exported assets; this travels across edits/renders and can be validated downstream. C2PA

5.7 For labels: solving the Suno/Udio problem with Trace

Problem: Labels cannot reliably identify when AI outputs are *influenced by*, *derived from*, or *confusingly similar to* their catalogs; nor can they monetize at scale. (Active suits illustrate both harm and intent to license going forward.) RIAA, Music Business Worldwide

Trace solution path:

1. **Onboard catalogs** via the **Content Partner Program**: bulk ingest masters and compositions; bind rights and payout rules.
2. **Pre-compute multi-depth references**: track/section/stem/motif embeddings + fingerprints; maintain label-private and industry-shared indexes.
3. **Deploy 1 to N scanning** across UGC/ DSP feeds and external AI outputs.
4. **Decisioning**:
    - ≥0.xx global similarity: *block & escalate* (litigation-grade packet).
    - Localized motif/section/stem matches: *auto-payout* per label policy or *negotiate license*.
5. **Commercial loop**: return usage analytics; labels adjust thresholds; issue proactive licenses to compliant AI vendors.

Why it works now: Trace stands on battle-tested primitives-fingerprinting, cover-song chroma matching, ANN vector search-extended to motifs, stems, and lyrics.

Complement with attribution R&D: When Soundverse's architecture-integrated attribution matures, labels can *augment* deep search-based payouts with influence-based splits (cohort shares, min-influence floors).

5.9 Interoperability & standardization roadmap

- **Short-term:** publish a **Similarity Score Schema** (score, feature space, depth, alignment, confidence) and a **Rights Action Vocabulary** (allow/warn/block/payout/escalate).
- **Medium-term:** align export provenance on **C2PA**; add **Audio-specific Assertions** (model id, prompts, stem lineage). C2PA

- **Industry forum proposal:** create a **Music AI Attribution & Similarity Forum (MAASF)** with labels, PROs/CMOs, AI vendors, DSPs, and research labs to steward thresholds, test sets, and TCRs (Trusted Corpus Repositories).

5.9 Implementation guide for a rights-holder

1. Readiness

- Select catalogs and rights rules; *review*, *payout* lanes; designate legal contacts.

2. Integration

- Bulk ingest (masters + compositions + lyrics), tag normalization, rights binding with Soundverse Trace.
- Compute multi-depth references

3. Operations

- Connect Trace scanning to internal A&R review, takedown desks, and licensing teams.
- Pilot with selected DSP/UGC partners; measure against counsel-approved test sets (covers, remixes, motifs).

4. Monetization

- For non-blocking overlaps, auto-payout per policy; expose dashboards to artists/publishers.
- Offer compliant AI vendors programmatic licenses using Trace verification.

6. Open Research Questions

- **Attribution granularity:** melody/motif/stem/section labeling-best practices for combining influence and similarity scores in one payout formula.
- **Robust watermarking:** real-world survival (TikTok transcodes, broadcast chains, on-device edits). Early claims exist (e.g., SynthID), but independent benchmarks are needed. The Verge
- **Global benchmarks:** public test suites for motif/section/stem similarity across genres and languages.
- **Lyrics paraphrase:** balancing near-duplicate detection (MinHash/SimHash) with semantic paraphrase models and fair-use boundaries. eecs.harvard.edu, research.google.com

7. Conclusion and Future Outlook

AI's expansion into music is a watershed moment for the creative industries. It holds the promise of boundless creativity - songs generated in seconds, personalized soundtracks for every listener, and new genres birthed from human-AI collaboration. Yet it also challenges the very framework of how we value and protect creativity. The **Ethical AI Music Framework across all stages**, as detailed in this paper, offers a blueprint for navigating this new landscape. By addressing ethical issues at each stage - **data sourcing**, **model training**, **generation**, and **distribution** - we can transform AI from a disruptive threat into a collaborative tool that amplifies human artistry.

