Generative Music AI's $350 Millon Problem: Compensating Creators for the Use of Copyrighted Materials in Training Sets

Josh Wagman
Queen's University
21jdw14@queensu.ca

Kay Yan
Queen's University
k.yan@queensu.ca

Rafael Costa
Queen's University
22xwyk@queensu.ca

Alex Levesque
Queen's University
Alex.levesque@queensu.ca

Armita Afroushe
Queen's University
24kb11@queensu.ca

***Abstract*—The rapid expansion of music AI technologies has led to the extensive use of large-scale datasets that often include copyrighted music without adequate oversight. Current legal and technical frameworks struggle to identify and quantify such copyrighted content, resulting in the under-compensation of copyright holders and potential violations of intellectual property rights. This study implements a unique approach to copyright detection. Utilizing federated learning (FL), our method trains models locally, preserving data privacy by keeping sensitive information on local servers while aggregating model updates centrally. Additionally, model fingerprinting assigns unique digital signatures to training data outputs, enabling precise tracking and verification of copyrighted material. Leveraging these techniques, our framework compiles a comprehensive catalog of artists and quantifies the number of songs present in the dataset, which is then integrated into our compensation mechanism to ensure fair remuneration for copyright holders. Our solution enhances transparency in data usage while delivering mutual benefits for both AI developers and creators, incentivizing a cooperative musical landscape where AI and creativity coexist.**

I. Introduction

In April 2023, an unknown Tik Tok user called Ghostwriter977 released a song on Spotify and Apple Music called "Heart on My Sleeve" that would greatly influence the music industry. Generative music AI claimed the spotlight in music innovation with the release of this song, featuring Drake and The Weeknd. The only complication, however, is that neither Drake nor The Weeknd ever sang a single note for this track. This was one of the first documented instances of generative music AI being used to create music that became a major worldwide hit, and many more have come since. The AI was trained on copyrighted music, and the artists and their record label, Universal Music Group (UMG) were never compensated for the use of their music in the training of this AI. AI creators and artists in music, visual arts, and other fields face the challenge of insufficient copyright laws governing AI. This lack of law and policy leads to intense legal battles, such is the case in Suno & Udio Vs. UMG, Warner Records, and Sony Records. The use of copyrighted music to train generative music AI by Suno and Udio, while not technically against any specific laws, has led three of the world's largest music licensers to sue them. If this discrepancy is not fixed, then those working in creative professions will continue to have their works used without their consent to train AI. This will inevitably lead to a loss of jobs due to this software, the program that was trained on their own works.

With the recent developments in AI, law and policy fall further behind, AI companies seem to be able to skirt the law for their own personal, monetary gain. The datasets are not monitored and are kept private by most AI companies, meaning that there is no way for the government or creative industries to get a hold of the datasets without legal action. On top of this, the creative professionals and those who own the works are not getting compensated for the use of their works. AI developers need huge amounts of data in order to properly train their AI, and this often leads to the usage of copyrighted works. If one AI company is using copyrighted works, then would it put the other AI company at a disadvantage to not do the same, especially when there is no specific law against this use? AI is a rapidly developing field, and every company is striving for the highest quality product to offer their users.

Both sides have an argument to be made, and until legal precedent is set, these two industries are poised to fight against each other. The questions this paper aims to answer are: How should copyright law be adapted to fit generative AI and should artists and record labels be compensated for the use of their works in the training of generative AI?

II. Background Information

"Heart on My Sleeve" was not the only case of AI-generated songs making their way into the public music scene [1]. UMG and other major record labels were starting to catch on that these AI used thousands of copyrighted music to train them [2]. This act, they claim, constitutes "copyright infringement on an almost unimaginable scale" [3]. To understand this claim, we will provide you with background context into the

training and usage of generative music AI software, copyright law, and the current legal landscape of relevant cases.

A. How Generative Music AI is Trained

Generative music AI begins by collecting datasets of musical elements, such as chords, melodies, rhythms, and timbres. These datasets often consist of music files in formats like MIDI or audio recordings in WAV or MP3. MIDI files are widely used due to their ability to represent music as a sequence of notes with information about pitch and velocity, which are easier for AI models to analyze and learn from [4].

The preprocessing stage of generative music AI involves converting audio waveforms into spectrograms. These graphs visually represent the frequency content of sound over time and almost act like a "map" that AI can use to analyze patterns or specific frequencies. Commonly, Short-Time Fourier Transforms (STFT) are used to convert audio waveforms into a time-frequency representation. As a result, the spectrogram's x-axis represents time, the y-axis represents frequency, and the color intensity represents amplitude, as seen in 3. For example, a bass-heavy song with deep male voices will have stronger signals in the lower-frequency range [5].

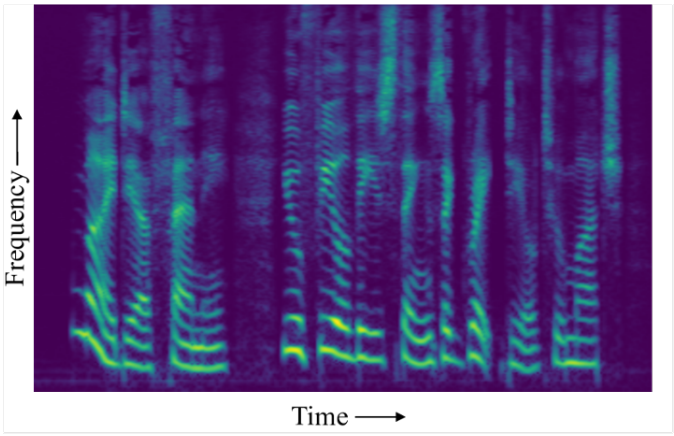

Fig. 1. Spectrogram Representation of a Four-Second Audio Signal [5])

Then, these spectrograms are used as input to convolutional neural networks (CNNs), which are specialized for recognizing spatial patterns in visual data. As spectrograms pass through a CNN, they are compressed into lower-dimensional feature vectors that represent musical characteristics like chord progressions and melodic motifs. Thus, instead of processing the entire spectrogram at once, CNNs process it frame-by-frame, generating sequential data corresponding to different time steps. These vectors retain essential information while reducing complexity, making them more suitable for training models that require sequential inputs.

Once the model is trained, the AI generates new music using algorithms that are based on the patterns it learned. Algorithms are essentially sets of rules or instructions that the AI follows to perform tasks. Once the music is made, the output can then be converted into formats for playback such as WAV or MP3.

B. Legal Landscape

The current legal landscape of AI is extremely volatile and still developing. Current cases on the court docket are bound shape the future of AI policy by setting important precedent. The current legal system is trying to determine how to adapt current copyright laws and policies to AI systems, specifically generative AI systems. To first understand the legal landscape, we will introduce you to the type of law that this paper deals with: copyright law.

*1) **Copyright Law**:* Copyright law exists to protect creative works around the world, but what exactly is defined under the scope of "creative work"? Creative works aim to define the expression of creative ideas but not the systems/processes used to derive these ideas: "Copyright protects expression, and never ideas, procedures, methods, systems, processes, concepts, principles, or discoveries" [6]. As such, copyright protects an instance of a creative work, for example, a song, painting, or distinct brand logos, but never to protect a process of creating a work. For example, Lord of the Rings is a copyrighted work by author J.R.R. Tolkien, however the brainstorming methods and writing techniques that Tolkien used are not protected under copyright. There are exemptions in which copyrighted works are allowed to be used without expressed consent from the copyright owner, and this is a staple of copyright law known as the Fair Use Doctrine. The Fair Use Doctrine promotes freedom of expression by allowing the use of copyrighted materials under specific circumstances. Copyrighted works may be used without consent in cases of criticism, teaching, research, and news reporting, to name a few [7]. The criteria in which Fair Use is evaluated are:

1. **Purpose and usage of the copyrighted work**: Copyrighted works used in a non-profit educational and non-commercial manner are much more likely to be classified as fair. Similarly, transformative uses (adding something new or changing something with a further purpose) are more likely to be classified as fair [7].
2. **Nature of copyrighted work**: Many types of works can be copyrighted, each with varying levels of creativity. Using works that are considered more creative such as songs, movies, and art are less likely to be classified as fair than factual work, as this relates to the premise of encouraging creative expression in copyright law [7].
3. **Proportional amount of copyright used**: If the use of copyright materials is found to have copied a large portion of the copyrighted work, then it is less likely to be classified as fair [7].
4. **Effect of the use upon the potential market for or value of the copyrighted work**: If the use of copyrighted works is competing in the same market as the original work, then courts must consider the potential effect to the market upon widespread use of the created work.

