



ARTIFICIAL INTELLIGENCE AND THE MUSIC INDUSTRY

PRESENTED BY:



CONTENTS

EXECUTIVE SUMMARY	03
01: DEFINITIONS	04
02: CURRENT TRENDS	05
03: TRAINING GENERATIVE AI	06
04: LEGAL CONSIDERATIONS & DEBATES	07
05: CHALLENGES & OPPORTUNITIES	11
ABOUT ONE STEP AHEAD	13

EXECUTIVE SUMMARY

Artificial intelligence (AI), particularly generative AI, is revolutionising industries, including the music sector, where there has been a lot of high profile discussion and debate in relation to AI that can compose and produce original music.

This technology, while impressive, is presently more of a tool that enhances human expertise rather than a standalone composer and producer of music. Several music generation services leverage a combination of AI and human input, using pre-recorded human-created loops for production. And even though significant players like Google and Meta are advancing in this field, current generative AI technologies face challenges in producing high-quality, comprehensive music outputs. Meanwhile, public perceptions of AI-generated music remain mixed, primarily due to the ambiguity surrounding what truly constitutes AI-created music.

Critical questions have also been raised regarding the legalities of music-making AI, and how AI models are trained and the licensing obligations of AI companies. The music industry is advocating for full transparency and proper licensing when copyright protected music is used for training purposes. Furthermore, there are legal uncertainties around whether AI-generated content can

enjoy copyright protection itself, with varying stances across different jurisdictions.

The emergence of generative AI that composes and produces original music creates both challenges and opportunities for artists, songwriters, record labels, music publishers and the wider music industry. Navigating this landscape requires continual adaptation, understanding and clarity to strike a balance between creativity, technological advancements, and legal fairness.

In this report, we review the latest trends in generative AI and consider how the technology will evolve, as well as outlining the legal questions. We explain how:

- Generative AI technologies, including in the music space, are rapidly evolving but still require human input for high-quality results.
- Some AI music platforms use pre-existing loops and samples manipulated by AI, blurring the line between human and AI creation.
- Google and Meta are developing their own generative AI technologies for music, but the output is a work in progress.
- Existing AI can produce short music snippets with limited memory hindering complex composition development.
- Copyright and transparency issues arise in training generative AI models due to the use of copyright protected music in datasets.
- Transparency from AI developers is crucial to address copyright concerns and shape future licensing agreements.

ONE STEP AHEAD: ARTIFICIAL INTELLIGENCE AND THE MUSIC INDUSTRY

Researcher: Sam Taylor, Chris Cooke | Produced by CMU Insights | Published by IMPALA in July 2023
IMPALA, Coudenberg, 70, 1000 Brussels, Belgium | +32 2 503 31 38 | info@impalamusic.org | impalamusic.org

ARTIFICIAL INTELLIGENCE AND THE MUSIC INDUSTRY

Artificial intelligence – and the use of AI by the music industry – is not new. However, AI technologies, and especially generative AI technologies, are rapidly becoming more sophisticated and much more widely used.

This has resulted in much more interest in AI, within the music industry and more generally, and the need to answer some of the questions posed by this technology now seems more urgent.

DEFINITIONS: “AI” AND “GENERATIVE AI”

Quite how ‘artificial intelligence’ – or even just ‘intelligence’ – is defined is open to some debate. Depending on your chosen definition, not all the technologies and platforms that are casually referred to or even formally positioned as “AI” necessarily are truly artificial intelligence.

Here are some useful definitions from Stanford University's Institute For Human-Centered AI...

“Intelligence might be defined as the ability to learn and perform suitable techniques to solve problems and achieve goals, appropriate to the context in an uncertain, ever-varying world. A fully pre-programmed factory robot is flexible, accurate and consistent, but not intelligent”.

“Artificial Intelligence, a term coined by emeritus Stanford Professor John McCarthy in 1955, was defined by him as 'the science and engineering of making intelligent machines'. Much research has humans program machines to behave in a clever way, like playing chess, but, today, we emphasise machines that can learn, at least somewhat like human beings do”.

In the music industry, most attention is currently focused on the application of a subset of AI known as generative AI, a technology with the ability to generate text, images, audio and video, and even mimic human voices.