The framework we explored is not merely theoretical; it is already taking shape in platforms like Soundverse, and in emerging industry standards. In Stage 1, we saw that sourcing data ethically through artist partnerships and licensing deals can turn antagonism into alliance, as evidenced by initiatives from Soundverse's partner program to ElevenLabs' deals with music companies. This lays the groundwork of **consent and compensation** that makes everything downstream more tenable. In Stage 2, we tackled the hardest puzzle: attribution and transparency. While perfect explainability may still be on the horizon, the efforts by Soundverse, and others to innovate in this space indicate a future where AI is not a black box but a **glass box** - one where influence flows can be observed and audited. Importantly, we recognized that technology alone isn't the savior; building **trust** through open communication and fair interim solutions is equally vital. In Stage 3, we confronted the necessity of guiding AI's use - implementing guardrails so that AI augments rather than imitates or replaces creative talent. As the Copyright Office and major lawsuits have signaled, **unbridled mimicry or flooding the market with AI content** is not a sustainable path. The ethical framework's answer is to encode respect for artists into the generation process itself: don't let the AI or users tempt fate by doing what a musician would consider crossing the line. And finally, in Stage 4, we described how distribution can be reimagined so that AI-generated music enters the world with **all rights accounted for and all contributors rewarded**, ensuring that human creators share in the prosperity of AI creativity rather than being left behind.

For record labels, publishers, and artists, this framework provides a path to engage with AI proactively. Instead of fearing AI music as a competitor, rights holders can become stakeholders - contributing to AI training in a controlled way, and earning from every AI song that draws on their repertoire. Regulators, on their part, can look to such frameworks when crafting policies: for instance, the idea of data licensing and attribution could inform new regulations requiring AI firms to document training data provenance and to compensate creators (much like compulsory mechanical royalties in songwriting, one could envision compulsory "AI training royalties" if negotiations fail). Already, the **"Fairly Trained"** certification movement is a step toward industry self-regulation, labeling models that respect rights. We may also see collective rights organizations extending their scope to AI - one can imagine PROs or mechanical licensing agencies managing AI data contributions and payouts, using the kinds of attribution tech discussed here.

The **trade-offs** we highlighted, such as data attribution vs deep search, will continue to be active areas of development. In the coming years, improvements in AI interpretability might reduce the current heavy reliance on proxy methods. Academic research is ongoing into more efficient influence estimation, watermarking techniques that survive generative transformations, and hybrid models that incorporate metadata tracking. One promising avenue is the concept of **"model provenance records"** - where an AI model could keep an internal log or watermark of which data influenced a particular output (some early research in image generation has attempted this). If such ideas mature for music, the balance could tilt further towards precise attribution. On the other hand, if those remain elusive, the industry might standardize on simpler methods (like cohort attribution or fixed-percentage data pools) as a pragmatic solution, as we saw with Landr and Human Native focusing on straightforward licensing deals.

Innovation vs. protection is a delicate balance. The framework presented strives to encourage innovation (AI can open up creative possibilities for independent artists, for example, by giving them infinite session musicians at their fingertips or allowing small content creators to have quality music scores) while ensuring protection (artists' existing works aren't exploited without reward, and their personas can't be hijacked by deepfakes). Achieving this balance means continuous dialogue between tech developers, musicians, and lawmakers. Ethical platforms like Soundverse can serve as experimental sandboxes for policies - demonstrating that it's feasible to have AI music that's both cutting-edge and compliant with artists' rights.

In conclusion, an ethical AI music ecosystem is not anti-technology; it is **pro-artist, pro-fairness, and pro-sustainability**. As AI becomes an inseparable part of music creation (as Soundverse's CEO predicted, "music creation and AI will be inseparable in the future" (musically.com)), the practices we establish now will determine

whether that future is one of mutual benefit or conflict. By implementing frameworks across all stages of AI music creation, we can ensure that the next generation of music is created in collaboration with AI *and* with respect for the human creators who inspire it. This will ultimately lead to a richer musical landscape - where human creativity is amplified by AI, and AI is guided by human values. The line between inspiration and infringement, between innovation and appropriation, can be maintained with clear ethics. In doing so, we safeguard the soul of music even as we embrace the power of machines to extend its possibilities.

Let's shape the future of sound together - one that echoes with creativity, fairness, and respect at every stage.