A key distinction to be made in the writing of these metrics is that every use must be evaluated on a case-to-case basis. The wording implies that even though specific uses are "more likely" or "less likely" to be considered fair use, each case is unique and requires specific interpretation of the law [7].

A central challenge in the current system is that standardized licensing agreements for AI-generated music is lacking. There is no clear framework for licensing rights, which given AI creators and user challenges in legally using or distributing AI-generated works. To add to that, the application of fair use to generative AI music is uncertain, especially regarding whether AI-generated compositions are "transformative" enough to qualify as fair use. Given that these models often are trained with extensive use of copyrighted materials, courts may have difficulty consistently applying fair use to these cases, as we will show in the two cases coming up: Suno-Udio and Stability AI [5] [8].

2) Relevant Cases: The legal landscape is bound to shape the future of generative music AI. As technology makes further progress, policy struggles to catch up. Current court cases will ultimately determine the future of the generative AI field, and as such, it is vital to understand past and present court proceedings. In this section, we aim to present what we believe to be the two most important cases, ongoing and present to paint a picture of the current legal landscape surrounding generative AI.

Case 1: Suno and Udio v. UMG, Sony Music Entertainment, Warner Records, et al. (2024)

The most relevant case to the discussion of copyright law regarding generative AI music is an ongoing case, between generative AI music companies Suno and Udio, and record label titans Universal, Warner, and Sony. First filed on June 24th, 2024, the three labels filed federal lawsuits in New York and Massachusetts against the two startups. They allege that Suno and Udio were involved in mass copyright infringement by using popular songs, such as The Temptations' "My Girl", Mariah Carey's "All I Want for Christmas Is You," and James Brown's "I Got You (I Feel Good)", to train their AI models. These models would then be able to create music on demand and can mimic iconic artists such as Michael Jackson, Bruce Springsteen, and ABBA [1].

The defense of the startups was that they allege that their AI training systems fall under "fair use," which would permit them limited use of copyrighted works without authorization. Suno's CEO, Mikey Shulman, emphasized the transformative nature of their technology, arguing that it generates "completely new outputs" rather than simply "regurgitating" existing songs. They also assert that their systems analyze patterns in music, rather than memorizing specific content, which allows users to create original music based on text prompts [9].

The record labels on the other hand, argue that this use of their music is unlicensed and amounts to willful infringement, potentially resulting in AI-generated songs that "cheapen" the original works by offering near-identical imitations. The labels are particularly concerned about the AI's ability to reproduce specific musical elements and even simulate artist-specific vocal styles. Mitch Glazier, CEO of the Recording Industry Association of America, criticized unlicensed services like Suno and Udio, arguing that they are exploiting artists' work without fair compensation, which he argues could hinder genuine innovation in AI [9].

The labels are seeking statutory damages of up to $150,000 for each song allegedly copied. According to the lawsuits, Suno is accused of copying 662 songs in training its AI model, while Udio allegedly used 1,670. This totals to a lawsuit of just under $350 million in damages. The labels also demand full disclosure of the training datasets used by the companies, accusing them of being "deliberately evasive" about the material, which, if revealed, they argue could constitute "willful copyright infringement on an almost unimaginable scale" [9].

Case 2: Getty Images v. Stability AI (2023)

Another case of importance in the realm of copyright law is the case between the stock photo provider Getty Images, and the creator of the AI model Stable Diffusion, Stability AI. Filed in February 2023 in Delaware, Getty Images alleges that Stability AI used more than 12 million of its copyrighted images to train Stable Diffusion without a license, which they argue constitutes copyright infringement. Stability AI allegedly did not seek or obtain a license to use these images, which Getty claims could have been acquired under established licensing agreements, as has been done by other technology companies [10].

Moreover, Stability's model sometimes generated images displaying Getty's watermark, which Getty argues could lead to consumer confusion and devalue its brand. This has prompted Getty to include watermark infringement alongside its copyright allegations. Getty is seeking both financial damages, including Stability's profits from the alleged infringement, and an injunction to stop Stability AI from using its images.

This case also provides more critical legal questions about whether AI companies need explicit licensing to use copyrighted material for training, especially as these companies compete with traditional creative industries. Getty's lawsuit also brings up risks of AI-generated content displaying watermarks, which is applicable to generative AI in music if these AI tracks are associated with specific artists or labels [10].

III. Key Issues and Challenges

The key issues and challenges in adapting copyright law to AI stem from an outdated legal framework that wasn't designed for machine-generated content and the rapid pace of technological change. Current laws struggle with defining authorship, determining originality, and reconciling the use of vast, often copyrighted datasets by AI developers. This section explores how these legal uncertainties intersect with the need to fairly compensate artists and protect intellectual property while still encouraging innovation in the AI space.

A. Adapting Copyright Law for AI

The complexity of AI and its uses have led to a complex implementation of copyright law in AI. At this time, judicial systems around the world struggle to adapt copyright laws to AI. Legal frameworks must adapt to these new technologies,

but they are struggling to address key topics such as authorship and originality. These topics cover questions such as:

1. Who is considered to be the author when AI creates content – the developer, the user, or the AI itself?
2. Traditional copyright laws require artists to use a degree of human creativity when creating a new product. Do AI-generated outputs meet that standard of creativity?

Policy has fallen behind the rapid development of AI systems and will likely continue to do so. If lawyers and litigators were able to predict how AI may evolve, they could pre-emptively address these solutions, allowing policy to have more dictation over how AI systems are created.

To address the complexities of AI-generated works, new laws must be introduced to handle this unfamiliar territory. There are many different mechanisms into which laws could be adapted, all of which come with their own challenges in implementation. The first of which is an AI Transparency and Attribution framework, where the disclosure of dataset sources and attribution for AI-generated outputs resembles those for specific copyrighted works. However, the issue arises while balancing transparency with protecting AI trade secrets, something which current AI companies are trying their best to maintain private.

Another option is introducing levy-based fees, where a levy is applied on AI tools or datasets with the goal of creating a compensation pool for copyright holders. This option also comes with some limitations as fair distribution among creators is challenging to determine, and AI developers may resist due to increased costs on their operations.

The third is to create licensing agreements where AI developers would need to obtain licenses for copyrighted works in training datasets. These could adopt collective licensing models similar to those for radio and streaming. This idea also has its downsides as these agreements could lead to complex negotiations, particularly for smaller developers. Alongside that issue, high licensing costs could help stifle AI innovation as higher operating costs due to these licenses could prevent AI companies from growing.

B. Compensating Artists

The recent spike in integrating generative AI to music production has presented many issues in compensating artists whose works are used as training data for these models. The primary concern is the lack of legal frameworks that would mandate compensation for artists when their music is used to train AI. This ambiguity allows AI companies to use copyrighted songs without proper licenses to develop their software. This has led to major record labels like UMG and Warner to file lawsuits against AI startups such as Suno and Udio, alleging illegal use of their music libraries for AI training [11].

However, implementing a fair compensation model faces several challenges and a fundamental paradox. For artists to be compensated, AI training data must be public and transparent, yet full transparency is impractical because of privacy risks and societal concerns. If datasets were to remain closed, artists cannot verify if their work has been used. If datasets were fully disclosed, this could pose significant societal risks. Publicly available datasets might enable the copying of AI models, leading to intellectual property theft. Too much transparency would also slow innovation and hurt the generative music AI industry by making it harder for companies to attract investment and stay competitive. This could also hurt their growth, as competitors could use their datasets, making it harder for companies to stand out in this niche market. This paradox reveals a no-win situation of developers, artists, and policymakers. Either prioritize accountability and fair compensation while putting innovation and security in jeopardy or protect AI companies and their datasets and prevent misuse at the cost of artists' rights. This dilemma underlines the need for a middle ground where artists can be compensated while AI can stay protected.

On the other hand, AI companies resist compensating artists by arguing that their use of copyrighted training data is transformative and is under the doctrine of fair use. For instance, companies Suno and Udio claim that their AI does not copy material but analyzes them for patterns. Some argue that fair use promotes freedom of expression by allowing companies and people unauthorized use of copyright-protected works under certain circumstances, with the main focus falling on transformative works [3]. Fair use is often a legal gray area, as each case is unique in its usages of copyrighted materials and thus, considerations change. This defense allows AI companies to justify their practices on uncharted ground, without knowing with certainty if what they are doing is legal or not.

C. Dataset Transparency

Generative AI models are generally trained on vast amounts of data. Suno has admitted to training on "tens of millions of recordings" [12]. Generally, the more data used to train the model, the better [13]. However, it is difficult to garner all the data you need while staying copyright free. This leads to a fundamental tension: AI companies need data but want to keep the challenging dynamic between transparency and protecting their competitive advantages. From a business perspective, the AI company's datasets represent significant competitive advantages and intellectual property, including the works they use and the specific data being utilized. Revealing detailed information about training data could potentially compromise their market position or expose them to competitors and those who believe their work does not constitute fair use. Currently, some companies have taken steps toward transparency by publishing limited information about their training data or working with specific rights holders [14] [6] [7]. Most other companies, however, maintain strict privacy. To deal with this, other methods of detecting the use of copyrighted music may be a better method to deal with the situation. These ideas are discussed in more detail below.