Again, there is some debate around the precise definition of generative AI, and not all the technologies and platforms currently referred to or positioned as being generative AI necessarily meet a stricter definition. Or, they may employ AI as part of a process that generates original content, but the entire process is not AI-based.

Here is a useful definition from the technology company Nvidia, who are the leading manufacturer of processors used to power artificial intelligence computing...

"Generative AI models use neural networks to identify the patterns and structures within existing data to generate new and original content.

One of the breakthroughs with generative AI models is the ability to leverage different learning approaches, including unsupervised or semi-supervised learning for training.

This has given organisations the ability to more easily and quickly leverage a large amount of unlabeled data to create foundation models. As the name suggests, foundation models can be used as a base for AI systems that can perform multiple tasks".

CURRENT TRENDS

Various music creation AI platforms assert their ability to compose and produce original music. While the reality is often not as impressive as the rhetoric, some of these tools, when combined with human input, can generate interesting musical compositions.

That said, without more significant human input, many generative AI tools for music creation often produce substandard outputs, far from the high-quality, well-arranged music that listeners are accustomed to. In essence, while generative AI can execute tasks like creating chord progressions, due to their mathematical nature, more complex music generation is still a largely human endeavour.

Given those current limitations, many services claiming to generate music through AI are suspected to use human musicians as part of the process. For instance, Soundraw seemingly utilises pre-recorded human-created loops, which are then manipulated by AI.

Similarly, it seems likely that Boomy – which has received a lot of attention recently – is also using pre-existing loops, samples, instrumental excerpts and other audio which is then stitched together using various transformations, which may or may not be AI.

It's significant that – in a highly technical, highly competitive sphere where patents on technology provide a significant advantage – Boomy does not currently appear to have any patents for music generation and AI registered to it.

Lifescore, a product of Abbey Road, openly combines sound loops crafted by human composers and musicians using AI, as does Memu. Neither is pretending to be a generative AI tool for music creation, but rather they position themselves very much as using AI as part of a wider toolkit.

These current limitations also mean that tracks that are presented as "AI-generated music" are often created with a considerable amount of human input, leading to blurred lines between human and AI creation.

While it is important to acknowledge the limitations of what is currently possible with AI when it comes to music creation, it is also important to appreciate the rapid pace at which generative AI is evolving, evidenced by the significant improvements in the realism of AI-generated images.

The ultimate goal for many in the field is for generative AI to be capable of fully composing and producing music that could feasibly be released, and with a level of quality comparable to human-created music.

However, there are a number of challenges that need to be addressed in AI models in order for fully AI-generated music to feel authentic, including: the underlying composition; the arrangement and choice of instruments; and the performer's interpretation of the piece. In addition, production – how these elements are combined into a cohesive track – is a crucial part of the process.

Google and Meta are both developing their own generative AI technologies – and both have demonstrated generative AI for music creation. While their advancements are considerable, the sophistication and quality of their output remain a work in progress.

Existing AI can typically only produce short snippets of music and has a limited memory, restricting its ability to develop complex compositions that span longer time periods.

Additionally, a significant aspect of what makes music so special is the human element – the interpretation and performance that a human musician brings to a piece. Replicating this nuanced aspect of music is something that AI has yet to achieve.

We've put together an interactive document that gives some examples of what AI can and can't currently do, which expands on this discussion. **You can access that document here.**

TRAINING GENERATIVE AI

Training an AI to create music involves analysing different types of music files – such as WAV, MIDI or sheet music – to learn about composing and producing music. This process can involve the analysis of hundreds of thousands, if not millions, of tracks.

That, of course, raises some important copyright and transparency questions, which are discussed in more detail below. A critical point to understand here is that generative AI does not directly copy content but instead generates new content inspired by its training data. However, the training process likely does exploit copyright where a dataset contains copyright protected works.

This puts the spotlight on what music has been used to train the generative AI models to date, and whether permission was required to use that music and, if so, whether a licence was obtained.

Some of the early datasets used for AI music training, such as the 'Lakh MIDI dataset', contained music files that might be protected by copyright. While other datasets, although not containing complete music files, may contain abstractions of key elements of the music. It could be argued that the value of such datasets is linked to the underlying songs and recordings.

Both Google and Meta used the 'Free Music Dataset' of public domain music alongside additional resources to train early iterations of their generative AI for music creation – and there's no suggestion that this in any way infringed copyright.