8. Glossary

- **AI-Generated Music** - Music composed or produced by an artificial intelligence system. In this context, it often refers to songs or audio created by generative models (like neural networks) based on patterns learned from training data, rather than being directly written or performed by humans.
- **Training Data** - The collection of music recordings, stems, MIDI files, or other musical material used to train an AI model. The model learns statistical patterns from this data. In ethical use, training data should be licensed or authorized by the rights holders of the music.
- **Content Partner Program** - An initiative (exemplified by Soundverse) where artists, labels, and rights holders can opt in to contribute their music to an AI platform's training dataset in exchange for compensation. It formalizes the relationship between creators and the AI, ensuring artists are paid when their content helps generate new music.
- **Data Attribution** (in AI Music) - The process or technology of tracing an AI-generated output back to specific pieces of training data that influenced it. Data attribution seeks to establish a causal link (who/what in the training data contributed to the creation of a particular melody, harmony, style, etc.). It is used for crediting and rewarding those contributors.
- **Influence Function** - A concept from statistics and machine learning used to measure the influence of a single data point on a model's output. In theory, it involves comparing model performance with and without that data point. In AI music attribution, influence functions (or approximations of them) are used to estimate how much a particular song in training affected the generated music.
- **Deep Search** - In this context, a technique to find and measure similarity between an AI-generated music piece and existing songs. It often relies on **embeddings** (numerical representations of audio) to compute how "close" two pieces of music are in terms of sound or structure. It's used as a proxy for attribution or to detect potential plagiarism by flagging if an AI output sounds very similar to a known track.
- **Embedding (Audio Embedding)** - A vector (a list of numbers) that encodes a piece of audio in a multi-dimensional space such that similar-sounding audio have vectors that are close together. Embeddings allow mathematical comparison of songs. For example, an AI might embed all training songs and a generated song, and then find which training song's embedding is nearest to the generated song's embedding to gauge similarity.
- **Black Box** (AI Black Box) - A system (often a complex neural network) whose internal workings are not interpretable or transparent to the observer. When an AI model is a "black box," we can see the input and output but not clearly explain *how* it made its decisions internally. A major goal of ethical AI in music is to turn this into a more transparent or "explainable" system (a "glass box").
- **Explainable AI (XAI)** - A field of AI focused on making AI's decisions and behavior interpretable to humans. In music AI, XAI methods would help clarify why a model produced a certain tune, e.g., by identifying influential training data or musical features. Attribution technology is a form of XAI applied to generative music models.

- **Watermarking** - In the AI training context, embedding a hidden signal or marker in the training data that the AI might carry into its output. For example, adding an inaudible high-frequency tone to all training songs so that if the AI output contains that tone, one knows those songs were used in training. It's a way to trace data usage. However, watermarks can degrade or be circumvented.
- **Cohort-Based Attribution** - An approach to attribution where, instead of tracking individual songs' influence, the system groups training data into cohorts (such as genres, styles, or some category) and attributes an output's influence to one or more cohorts. All contributors in that cohort share credit or royalties. This is less granular but more scalable, ensuring groups of similar music get collective recognition when an AI output is in their style.
- **Pro-rata Splits** - A method of revenue sharing where contributors are paid proportionally to something like their share in the dataset or number of items contributed, rather than the specific usage. For instance, if an artist's songs constitute 5% of the training library, they get 5% of the royalty pool from the AI model's earnings, regardless of exactly which outputs were influenced. It's a simpler attribution alternative, used by some platforms (e.g., LANDR's 20% pool split by data volume).
- **Guardrails (in AI)** - Policies or technical filters put in place to prevent undesirable or unethical outcomes in AI generation. Examples in AI music include banning certain prompt inputs (like artist names), preventing the AI from outputting someone's exact voice or a near copy of a song, or limiting the length/quality of output to avoid misuse. Guardrails are safety mechanisms to align AI behavior with ethical and legal standards.
- **Fair Use** - A legal doctrine (in the US and some other jurisdictions) that allows limited use of copyrighted material without permission under certain conditions (like for commentary, criticism, transformative use, etc.). Suno and Udio attempted to claim their AI's training on and output of music was "transformative" and thus fair use. However, regulators and many experts are skeptical that wholesale AI training on music qualifies as fair use, especially if it causes market harm. Fair use is determined by factors including purpose, nature, amount used, and effect on the market.
- **Market Dilution Theory** - In the context of AI, the idea that even if an AI output doesn't copy a specific work, the sheer volume of AI-generated content could flood the market and reduce economic opportunities for human creators. It's essentially a macro-level harm where AI music oversupply dilutes streaming revenues, attention, or value of original music. The Copyright Office has acknowledged this as a potential concern in analyzing AI impact.
- **Rights Holder** - An individual or entity that owns legal rights to a creative work. In music, this could mean songwriters (and their music publishers) who hold composition rights, or record labels/artists who hold sound recording rights. In our discussion, rights holders are those who must be considered when using their music in AI training or when AI outputs might affect their copyrighted works.
- **Opt-in vs Opt-out** - Models for consent regarding data usage. Opt-in means data is only used if the owner explicitly gives permission (Soundverse's approach with partner programs). Opt-out means data might be used by default unless the owner explicitly requests removal (some tech companies have argued for this in AI training, but it's controversial). Ethical frameworks favor opt-in to respect creators' choices.
- **Deepfake (Voice)** - An artificial simulation of a real person's voice (or likeness) produced by AI. In music, this can mean generating singing or speech that sounds like a specific artist who never actually performed those vocals. Deepfakes raise issues of consent, authenticity, and could violate publicity or trademark rights if done without permission.
- **Royalty (Music Royalty)** - Payments made to rights holders for the use of their music. In the AI context, royalties could be due to artists whose works were used to train an AI (data royalties) or artists whose style was imitated. Soundverse's system of paying content partners when AI generates music using their data is a form of royalty payment in the AI setting.
- **License (Licensing)** - A license is permission to use a copyrighted work under defined conditions. In AI music, licensing applies both to training data (licensing music to use for AI training) and to AI outputs