D. Detecting the Use of Copyrighted Music

Instead of having to detect the use of copyrighted music in generate music AI, many methods revolve around protecting the original music. Audio fingerprinting adds a unique identifier to a copyrighted piece of music, very much comparable to how humans each have unique fingerprints identifying themselves. In music, these identifiers are used by services like Shazam to identify songs. In the case of generative music AI, if an original piece and its identifier are used in training data, it is possible that the fingerprint gets carried over to the newly generated music. Then, this would be an indication that copyrighted music had been used to generate music [15]. Although, if the generative music AI startups act according to their affirmations and only incorporate subtle patterns from source material, then this makes fingerprinting less effective. Similarly, a watermark code can be embedded into copyrighted music to protect it for misuse. But models can easily transform input data to obscure the watermark or even learn to ignore such codes.

Alternatively, algorithmic similarity analysis is a technique that uses machine learning to analyze patterns and any similarities between generative music AI and copyrighted songs. This method starts by extracting features from both works such as melody, harmony, and rhythm. Once extracted, features are then represented in a high-dimensional space using embeddings. Think of an embedding as a unique location or "coordinate" for each piece in a multi-dimensional space. An easy application of this topic can be made to books in a vast library that are placed across a giant map. Their "coordinates" could be defined by their characteristics, such as length of the book, writing style, language, etc., and each trait is represented numerically. Similarly, for algorithmic similarity analysis in music, their features are numerical represented as embeddings using machine learning algorithms and are then placed on this giant high-dimensional space. Finally, this would mean that music with similar features, thus similar "coordinates", naturally cluster together. In further detail, the distance between two embeddings—such as an AI-generated song and an original piece—can be measured. Common measurements in a similarity metric are cosine similarity (which compares angles), and Euclidean distance (which measures straight-line difference). If the distance between them passes a set threshold, it may suggest that the AI-generated song was influence by or trained on the original piece. However, this method has limitations. Think of a ChatGPT detector that might detect a paragraph having a 70% chance it was generated using AI. This detector operates on probability rather than certainty. Algorithmic similarity analysis shares these philosophical issues. If a system determines there is a 70% similarity between an AI-generated song and a copyrighted song, does that draw any certain conclusions? Music inherently shares common structures, like commonly used chord progressions or similar piece structure. As a result, using this technique in lawsuits remains controversial.

IV. Case Study: [Suno v. UMG, Sony Music Entertainment, Warner Records, et al.]

Suno v. UMG, Sony Music Entertainment, and Warner Records stands as one of the most pivotal cases in shaping how copyright law applies to AI-generated music. First filed in mid-2024, the lawsuit revolves around allegations that Suno used copyrighted tracks to train its generative music AI without proper licensing. The dispute underscores core questions about fair use, authorship, and the extent to which AI can transform existing creative works. By analyzing both the plaintiffs' and defendants' perspectives, this case study highlights the legal complexities and potential industry-wide repercussions of AI-driven content creation.

A. Case Timeline

On June 24, 2024, the Plaintiffs, comprised of major record labels such as UMG, Sony Music, Warner Records, Capitol Records, Atlantic Records, and more, filed their initial complaint against Suno. The initial complaint included evidence that Suno was using copyrighted music, along with demands of $150,000 per song used in the training of Suno's generative AI software [16]. On July 9, a lawyer by the name of Shlomo Fellig from the firm Latham & Watkins officially announced that he would be handling this case for Suno, and on August 1, he filed an answer to this complaint [16]. At this time, this is the most recent update in the case. It is presumed that Suno and the major record labels are in their litigation stage, trying to reach a settlement with these record labels. If this litigation is not able to reach a settlement, then this case will proceed to a jury trial. No date has currently been set for this trial at this time.

B. Plaintiff Perspective

In the ongoing dispute between Suno and major record labels, the plaintiffs argue that the AI startup relied on unlicensed, copyrighted music to train its generative model. They claim this constitutes willful infringement, depriving creators of due compensation and control over their works. This section outlines the evidence and legal rationale underpinning the labels' stance, shedding light on how they plan to prove unauthorized use of their music. As discussed in the background information section, AI companies typically do not disclose datasets, as they can be considered to be trade secrets [17]. In this section, we will detail the type of evidence that is being used against Suno and other musical generative AI companies. The following points were taken from the response filed in the District Court of Massachusetts and should be understood as the opinion of a law firm:

1. In pre-litigation correspondence, it was stated that "Suno also claimed that its large-scale copying of sound recordings is "fair use," which was telling because fair use only arises as a defense to an otherwise unauthorized use of a copyrighted work." [3].
2. An early investor in Suno admitted that "if [Suno] had deals with labels when this company got started, I probably

wouldn't have invested in it. I think that they needed to make this product without the constraints" [18]. The constraints, of course he is referring to are implied to be copyrighted music.

3. Using targeted prompts, the plaintiffs were able to create AI-generated songs that were almost identical in output to that of their own works. The approach was to specify key identifiers from the song such as the decade of release, topic, genre, and description of artist. An example of this is "Johnny B. Goode" by Chuck Berry (copyright owned by UMG). Suno was given the prompt "1950s rock and roll, rhythm & blues, 12 bar blues, rockabilly, energetic male vocalist, singer guitarist" and fed the lyrics for "Johnny B. Goode" [3]. The result was an output entitled "Deep down in Louisiana close to New Orle," with quite a few similarities to the original, mainly in the rhythm of the melody and key.

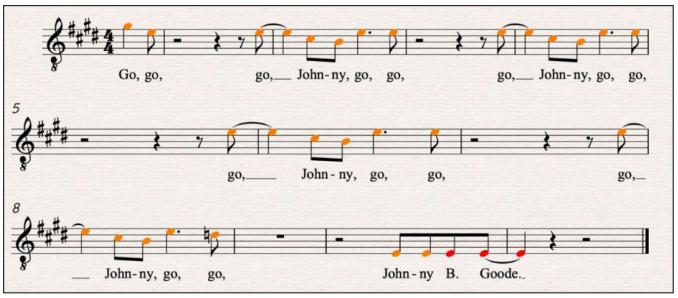

Fig. 2. Deep down in Louisiana close to New Orle (Suno generated tune) [3]

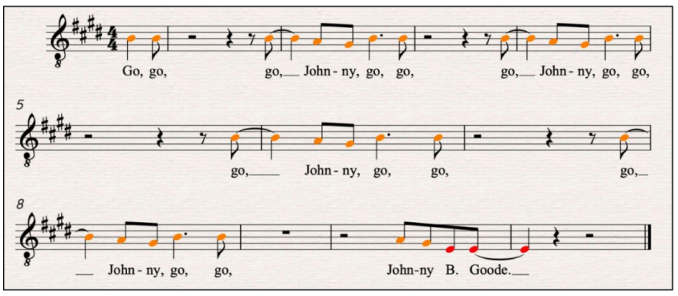

Fig. 3. Johnny B. Goode (Chuck Berry) [3]

The plaintiff was able to create 29 total responses sharing some kind of similarity (melodies in verses/choruses, rhythms, structure, etc.) to "Johnny B. Goode" using the same prompt.

4. Another common output that the plaintiff was able to generate was the AI's use of producer tags. Common in rap music, "A producer tag is a short audio clip that typically contains the producer's name or a catchphrase, used to identify their work and assert ownership over a track" [19]. One instance of the use of producer tags is in an output entitled "Rains of Castamere", which begins with the producer tag of CashMoneyAP, who is most famous for his recordings with artists Dababy and Pop Smoke.

To summarize, Suno did not disprove the claims of using copyrighted music and instead invoked "fair use," key investors have implied that they knew the AI was trained on copyrighted music, multiple songs have been recreated almost note for note through specific prompts, and producer tags from copyrighted artists are making their way into the generated songs.

As discussed in previous sections, 'fair use' has been a cornerstone of Suno's defense, as a protective measure to allow them to use copyright-protected music in AI training. The plaintiff, however, believes that "fair use" cannot be invoked in this circumstance. Here are their arguments why:

1. Suno is using copyrighted works for commercial gain. As stated previously, uses of copyrighted material for commercial gain, especially in the same market as the original, is much less likely to be considered under fair use [7].

2. The fair use doctrine describes certain use cases that can be considered fair, such as "criticism, comment, news reporting, teaching . . . scholarship, or research" , however Suno does not fall under any of the listed categories, stating "Suno's service does not offer "commentary" or "scholarship" or promote human authorship" [3].