Google supplemented the training with 5500 fragments of music created by human musicians.

Meanwhile, Meta has recently licensed a large amount of library music from Shutterstock – and OpenAI has done a similar deal.

However, determining what has been used to train different models can be challenging and requires transparency from the companies developing these models. Some players are providing some transparency, often through the academic papers they publish alongside their work.

But fuller transparency is arguably required, partly to assess whether copyright has been infringed in the past, and partly to inform how future licensing deals with the music industry might be structured and managed.

Despite the complexity, transparency in AI is possible.

For instance, Memu openly discloses the content used to train its AI model, and there are various other technical solutions – including some working with existing image creation AI platforms – that allow for greater transparency and attribution of sources. Furthermore, an AI's creative process can potentially be trained to note the influences from its training data.

LEGAL CONSIDERATIONS & DEBATES

The training and use of generative AI poses a number of important legal questions for the music industry, and the wider copyright industries. As these technologies become more sophisticated and much more widely used, the need to answer those questions is becoming more urgent.

Policy-makers in many countries are also now actively considering how to regulate the development and use of AI, including generative AI, which provides a forum in which the music and other copyright industries can push for the legal protection and certainty that creators and their business partners need.

In the European Union, a new AI Act is in the final stages of negotiation as of summer 2023. Although this act did not originally intend to regulate generative AI, the final draft does include some specific obligations for those developing and using this kind of artificial intelligence.

Meanwhile, policy-makers elsewhere are actively consulting both technology companies and copyright owners. For example, a number of Congressional sessions have been held in the US, and in the UK the Intellectual Property Office has convened a working group to consider a code of practice. Cross-sector trade group UK Music recently published a policy paper to inform that process.

While music industry companies and trade organisations will make representations on a country by country basis, a lot of that work is being informed by the globally focused Human Artistry Campaign. Launched in March 2023, this sets out a number of objectives for how generative AI should be used and regulated.

This campaign was originally conceived by the US music industry, but has been backed by music organisations across the world, as well as organisations representing creators and copyright owners beyond music.

The key legal questions and debates are as follows...

Does and should AI-generated content – including music – enjoy copyright protection?

Copyright law primarily seeks to provide human creators with control over the outputs of their creativity, either as a point of principle, and/or to encourage human creativity by allowing creators to utilise the control and protection provided by the law to generate revenue. Meanwhile, 'originality of expression' is usually a requirement in order for copyright protection to apply.

With this in mind, many have argued that entirely AI-generated content does not currently enjoy copyright protection. That said, there are some jurisdictions where protection may be provided for such works.

For example, UK copyright law includes a rule stating who the presumed first owner is for 'computer-generated works', which implies that those works enjoy copyright protection. Although some argue that the 'originality of expression' requirement would still need to be met and question whether that is possible if an AI fully generates a piece of content.

In the US, where works need to be registered to enjoy full protection under copyright law, the Copyright Office has had to decide whether it will allow AI-generated works to be registered.

Its current position is that fully AI-generated works do not enjoy protection and therefore cannot be registered. However, "AI-assisted" works can be registered. Which then poses the new question: what is the dividing line between "AI-generated" and "AI-assisted"?

Of course, copyright law could be revised to specifically provide protection for fully AI-generated works. But would that be desirable? Many people believe that fully AI-generated works should not enjoy copyright protection. Though in the context of the music industry, there are both advantages and disadvantages to AI-generated works being protected by copyright.

Does the training of generative AI require a licence?

The music industry is adamant that, in order to train an AI model with existing copyright protected songs and recordings, an AI company must first seek licences from the relevant copyright owners. And assuming the training process involves making copies of music files onto a server – which it usually does – it seems pretty clear that consent must be sought in this way.

However, some AI companies have argued that training activity might be covered by a copyright exception. Copyright law does provide a list of scenarios where protected works can be used in certain ways without seeking consent or securing a licence, which poses the question: could any one exception be inferred in a such a way to cover the training of AI?

In most countries, it's hard to see what current exceptions could be used in this way. Though in the US, where copyright law includes the related but wider and more ambiguous concept of 'fair use', that probably still needs to be confirmed via a test case in court.

Meanwhile, some AI companies have claimed that there are some jurisdictions where an exception could be inferred in a way that means it applies to the training of their models. And if that is correct, and those companies based their servers in those countries, that might mean they could legally train their AI without first seeking a licence.