(giving users licenses to use the AI-generated music). An **integrated licensing model** means the platform directly provides a license to the user for the output, and ensures underlying rights are handled.

- **DNA (Artist trained AI Models)** – An AI model trained on a specific artist's catalog that captures their stylistic essence (timbre, melody, production traits, etc.). DNAs allow artists to monetize their *style* in the AI era. They can be private (for the artist's own use) or public (published on a DNA Marketplace for fans and creators to generate in that style).
- **DNA Marketplace** – A digital storefront within Soundverse where artists publish their DNAs for public use. Users can generate music with DNAs and pay export fees (set by the artist, typically $0.99–$9.99). Revenue is shared between the artist and Soundverse, ensuring ethical monetization of AI-driven creativity.
- **DNA Export** – The process of purchasing and downloading music created using a DNA. Exports are the primary monetization event in the DNA ecosystem, triggering automatic royalty splits between the artist and Soundverse.
- **DNA Generation** – The act of creating a musical output using an artist's DNA inside the platform. Generation is powered by Soundverse tokens. Users may experiment freely in-app, but exports (downloads/clearances) require payment.
- **DNA Sensitivity Selector** – A feature that clusters an artist's catalog into multiple distinct DNAs based on stylistic eras or sonic variation. For example, an artist's "acoustic period DNA" and "electronic period DNA" can be separated automatically. The artist can merge, split, or fine-tune these groups before publishing.
- **DNA Profile** – The metadata bundle provided by an artist when publishing a DNA. It includes DNA name, description, cover art, pricing (per export), licensing rules (UGC only, sync, commercial), and attribution terms.
- **DNA Royalty Split** – The default 50/50 revenue sharing mechanism for DNA exports: 50% to the artist, 50% to Soundverse. Artists may adjust splits in cases of collaboration or label partnerships.
- **Fairly Trained** - A term emerging to describe AI models that have been trained only on data obtained ethically (with consent and/or compensation). There is even an initiative to certify models as "Fairly Trained" if they meet these criteria. It's analogous to "fair trade" in other industries - assuring consumers and partners that the product (AI model) was created without exploiting creators.
- **Audio AGI** - Stands for "Audio Artificial General Intelligence" in Soundverse's context. It refers to an advanced AI that can understand and generate music with a high degree of generality and intelligence, potentially across various tasks (composition, production, mixing, etc.). Soundverse's mention of an Audio AGI roadmap implies developing very advanced, explainable music AI capabilities. It's a forward-looking concept towards AI that might eventually match human-level versatility in music creation.
- **Stakeholders (in music AI)** - All parties who have an interest or stake in AI music. This includes AI developers/platforms, musicians, songwriters, record labels, music publishers, listeners, distributors (like streaming services), and regulators. An ethical framework tries to balance the interests of stakeholders: innovation and profit for developers, compensation and respect for creators, quality and variety for listeners, and compliance with legal standards for regulators.
- **ANN (Approximate Nearest Neighbor).** Fast indexing (e.g., FAISS, HNSW) for high-dimensional embeddings. arXiv
- **Attribution (influence).** Mapping generated outputs to specific training items that influenced them, ideally via model-integrated signals. Water & Music
- **Chroma features.** Pitch-class energy per semitone; key/tempo invariant signals for cover detection.
- **C2PA / Content Credentials.** Open standard for cryptographically verifiable media provenance.
- **DTW (Dynamic Time Warping).** Alignment method to compare time series that vary in tempo.
- **Similarity score.** Scalar in [0,1]; Trace default policy treats ≥0.xx as "block & review."
- **SynthID.** Google DeepMind watermarking for AI-generated media, including audio pilots. The Verge