3. The use of Suno's AI is non-transformative, and the only use for this software is to generate competing music for monetary gain, "directly proportional to the number of music files it generates". Citing the fair use doctrine "If an original work and a secondary use share the same or highly similar purposes, and the secondary use is of a commercial nature, the first factor is likely to weigh against fair use, absent some other justification for copying" [3].

C. Defendant Response

On June 24th, 2024, Hueston Hennigan LLP filed a complaint against Suno, placing the company under intense legal scrutiny. In response, Suno immediately engaged Latham & Watkins, a leading law firm with a strong reputation in AI and technology litigation, to spearhead its defense. This move underscored Suno's commitment to addressing the allegations head-on while highlighting the case's potential impact on the AI music industry. Latham & Watkins quickly filed a formal answer to the complaint, setting the stage for a high-profile legal battle. The following points were taken from the response from Latham & Watkins filed in the District Court of Massachusetts, here are their main arguments against the points made in the original complaint:

1. Suno is a tool used to make new music, designed for originality, to see how people around the world can create new songs. Suno, built from extensive analysis of all genres and styles of music, intends to mimic these styles of music, not directly copy any song. The act of generating a song in a genre violates no copyright or intellectual property (IP) laws, stating "IP rights can attach to a particular recorded rendition of a song in one of those genres or styles" [20]. The act of generating a song in a genre violates no copyright or intellectual property (IP) laws, stating "IP rights can attach to a particular recorded rendition of a song in one of those genres or styles" [20].

2. The major record labels frame their concern as creating copies of pre-existing music, but what record labels are really after is to shut Suno down, effectively eliminating competition from the market. “Where Suno sees musicians, teachers, and everyday people using a new tool to create original music, the labels see a threat to their market share.” [20].

3. Suno has constructed multiple “guardrails”, specifically to ensure that no Suno generated output related too closely to a particular song used in the training process. This includes but is not limited to using industry standard software to ensure that user inputted audio clips are owned by the user, and not commercial. The software they referenced is most likely similar to Content ID or Shazam, audio fingerprinting software used by Youtube and Apple Music respectively [20].

4. “It is fair use under copyright law to make a copy of a protected work as part of a back-end technological process, invisible to the public, in the service of creating an ultimately non-infringing new product.” [20]. This statement is true; however each case is unique and must be weighed against four main factors of fair use: purpose and character of the use, nature of the copyrighted work, amount and substantiality used, and effect on the market.

These arguments aim to establish Suno as a tool used to make new music, which is not copyrightable under Fair Use and copyright laws. They claim that Suno is being targeted as “competition to the market” as these major record labels have established a monopolistic hold on the music sphere. The defendants ensure they have installed the proper protection to ensure their outputs are unique and argue that the usage of copyrighted music is protected under fair use laws as a part of a back-end technological process.

D. Looking Forward

Ultimately, what this case boils down to is an application of fair use and copyright law. Both sides have made arguments as to why fair use applies or does not in this circumstance. Due to the uniqueness of AI systems, courts will have to carefully consider the behaviour of this software, to determine its impact on the field. Copyright law protects creativity and ingenuity, so there is a main question that courts have to answer: is there a difference between an AI using a song from its dataset as inspiration versus a musician taking inspiration from an artist? The outcome of this case is set up to redefine the intersection of AI innovation and copyright law, setting critical legal precedents that will influence the entire generative music landscape. As the courts continue to discuss the applications of fair use in the context of AI training datasets, a ruling in favor of either party could catalyze significant shifts in industry practices. This will affect how data is sourced, utilized, and disclosed. This case, therefore, not only impacts the music sector but also offers an application for copyright law in AI in multiple sectors. Looking to the future, companies can mitigate legal risks and heighten ethical practices by using more transparent data practices and fair compensation frameworks. If more companies tried to implement ethical practices, such as licensing agreements or non-copyrighted datasets, we would see a trend of less and less cases and civil suits.

V. Proposed Solution

This section outlines our proposed solution, which is designed to safeguard the rights of copyright holders while facilitating innovation in generative music AI. It is important to note that the implementation of this solution is contingent upon the assumption that legal precedent has not yet established that the use of copyrighted works qualifies as fair use. Our approach is built on two core components: a detection system for identifying instances where copyrighted music is used during AI training, and a compensation mechanism to ensure that creators are fairly remunerated for such usage. The detection component employs advanced watermarking and fingerprinting techniques to accurately flag any unauthorized replication of original works. Meanwhile, the compensation component aims to establish a transparent, levy-based framework that directly channels revenue to the rightful copyright holders. By integrating these two components, our solution seeks to strike a balance between fostering technological innovation and upholding the integrity of intellectual property rights.

A. Detecting Copyrighted Music

1) ***Legal and Regulatory Compliance Rules****:* A solution that aims to handle a large amount of secure data from top AI companies and record labels can break laws and regulations, leading to counter suits. This section aims to address the main regulations before discussing how we plan to circumvent them. IP laws, such as the Copyright Act of Canada [21] require proof of ownership and unauthorized use to initiate legal proceedings which makes this capability especially relevant under such intellectual property laws. Eliminating the need to share raw data between companies becomes a core principle in our solution in order to comply with key IP laws. Proving ”similarity” between items in datasets does not consistently hold up in a court of law as a substantial amount of proof [22]. As such, an optimal solution could implement some kind of test that does not rely solely on a similarity score, but instead can tell with absolute certainty if an item in the dataset is copyrighted. Our solution aims to establish a legal framework capable of detecting and stopping copyright infringement. These methods include:

1. Protecting personal data: Training data must stay within local boundaries to meet privacy law requirements.

2. Verifying ownership: Organizations can establish unambiguous AI model ownership proof through the use of adversarial fingerprints and distinctive watermarks.

3. Tracing infringers: Legal accountability for unauthorized AI-generated music distribution can be achieved through tracking models with embedded identifiers.

4. Enhancing copyright enforcement: Ensure AI-generated content compliance with intellectual property laws through

auditable verification processes that produce legally admissible evidence.

*2) **Introduction to Split Learning**:* We selected the split learning paradigm because it best aligns with the requirements of this compliance check. The neural network architecture divides across different parties when using split (or vertical) federated learning. The client stores the initial model components and input data while the server holds the rest of the model layers and produces the outputs. The cut layer in split learning exchanges intermediate activations ("smashed data") between parties while raw inputs and complete model parameters remain undisclosed. This setup matches our scenario: The record label (client) feeds its music data into the initial layers of a compliance-check model while the AI company (server) processes this data using its own model weights to complete the forward pass. Sensitive information remains confidential because the record label cannot access the company's model details while the company cannot access the label's raw audio.

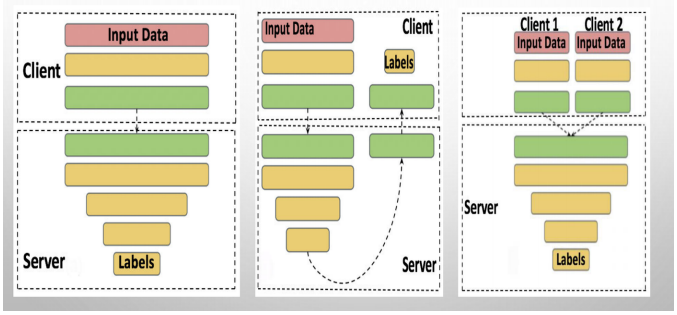

Fig. 4. Illustration of the split learning configuration [23]

For instance, in this setting, each client trains a partial model up to a specific layer called the 'cut layer'. Only the intermediate features at the cut layer (boundary between yellow and green) are sent to the server, and gradients at that cut layer are returned to the client for training. This allows joint model processing without sharing raw data.

Structure: For compliance checking, we configure a two-part neural network: The record label runs a feature extractor on its audio data locally, up to a cut-layer. The extractor could be a lightweight CNN or audio encoder that transforms the song into an intermediate feature representation. The AI company attaches a corresponding detector head (the remaining layers) that takes the intermediate feature and produces a compliance result (for example, a likelihood score or an "infringement" prediction). During an infringement check, the process works as follows:

- The record label's system takes one of its songs (or a unique fingerprint derived from it) as input and forward-propagates through the local cutlayer of the network. This yields an encoded representation of the song. Crucially, this encoding is abstract — it does not reveal the raw audio, but it captures patterns the later layers can analyze [23].
- The "smashed" features at the cut layer are securely transmitted to the AI company's server. The AI company then forward-propagates that activation through the rest of the model (or a special compliance-check subnetwork) using its private model parameters. For example, the AI company's model (or detector head) might compute a similarity or likelihood that the input audio was part of its training data.
- The final output (which might be an encrypted or blinded result – see privacy measures below) is sent back to the record label. This output could be, for instance, a probability or an error metric indicating how closely the AI model's knowledge matches the input song. If the value is above a certain threshold, it suggests the model was likely trained on that song (possible copyright infringement), if not, it suggests no memorization of that content.