The music industry – in alliance with the other copyright industries – is currently seeking clarity on this point, ie confirmation that AI companies must seek consent and negotiate a licence before training their models on existing music or any other copyright protected works.

There are already some test cases working their way through the courts. These are not music related, but could nevertheless set precedents of relevance to the music industry. And in addition to the litigation, the copyright industries are putting pressure on policy-makers in multiple countries to more clearly set out the obligations of AI companies in this domain.

If it does transpire that relevant exceptions do exist in a small number of jurisdictions, arguably creating a 'copyright safe haven' for AI companies, the copyright industries will seek to ensure that those companies will still be restricted from commercially exploiting their models in the majority of countries where there is no such copyright exception available.

From the Human Artistry Campaign

"The use of copyrighted works requires permission from the copyright owner. AI must be subject to free-market licensing for the use of works in the development and training of AI models. Creators and copyright owners must retain exclusive control over determining how their content is used. AI developers must ensure any content used for training purposes is approved and licensed from the copyright owner, including content previously used by any pre-trained AIs they may adopt".

From the UK Music Policy Position Paper

"The creator, or their chosen rightsholder, should be able to decide if and how they want to use their creative talent. This certainty underpinned by legal rights (copyright) should not be undermined by any exception to copyright or compulsory licensing during the input stage. Users need to respect creators' choice as baseline for any discussions".

How will the industry license the training of AI?

Assuming it is agreed that licences must be sought in order to train an AI with existing music, the industry then needs to develop licensing models that meet the needs of the AI companies and all the different stakeholders in the music industry.

In the main, the industry argues that such licensing should be led by copyright owners in the open market, ie no compulsory or statutory licences should be introduced that set the rules and rates.

Quite how this licensing will work does depend to an extent on what the commercial opportunities are. As always with music licensing, you have the added complexity that both songs and recordings need to be licensed, and that usually happens separately. Other questions to be asked include...

- Will deals be lump sum fee or revenue share based? Or both?
- How will any one catalogue be valued, given the popularity of a track doesn't necessarily impact on its value in the training process?
- What role, if any, should the music industry's collecting societies play? On the songs side in particular, are rights actually controlled by the societies being exploited?
- How will revenue be allocated to individual works in any one catalogue – and how will it be shared with artists and songwriters?

Can artists control the use of AI to imitate their voices?

Generative AI is also being used to create 'cloned vocals' that imitate the voices of specific artists. In recent months, there have been some high profile examples of this use of AI. This raises another legal question: is the specific permission of an artist required in order to clone their voice?

Obviously, if the AI needs to be trained using that artist's existing songs and recordings, then copyright provides a certain level of control.

However, the artist may not own the copyrights in their music. And, as noted, there is still some clarity needed regarding the copyright obligations of those developing and using generative AI models.

It is thought that the separate legal principle of publicity or personality rights may also provide some protection here – ie an artist could enforce these rights against any AI company that imitated their voice without permission, therefore requiring the AI company to seek consent and negotiate a deal before generating any music featuring cloned vocals.

That said, how publicity or personality rights work varies greatly from country to country. In the US, where the right exists at a state level, the publicity right does primarily exist to stop third parties commercially exploiting an individual's image or persona.

In some other countries, for example in some European jurisdictions, the right originates more in privacy law. And in some countries, like the UK, the right is not currently available.

With this in mind, the music industry needs to seek clarity on the extent to which publicity or personality rights can empower artists in this domain.

There is also arguably a need to harmonise this right to an extent around the world. And, of course, in countries where it doesn't currently exist, the industry should be campaigning for its introduction.

From the Human Artistry Campaign

"Performers' and athletes' voices and likenesses must only be used with their consent and fair market compensation for specific uses".

From the UK Music Policy Position Paper

"A new personality right should be created to protect the personality/image of songwriters and artists".

What are the transparency obligations of AI companies?

This is another big area of debate. The music industry is calling for AI companies to be fully transparent about what datasets they are using to train their models.

And there is also a call for companies and individuals who are using generative AI in the creation of content to declare that fact, so that fully AI-generated content is clearly labelled.

These transparency obligations arguably need to sit in law, and indeed the EU AI Act does have some provisions around the labelling of AI-generated works.