The split learning method enables every participant to access only the information they need. The record label reveals only derived features from its music files while the AI company keeps its model weights private and does not get the raw input. Split learning offers enhanced model privacy protection compared to horizontal FL because each party has access to only parts of the model which prevents full visibility of the entire model to any single entity [24]. It is ideal when parties hold different modalities of data (here, one had a model or model updates, while the other has data to test), effectively creating a form of federated inference on combined inputs. Through split federated learning record labels and AI companies engage in a compliance check workflow that mirrors model training/inference activities while preventing data pooling. Using a shared split model the organizations process data collaboratively to identify copyright overlaps while maintaining local storage of proprietary information.

*3) **Privacy-Preserving Techniques for Secure Compliance**:* Ensuring privacy is paramount: The music owned by record labels and the parameters of the AI company's model must remain hidden throughout the federated compliance check. Our system employs multiple cryptographic and privacy methods to protect the federated process:

Secure Multi-Party Computation (SMPC): SMPC enables multiple participants to perform joint computations over their input data while keeping those inputs hidden from each other [25]. Our design uses SMPC protocols to enable collaborative analysis between the AI model and the record label's data. For example, the record label and AI company can employ an SMPC framework (like Facebook's CrypTen or Microsoft SEAL) to evaluate the AI's model on the label's song in a secret-shared manner. The model's computations (matrix multiplications, etc.) are performed on encrypted or secret-shared values, so the AI company never sees the actual audio features and the record label never sees the raw model outputs. This could be implemented with an additive secret sharing scheme: the record label secret-shares the intermediate features with the AI company (or a neutral server), and the AI company secret-shares its model parameters. They then perform the forward-pass computation by exchanging masked values. At the end of the computation, only the final result (e.g., a risk score) is revealed (and only to the party authorized to see it). SMPC essentially functions like a virtual trusted calculator:

for instance, it can let the AI company privately evaluate its model on the record label's data [25], or allow multiple labels to perform a joint aggregate audit without sharing individual data. The process ensures that proprietary data and models stay secure and undisclosed throughout the compliance check.

Homomorphic Encryption (HE): Homomorphic encryption [26] enable users can perform calculations on encrypted data to receive encrypted results that decryption is possible only via owner of the secret key. HE can be used in selected areas of the compliance pipeline to improve data protection. For example, a record label can encrypt an audio feature vector with their public key and send it to the AI company. The AI company then runs its model on the encrypted data without ever decrypting it, using partially homomorphic operations (addition, multiplication on ciphertexts). The outcome is an encrypted infringement indicator, which only the record label can decrypt to see the result. Throughout this process, the AI company learns nothing about the input or the output, since all intermediate data remains encrypted. Fully Homomorphic Encryption (FHE) schemes (which allow arbitrary computations on ciphertexts) can be heavy for deep learning, hence we use optimizations like leveled HE or Partially Homomorphic Encryption (PHE) for specific operations to keep overhead reasonable [27]. For instance, computing a simple dot-product similarity or a reconstruction error between the model's output and a target song can be done under HE if we linearize the operation. Non-linear operations (like activations) can be handled either by the split learning approach (so that they occur on the AI's side in plaintext on already encrypted inputs) or by efficient garbled circuits if needed. By carefully choosing which parts of the computation to encrypt, we ensure a balance between privacy and performance. The cryptographic strength of HE means even if communications are intercepted, the content (songs or model responses) remains unintelligible without the decryption key.

Differential Privacy (DP): We incorporate differential privacy to protect against information leakage in any shared outputs or updates. DP works by adding carefully calibrated random noise to results so that the presence or absence of any single data record is indistinguishable. In our context, the AI music company could train its model with DP-SGD (differentially private stochastic gradient descent), which would ensure the model does not memorize specific training examples (like a particular song) too exactly. This pre-emptively protects against infringement, because a DP-trained model is unlikely to regurgitate any one song verbatim. Even during the federated compliance check, DP can be applied. For example, if multiple record labels participate in a joint audit, the aggregated compliance metrics can have noise added before being revealed, so that no single label learns specifics about another label's queries. The record label's query results themselves could be noised if we only need a yes/no answer with high confidence. Importantly, the noise levels are set such that they do not obscure true infringements but hide minute details of the model's behavior. DP ensures that any one song (even if it were in training) has a limited influence on the output, preventing the exposure of exact memorized content. This technique is computationally cheap (just noise addition) and scales well, complementing heavier cryptographic methods by reducing how much sensitive information even exists in the computed results.

Zero Knowledge Proofs (ZKP): Zero-knowledge proofs allow a party to prove a statement about data or computations without revealing the data itself [28]. We utilize ZKPs to make the compliance process verifiable and legally defensible. For instance, after training, the AI company can generate a zeroknowledge proof of training that attests "This model was trained only on licensed data and did not include Record Label X's songs" without revealing anything about the training data or model parameters. Recent advances in ZKPs for ML (zkML) enable proving properties of models, such as training steps or dataset membership, in a computationally feasible way [29]. Concretely, the AI company could commit to the dataset it used (e.g., via a cryptographic hash or Merkle root of all training data) and then provide a ZK-SNARK proof that none of the record label's song hashes are in the committed dataset. This is akin to a zero-knowledge set membership test – proving a set intersection is empty without revealing the actual sets. Alternatively, the AI company can prove that it followed a prescribed training procedure (for example, a training run with differential privacy enabled, or only using a specific approved dataset) [29]. The record label (or a regulator) can verify this proof and be mathematically assured of compliance. Another use of ZKP in our system is for the infringement test itself. If the record label gets a negative result (no infringement detected), the AI company could output a ZKP that the test was carried out honestly on the model in question (preventing an AI company from swapping in a different "clean" model just for the test). Although generating ZK proofs for deep learning computations can be resource-intensive, we confine their use to periodic checks or final audits to keep it tractable. The outcome is that compliance checks are trustless – the record label doesn't have to trust the AI company's word, they have cryptographic proof of either compliance or violation, which is crucial for legal defensibility. By combining these measures – SMPC/secure computation for processing data, HE for data encryption in transit and compute, DP for output privacy, and ZKP for process verification – we create a robust privacy-preserving compliance system. Each technique is chosen to minimize performance hits. For example, we use partial HE and secret-sharing (which are faster than full FHE), we add only small noise for DP, and we generate ZK proofs for high-level properties rather than every single operation. The overall design ensures that at no point is sensitive information exposed in plain form, yet all parties can collaboratively achieve the goal of detecting unlicensed training. When the model training is complete, the AI company can package a compliance report: this might include the ZK proofs of training, differential privacy parameters used, and summary of any internal compliance tests. The record labels, through the federated system, get the ability to verify this report and test the model themselves, yielding high assurance that if the model passes, it truly did

not use unlicensed music. To avoid high financial cost for any single entity, the federated compliance system can be managed by a neutral third-party service or consortium of stakeholders. This service can maintain the secure aggregation server and coordinate cryptographic key management. Using cloud computing with hardware acceleration (like FPGAs for HE, or SGX secure enclaves as a backup option) can speed up cryptographic operations, reducing runtime and therefore cost. Also, many cryptographic libraries are open-source and optimized, meaning the main expense is computing time. With model compression and batching, we ensure that even large models can be handled with commodity hardware given some time (hours, not days, for a thorough audit of a big model against millions of song samples, for instance). The design favors one-time heavy computations (like proof generation or full-catalog scan) only when absolutely necessary (e.g., a legal dispute), whereas routine compliance checks can be much lighter (sampling a subset of songs, using partial evaluations, etc., to get a quick assurance).

*4) **Workflow**:* **Step 0:** Initial Setup and Key Exchange – All participating entities (the AI company and one or more record labels, or an auditor) set up the cryptographic environment. This involves generating encryption keys (public/private key pairs for HE for each label, key shares for SMPC, etc.) and exchanging any public parameters. They also agree on the model split architecture and the protocol (which cut layer, what format features will be, what threshold constitutes a violation, etc.). For example, the AI company publishes the architecture of the compliance model or the fingerprinting method it will use. A central coordination server (could be run by a neutral party or consortium) may exist to facilitate scheduling and key management, but it will not see any raw data or models.

Step 1: Registration of Data Commitments – The AI company commits to its training dataset and model. It computes a commitment hash (or Merkle root) of all training data it used. This is submitted to a smart contract or to the record labels in a ledger so that it's fixed (the company can't later change it). Likewise, each record label prepares a fingerprint database of their copyrighted songs – e.g., a set of audio hashes or embeddings – and commits to those (so that they can't maliciously add more songs later just to trap the AI company). These commitments will be used in ZK proofs later. This stage ensures both sides "lock in" the items of interest without revealing them.

Step 2: Local Model Training (AI Company) – The AI company trains its AI music model on its own data (e.g., publicly licensed music, user-generated music, etc.). This is done using its standard pipeline, possibly with differential privacy and logging as described. No external interaction is needed during core training, so no overhead is incurred here aside from any self-chosen privacy technique. Once the model is trained (or at certain checkpoints), it is saved for audit. Let's assume the model is now ready to be checked for compliance.

Step 3: The coordinator server notifies the AI company and relevant record label(s) that a check will happen. They establish a secure session. The AI company provides the server-side model for the split learning inference – typically, this means loading the second part of the model on a secure computation server. If using SMPC, the AI company secret-shares or encrypts its model weights with the computation service (or among multiple servers). If using a TEE (Trusted Execution Environment) as an aid, the model could be loaded into an enclave. In any case, the AI company does not give the model in plaintext to the label, it only makes it available in the secure protocol. The record label in turn prepares its input data for the check. For instance, it selects a batch of 100 songs (or segments) that it strongly cares about. The label either keeps these on its local machine (for split learning) or encrypts them with homomorphic encryption (if the model will process them directly in encrypted form). All parties confirm readiness.

Step 4: Federated Inference/Processing – The record label's client-side application now goes through the selected songs one by one (or in batches). For each song, it does the following:

- Compute the feature representation (e.g., passes it through the local cutlayer of the model or simply prepares the raw input if using HE directly).
- Send the intermediate activation to the AI company's model server over an encrypted channel (TLS + the values might already be secret shares or encrypted numbers). If using pure HE, send the encrypted audio/features to the server.
- The AI company's server (or the joint MPC nodes) then perform the forward pass on the encrypted/secret shared data through the remaining network layers. For example, it computes the output logits or reconstruction of the input. Since the model is large, this computation is optimized as discussed (maybe using GPU, etc.).
- The server returns the encrypted result of the inference back to the record label. This might be the log-likelihood of the sequence, a set of output audio tokens, or a high-level "yes/no" flag in secret-shared form.

Step 5: Compliance Metric Computation – The record label now decrypts or reconstructs the results from Step 4. If the result was an encrypted likelihood score, the label decrypts it with its HE secret key. If it was done via MPC, the label combines its share of the result with the shares from the server to obtain the final number. Now the label has, for each song tested, a metric indicating how strongly the model reacted. The label compares these metrics to the expected range for non-members. For instance, if a certain song has a model likelihood far above a threshold (meaning the model highly likely has seen it [30], the system flags this song as a potential infringement. In practice, the label might set a threshold based on a statistical confidence (e.g., "if probability that the item is in the data set is greater than 0.9, flag it"). They could also use an internal classifier on the outputs – for example, if the output was the model trying to continue the song, the label can measure similarity between the continuation and the original. Some systems might automate this: e.g., compute a cosine similarity between audio embeddings of the original and the generated continuation. High similarity would yield

a flag. These calculations are done on the label's side, so no privacy issue arises. The outcome of this step is a compliance report: perhaps a table of songs vs. scores, highlighting any that exceed the infringement threshold.

Step 6: Result Sharing and Proof Generation – Now the record label has preliminary results. If all songs are in the clear (no suspicious scores), the AI model likely did not use any of the label's data. The record label can then cryptographically sign an attestation that "We, Label X, have tested Model Y on [date] and found no evidence of training on Label X's catalog." This attestation can be shared with the AI company as part of a compliance certificate. On the other hand, if any song was flagged, the system can escalate. The record label can notify the AI company (most likely through the protocol, without revealing which song in plaintext, at least initially). They might say "Song ID #5 from our hashed list appears to have been used in training. We request an explanation or remediation." At this stage, the AI company has the option to contest or accept. If contesting, this is where zero-knowledge proofs or additional verification come in. The AI company might invoke the previously computed commitment of its training set and perform a private set intersection (PSI) with the label's song in question. PSI can definitively show if that song (or its fingerprint) was in the training set, without the AI co learning which song it is (if done properly). If PSI comes out positive, it's proof of infringement. If PSI is negative but the model's behavior was still highly suspect, it could indicate the model learned something very close to the song (e.g., an overfitted surrogate). In either case, the parties now have cryptographic evidence. Optionally, they can involve a neutral auditor who reviews the evidence (the auditor could be given access to the song under NDA and maybe run a targeted test themselves for confirmation).

- The AI company can produce a ZK-proof that the model tested was indeed the one corresponding to the committed training hash. This prevents a scenario where the company trained a second "decoy" model without the label's songs just to pass the test. The proof would show that the weights of the deployed model are a result of training on the committed dataset (or at least that they match a certain hash that was committed). Such a proof might use zk-SNARKs as described in the proof-of-training concept [29].
- If the result is clean, the AI company might also produce a ZK-proof that none of the label's songs (from a committed list) appear in its training set. This could use a zk-proof of set disjointness, which might be heavy, but perhaps they only do it for a small set of top songs.

Step 7: Compliance Outcome – After analysis, one of two outcomes occurs. No Infringement Detected: All checks pass. Or, one or more labels detected their content in the model. In this case, the system can automatically provide evidence to the AI company and a regulator.

Step 8: Ongoing Monitoring – The federated system remains available for future checks.Throughout this process, all actions (from key exchange to final verification) are designed to be auditable and repeatable. Each cryptographic message or proof can be logged (in encrypted form) to provide a trace in case of disputes. The combination of federated learning structure and advanced privacy techniques ensures that compliance verification is done scientifically and rigorously, minimizing trust and subjectivity. The result is a feasible, efficient, and privacy-preserving federated system that upholds copyright law without stifling the development of AI models. By balancing the load between parties and using cutting-edge cryptography, the solution scales to real-world industry usage – enabling record labels to defend their intellectual property and AI companies to innovate with accountability.

B. Compensation Model

One of the most pressing issues in AI-generated music is how to fairly compensate artists and record labels for the use of their works in training models. As AI technology's popularity skyrockets, debates over licensing and copyright have grown. This section explores the complexities of crafting effective compensation frameworks that balance innovation with fair treatment of creators. First, we introduce two methods we considered using as our compensation model. We then outline our full model, along with an example of our system using a mock dataset.

1) ***Royalty-Based System****:* The first method we considered was a royalty-based system. Music streaming services utilize royalty-based systems in their compensation methods. It works by giving artists money based off the number of streams or plays, usually a fraction of a cent per stream. So how can we implement this system? Well, similar to streaming services, when models are trained, they must "play" a song in order to for the model to learn the properties of the song and genre such as rhythms, melodies and chords. Models usually require multiple full passing of training data through the model (known as epochs). In an epoch, some data can be passed through more than others, known as oversampling and undersampling. For example, if you are training a musical AI to create songs from multiple genres but your dataset has much more rock music than it does classical, your epoch can reuse some classical songs and leave some rock songs out of the training to represent the two genres equally. The royalty system aims to compensate artists for each play that their song has in the training process, acting similarly to how streaming services such as Spotify work. The cons of implementing this method is the information needed to implement this solution. Being able to see the source code and determine exactly how a model was trained is the only way to implement a royalty-based system. Many AI companies would consider this to be proprietary knowledge, and as such would not be open to disclosing this information. The main pros of a royalty-based system are that artists are compensated by plays, which means that if a song was used more in the training of the system, they are entitled to a bigger cut. If a song is used more in training, then it is more likely for the output of the generative AI to share key properties with that song.

2) ***Levy-Based System****:* A levy is defined as "(of a government or organization) to demand an amount of money, such

as a tax, from a person or organization" [31]. In a levy-based system, AI companies would be demanded to give an amount of money, which would then be pooled out amongst the artists in the dataset. This requires knowing which artists are in the dataset and also knowing how many of each of their songs are represented. Once this is done, we can divide the number of songs each artist has in the dataset by the total amount in the fund to evenly distribute the amount that each artist gets from this levy fund. The cons of implementing this method are that the data is not as fairly represented, as we have discussed with the oversampling and undersampling case above, thus some artists may get more than they actually deserve. The pros of this method are that it is easier to retrieve the data such as the number of songs that each artist has in a data set over trying to determine the specific training process.

*3) **Compensation Framework**:* Our compensation framework first operates by receiving the data from our FL framework, telling us which songs are above the threshold of probability, and therefore, are very likely to be in the dataset. Using this data, we can then generate a list of how many songs each creator has in the dataset. Due to our FL framework, we have decided that a levy-based system better fits the compensation model, but we still need to decide how to generate our levy (i.e. how much money each company will be paying to copyright holders).