From the Human Artistry Campaign

"Complete recordkeeping of copyrighted works, performances and likenesses, including the way in which they were used to develop and train any AI system, is essential. Algorithmic transparency and clear

identification of a work's provenance are foundational to AI trustworthiness. Stakeholders should work collaboratively to develop standards for technologies that identify the input used to create AI-generated output. Content generated solely by AI should be labeled describing all inputs and methodology used to create it – informing consumer choices, and protecting creators and rightsholders".

From the UK Music Policy Position Paper

"It is important that in the input stage, the technology providers keep an auditable record of the music ingested before the algorithm generates new music. Music generated by AI should be labelled as such".

CHALLENGES & OPPORTUNITIES

While the rapid evolution of AI – and especially generative AI – creates many challenges for the music industry, it also clearly creates many opportunities for music-makers and their business partners.

In terms of the challenges, as discussed above, the industry needs clarity on the copyright and transparency obligations of the companies and people developing and using AI. It also needs to work on enhancing publicity and personality rights to ensure artists can protect their voices and personas. And it needs to develop licensing models that are practical and equitable.

There are also legitimate concerns about the commercial challenges – and increased competition – posed by ever more sophisticated generative AI. The technology will inevitably replace some of the people currently involved in the music-making process.

And high volumes of AI-generated music will likely flood the streaming services, making it even harder for consumers to navigate the catalogue and potentially taking money out of the royalty pool. Some are also concerned that AI-generated music will exacerbate streaming fraud, by making it easier for scammers to create large quantities of tracks with which to game the system.

However, the challenges are balanced by the opportunities. Some artists and labels are already employing AI as part of the music-making process, and these technologies will continue to make certain aspects of that process easier and more cost efficient. And that also applies to the creation of the large quantities of marketing content that is now required for the promotion of any one release.

Generative AI will also enable new kinds of music and fan experiences to be developed. Those will create both licensing opportunities for the music industry, and also allow artists to further enhance and monetise the fan relationship. Indeed, some of the biggest opportunities may lie in artists being able to offer superfans more personalised and interactive experiences, and labels should consider how they can work with and support their artists in pursuing those opportunities.

AI – both generative AI and beyond – will also help labels more efficiently manage their businesses more generally, including in areas like data analysis, and legal and business affairs. These developments will be particularly empowering for independent music businesses, for example allowing labels to analyse and manage data at a scale that previously would not have been commercially viable.

OWN
YOUR
FUTURE

We Proudly Support
IMPALA's One Step Ahead

15 YEARS OF
MERLIN

Find out more at merlinnetwork.org



ABOUT ONE STEP AHEAD

One Step Ahead is a digital intelligence facility from IMPALA to enable independent music companies across Europe to navigate and identify trends, developments, challenges and opportunities in the digital music market. Delivered in partnership with music business consultancy CMU Insights and with support from Merlin, One Step Ahead provides reports, tools and training to help independent music businesses of all sizes to get the most out of today's digital opportunities in Europe.



ABOUT IMPALA

IMPALA was established in 2000 and now represents over 5000 independent music companies. 99% of Europe's music companies are small, micro and medium businesses and self-releasing artists. Known as the independents, they are world leaders in terms of innovation and discovering new music and artists – they produce more than 80% of all new releases and account for 80% of the sector's jobs. IMPALA's mission is to grow the independent music sector sustainably, return more value to artists, promote diversity and entrepreneurship, improve political access, inspire change, and increase access to finance. Read more about IMPALA's story in 20 milestones on our website. impalamusic.org



ABOUT CMU INSIGHTS

CMU helps people navigate and understand the music business through media, education, research and events. CMU Insights is the company's business intelligence unit. It supports the music industry – both music companies and companies working with music – through training, consulting, research reports and conference sessions. cmuinsights.com



ABOUT MERLIN

Merlin provides independents the means to own their future. Members work directly with key digital partners, but benefit from Merlin's premium deals with over thirty digital music services around the world. Merlin's membership consists of independent labels, distributors, and other rights-holders, representing tens of thousands of labels and hundreds of thousands of artists from every country in the world, representing 15% of the global market share. As a member-led, artist-focused organisation, Merlin's mission is to deliver best in class services and work collaboratively with its digital partners to drive incremental benefits to members, while delivering value back to digital music services. merlinnetwork.org