ProRata.AI, an artificial intelligence start-up, has established a revenue-sharing model designed to fairly compensate content creators for their work when used by AI systems. The company, unlike other traditional AI models which tend to scrape online content without compensation publishers, has claimed to share half the revenue from subscriptions to its platform with its licensing partners. This already includes Universal Music, Axel Springer, Financial Times, The Atlantic, and Fortune. The main goal of the company is to license the technology behind its search engine to other generative AI companies. If an AI company were to adapt this business model, all of their lawsuits would come to an end, states ProRata CEO Bill Gross; "If you adopt this business model, this will end your lawsuits, because now you'll be sharing revenue properly" [32].

ProRata generates revenue mostly through an AI-powered ad platform that places relevant ads within AI search results and digital content. They also have revenue streams from their proprietary attribution technology which can be licensed to other AI companies as a service, monitoring services that track content usage by AIs, and potentially subscription options through their Gist.ai search engine that showcases their attribution tech, although this technology is still in its Beta-testing stage.

ProRata.AI is currently valued at over $130 million, and off the backs of a successful and ethical business model, we plan to adapt this to music AI companies. Taking Suno as an example, their current business model works almost entirely off subscriptions. Suno has three different subscription tiers, each with different features and pricing.

• Free Plan: 50 credits/day (equivalent to 10 songs), Suno retains copyright.
• Pro Plan: $8/month, 2,500 credits/month (500 songs), users hold the right to their creations.
• Premier Plan: $24/month, 10,000 credits/month (2,000 songs), users retain full rights [33].

Suno has not released their revenue model to the public. However, we know that in 2023, Suno partnered with Microsoft Copilot to introduce AI music generation in the Copilot software. This partnership does not include any financial agreement, rather a win-win situation that brings more users to Suno, while adding additional features to Microsoft Copilots AI software [33]. As such, we currently estimate that Suno's revenue come 100% from subscription fees. As such, half of all Suno's revenue would be sent to the levy to then be distributed to artists. To show our compensation framework in action, we have generated a "mock company" to demonstrate. By taking some of the top musical generative AI companies and modelling them based on their revenue, we can show what a major AI company would have to pay in our framework. Suno has not released their revenue model to the public, but we know that in 2023, Suno partnered with Microsoft Copilot to introduce AI music generation in the Copilot software, however this partnership does not include any financial agreement, rather a win-win situation that brings more users to Suno, while adding additional features to Microsoft Copilots AI software [33]. As such, we currently estimate that Suno's revenue come 100% from subscription fees. As such, half of all Suno's revenue would be sent to the levy to then be distributed to artists. To show our compensation framework in action, we have generated a "mock company" to demonstrate. By taking some of the top musical generative AI companies and modelling them based on their revenue, we can show what a major AI company would have to pay in our framework.

Company	**Revenue ($USD)**
Aiva Technologies	$1.5M
Beatoven.ai	$37.8k
Amper Music	$5.1M
Boomy	$5.8M
Suno	$8M
Music.AI	$22.1M

TABLE I
COMPANY REVENUES [34], [35], [36], [37], [38], [39]

The average revenue of $13.26 million USD across major AI music companies (Aiva Technologies, Beatoven.ai, Amper Music, Boomy, Suno, and Music.AI) provides a useful baseline for projecting potential artist compensation frameworks. If these companies were to adopt a ProRata.ai-style revenue-sharing model, allocating 50% of subscription revenue to rights holders, and assuming an 100% subscription-based revenue structure approximately $6.63 million would be distributed to artists and copyright holders annually per company. We have generated a mock dataset, comprised of just under 5,000 songs that a generative music AI company could use to train their system [40]. This dataset represents a multitude of genres, artists, and time periods so a user could realistically generate a song from almost every genre and style. Our goal

in this model was to showcase how a levy-based system could work based upon knowing the number of songs that each artist has in a given training set. Based upon this model, we will show you some of the top performers from this dataset, and how much they are owed based upon our compensation formula.

Track Artist	Unique Count	Compensation
Bad Bunny	30	$8,816.10
Ren Avel	26	$7,640.62
Asake	21	$6,171.27
Bnxn	19	$5,583.53
Seyi Vibez	18	$5,289.66
LoFi Waiter	18	$5,289.66
Wizkid	16	$4,701.92
Linkin Park	14	$4,114.18
Hozier	13	$3,820.31
Sabrina Carpenter	11	$3,232.57
Burna Boy	11	$3,232.57
Zinoleesky	11	$3,232.57
Billie Eilish	10	$2,938.70
Red Hot Chili Peppers	10	$2,938.70
Central Cee	10	$2,938.70
Yume.Play	10	$2,938.70
Green Day	9	$2,644.83
Celine Dion	9	$2,644.83
Metallica	8	$2,350.96
Gunna	8	$2,350.96
Lil Baby	8	$2,350.96
Brent Fiyaz	8	$2,350.96
Bruno Mars	8	$2,350.96
J Balvin	8	$2,350.96
Paramore	8	$2,350.96
My Chemical Romance	8	$2,350.96
Zhao Ying	8	$2,350.96
Hao Yu	8	$2,350.96

TABLE II
TOP ARTISTS COMPENSATION

Based on our calculations, each song in the dataset is entitled to $1,372.10. Pictured above, we can see artists of many genres (latin, rap, pop, metal, rock, etc.) that have all made a significant impact on the dataset. Once these values are generated, AI companies are made aware of these values and will distribute the money to the necessary labels and copyright holders. We believe this method to be best implemented as a third party software to best ensure the protection of both parties' data. Record labels would enlist the help of the third party, who would then be responsible for compiling all required data, and only sending vital information when needed. The information received by the both companies (record label and AI company) during this process would be the number of songs by each artist that (our system believes) appears in the dataset, along with the breakdown of how much each artist receives, in a format similar to Table II. We designed our compensation framework to create the best possible outcome, a win-win scenario for both companies that avoids legal fees and time spent dealing with legal issues. Currently, the case against Suno is worth $150,000 per song, and totals to over $350 million. If suno were to implement our framework, this number goes down to $22.5 million and saves valuable company time and resources by not having to fight a litigious lawsuit. At the same time, record labels are also avoiding legal fees and court time, while receiving a steady, annual payout from AI companies. This means that record labels are actually given incentive to help AI companies grow, and thus could lead to a less competitive music culture.

VI. THE FUTURE OF MUSIC & AI

As the future of copyright law and music AI is still unsure, we can discuss two possible futures. One in which all materials are indeed fair use, and another does not believe these materials are fair use. If precedent is set in the copyright AI world that all copyrighted materials are indeed fair use, then this solution will no longer be viable to implement, as AI companies will not have to compensate copyright holders. Since the future of copyright law and AI remains unsure in the public eye, we set out to address all possible outcomes of these important cases. Our solution introduces ethical business practices by integrating advanced watermarking and fingerprinting techniques into AI training processes, ensuring that artists receive fair compensation while safeguarding intellectual property rights. This approach not only fosters transparency in data usage but also creates a win-win scenario for both AI developers and creators. Experts predict that future copyright law will evolve to address the unique challenges posed by AI, potentially leading to new legal precedents that better recognize the transformative nature of AI-generated music [41]. As ethical practices gain traction, we anticipate a shift toward revenue-sharing models that empower artists and encourage responsible innovation. Ultimately, this balanced framework sets the stage for a sustainable music ecosystem where technology and creativity coexist harmoniously.

REFERENCES

[1] Soundraw, "Top 8 ai-generated songs you need to hear in 2025," *Soundraw Blog*, 2025, [Online]. [Online]. Available: https://blog.soundraw.io/post/ai-generated-songs-you-need-to-hear

[2] V. Yurkevich, "Universal music group calls ai music a 'fraud,' wants it banned from streaming platforms. experts say it's not that easy," *CNN*, 2023, [Online]. [Online]. Available: https://www.cnn.com/2023/04/18/tech/universal-music-group-artificial-intelligence/index.html

[3] D. J. Cloherty, I. UMG RECORDINGS, L. CAPITOL RECORDS, S. M. ENTERTAINMENT, A. R. CORPORATION, A. R. G. LLC, R. E. LLC, I. THE ALL BLACKS U.S.A., W. M. I. S. LIMITED, and W. R. INC., "Umg recordings, inc., et al. v. suno, inc." 2024, [Online]. [Online]. Available: https://regmedia.co.uk/2024/06/24/suno-complaint.pdf

[4] "How does ai music work? from machine learning to viral hits," 2025, [Online]. Accessed: Jan. 27, 2025. [Online]. Available: https://rareconnections.io/how-does-ai-music-work/

[5] ar5iv, "Melnet: A generative model for audio in the frequency domain," 2025, [Online]. Accessed: Mar. 01, 2025. [Online]. Available: https://ar5iv.labs.arxiv.org/html/1906.01083

[6] Copyright.gov, "What is copyright?" accessed: Nov. 06, 2024. [Online]. Available: https://www.copyright.gov/what-is-copyright/#:~:text=U.S.%20copyright%20law%20provides%20copyright,rental%2C%20lease%2C%20or%20lending

[7] U.S. Copyright Office, "U.s. copyright office fair use index," accessed: Nov. 06, 2024. [Online]. Available: https://www.copyright.gov/fair-use/index.html

[8] A. Academy, "The rise of ai in audio engineering: How machine learning is revolutionizing music production," 2025, [Online]. Accessed: Jan. 15, 2025. [Online]. Available: https://audioacademy.in/the-rise-of-ai-in-audio-engineering/#:~:text=For%20example%2C%20companies%20like%20iZotope, extensive%20technical%20knowledge%20or%20experience.

[9] B. Brittain and B. Brittain, "Music labels sue ai companies suno, udio for us copyright infringement," https://www.reuters.com/technology/artificial-intelligence/music-labels-sue-ai-companies-suno-udio-us-copyright-infringement-2024-06-24/, Jun. 2024, accessed: Nov. 06, 2024.

[10] B. Brittain, "Getty images lawsuit says stability ai misused photos to train ai," https://www.reuters.com/legal/getty-images-lawsuit-says-stability-ai-misused-photos-train-ai-2023-02-06/, Feb. 2023, accessed: Nov. 06, 2024.

[11] S. Dhameliya, "Record labels sue ai platforms making music: Universal, sony, warner v/s suno, udio," 2025, accessed: Jan. 29, 2025. [Online]. Available: https://iprmentlaw.com/2024/07/13/record-labels-sue-ai-platforms-making-music-udio-v-universal-sony-warner/

[12] J. Koebler, "Ai music generator suno admits it was trained on 'essentially all music files on the internet'," *404media*, 2025. [Online]. Available: https://www.404media.co/ai-music-generator-suno-admits-it-was-trained-on-essentially-all-music-files-on-the-internet/#:~:text=The%20AI%20music%20generator%20company,tens%20of%20millions%20of%20recordings.%E2%80%9D

[13] E. Balla, "Here's how much data gets used by generative ai tools for each request," *Data Science Central*, 2025, accessed: Feb. 17, 2025. [Online]. Available: https://www.datasciencecentral.com/heres-how-much-data-gets-used-by-generative-ai-tools-for-each-request/

[14] Musette, "Why artificial intelligence will never replace musicians," *Medium*, 2025, [Online]. [Online]. Available: https://musettedc.medium.com/why-artificial-intelligence-will-never-replace-musicians-a03600b0310f

[15] "Ai in music copyright detection — restackio," 2025, accessed: Feb. 01, 2025. [Online]. Available: https://www.restack.io/p/ai-in-music-answer-copyright-detection-cat-ai

[16] "Umg recordings, inc. et al v. suno, inc. et al," [Online] Available:. [Online]. Available: https://dockets.justia.com/docket/massachusetts/madce/1:2024cv11611/272063

[17] B. Chapman, C. Taylor, and B. Husband, "Protecting training data for ai innovations in the medtech space," *Carpmaels & Ransford*, Apr. 2024, [Online] Available:. [Online]. Available: https://www.carpmaels.com/protecting-training-data-for-ai-innovations-in-the-medtech-space-part-1/#:~:text=If%20there%20has%20been%20a,may%20constitute%20a%20trade%20secret.

[18] B. Hiatt, "A chatgpt for music is here. inside suno, the startup changing everything," *Rolling Stone*, 2024, [Online]. [Online]. Available: https://www.rollingstone.com/music/music-features/suno-ai-chatgpt-for-music-1234982307/

[19] A. Jawed, "What is a producer tag: All you want to know," https://www.hollyland.com/blog/tips/what-is-a-producer-tag, accessed: 16 March 2025.

[20] S. Fellig, A. M. Gass, B. N. Lovejoy, S. N. Feldman, N. Taylor, and S. V. Damle, "Answer of defendant suno, inc. to complaint," 2024, [Online]. [Online]. Available: https://www.musicbusinessworldwide.com/files/2024/08/SUNO-response-to-copyright-suit.pdf

[21] G. of Canada, "Copyright act (r.s.c., 1985, c. c-42)," November 2024, [Online; accessed 2025-01-15]. [Online]. Available: https://laws-lois.justice.gc.ca/eng/acts/C-42/Index.html

[22] S. Vondran, "Proving "substantial similarity" in copyright infringement actions," September 2024, [Online; accessed 2025-02-15]. [Online]. Available: https://www.vondranlegal.com/proving-substantial-similarity-in-copyright-infringement-actions

[23] MIT Media Lab, "Split learning: Distributed and collaborative learning," https://www.media.mit.edu/projects/distributed-learning-and-collaborative-learning-1/overview/, online; accessed March 2025.

[24] S. C. C. Thapa, M. A. P. Chamikara and L. Sun, "Splitfed: When federated learning meets split learning," *arXiv preprint*, vol. arXiv:2004.12088, Apr. 2020.

[25] A. H. S. S. M. I. B. Knott, S. Venkataraman and van, "Crypten: Secure multi-party computation meets machine learning," *arXiv preprint*, vol. arXiv:2109.00984, Sep. 2021.

[26] S. S. J. Ma, S. Naas and X. Lyu, "Privacy-preserving federated learning based on multi-key homomorphic encryption," *International Journal of Intelligent Systems*, vol. 37, no. 9, pp. 5880–5901, Apr. 2021.

[27] S. A. R. et al., "An advanced data fabric architecture leveraging homomorphic encryption and federated learning," *Information Fusion*, vol. 102, Feb. 2024.

[28] R. Nguyen, "Zero knowledge proofs in machine learning: A comprehensive guide," SotaZK, Oct. 2024, https://sotazk.org/insights/zero-knowledge-proofs-in-machine-learning-a-comprehensive-guide/.

[29] V. G. et al., "Experimenting with zero-knowledge proofs of training," in *CCS: Computer and Communications Security*, Nov. 2023, pp. 1880–1894.

[30] A. S. S. Antebi, E. Habler and Y. Elovici, "Tag&tab: Pretraining data detection in large language models using keyword-based membership inference attack," *arXiv*, vol. 2501.08454, Jan. 2025.

[31] "Levy," [Online]. [Online]. Available: https://dictionary.cambridge.org/us/dictionary/english/levy

[32] D. Thomas, "Start-up prorata.ai valued at $30mn after signing up uk publishers," *Financial Times*, 2025, [Online]. [Online]. Available: https://www.ft.com/content/c917a1e1-60a5-42c5-9158-6199f8a1f9ab

[33] J. Handy, "Looking closer: Suno ai," *Handy AI*, 2025, [Online]. [Online]. Available: https://handyai.substack.com/p/looking-closer-suno-ai

[34] "Suno revenue, growth, and valuation," Sacra, 2025, https://sacra.com/c/suno/.

[35] "Aiva technologies: Revenue, competitors, alternatives," 2025, [Online]. Accessed: Mar. 06, 2025. [Online]. Available: https://growjo.com/company/Aiva_Technologies

[36] "Amper music - overview, news similar companies — zoominfo.com," 2025, [Online]. Accessed: Mar. 06, 2025. [Online]. Available: https://www.zoominfo.com/c/amper-music-inc/414574746

[37] "Beatoven.ai - company profile - tracxn," 2025, [Online]. Accessed: Mar. 06, 2025. [Online]. Available: https://tracxn.com/d/companies/beatoven.ai/__GqrPyXIL1JWKE1nyCfOnii9m9JI0oifmlTTOCIewlI0

[38] ZoomInfo, "Boomy - overview, news & similar companies — zoominfo.com," https://www.zoominfo.com/c/boomy-corp/470700868, 2025, accessed: Mar. 06, 2025.

[39] "Music.ai: Revenue, competitors, alternatives," https://growjo.com/company/Music.AI, 2025, accessed: Mar. 06, 2025.

[40] "Spotify music dataset," https://www.kaggle.com/datasets/solomonameh/spotify-music-dataset?resource=download.

[41] J. Hutson, "The evolving role of copyright law in the age of ai-generated works," *Journal of Digital Technologies and Law*, 2024. [Online]. Available: https://www.lawjournal.digital/jour/article/view/486#:~:text=Therefore%2C%20a%20revised%20copyright%20framework,as%20a%20replacement%20for%20it.