

Your Inside Track Newsletter

Description

Fifth Edition

Interview with Manaswi Mishra

Part II

By Don Franzen and Judith Finell

Don Franzen

A lot of the pending lawsuits are focused on using copyrighted material to train AI. Can you explain what is involved in training AI, and how the source material is fed into it? What's that process like?

Manaswi Mishra

Again, AI is a very broad term. I think within the scope of these lawsuits, we are specifically referring to deep neural networks, which is a specific class of more recent breakthroughs in AI. We have used algorithms to understand and generate music from a very early time. Even before computers, we were trying to make probabilistic descriptions of music like John Cage's process and chance-based music, dadaism movements, etc. So the idea of using an algorithm to generate ideas is not new. But in the context of this deep neural network that requires large amounts of data to learn patterns, we have these lawsuits because these large amounts of data have to come from an archive. And I think right now the legal arguments address questions like,

is it fair use to make copies and derivatives? At what point does the copy or derivative become completely different from the original? In what ways is it being stored? It's going to be very interesting to see how the litigation evolves, and what conclusions we come up with as a society. But as far as the scientific method, you need a collection of musical knowledge to learn these representations from. These representations might be basic ideas of music like, tonal harmony, or timbre of a certain instrument but also more complex ideas like a particular voice or genre. (analogous to learning words and grammar all the way to unique writing styles in language)

The lawsuits you mention are referring to foundation models, these large AI models built from a compendium of text/image/audio knowledge. A foundation model is something that has to learn probabilistic representations for all of language and all of sound and music in our case. So these foundation models have the following very real problems, are they trained on copyrighted music? Should they have gotten a special license beforehand or is it a problem in the way they are made available for people to use? But to draw a contrast to foundation models, artists and musicians are also able to make smaller tunable AI models and systems that they train on data sets that they have access to, and have copyright licenses for. Every musician has recordings, every production house has samples, a collection from which they can build certain AI models as well. It is going to be interesting if the future law makes a distinction between foundation models and who has the license to train a foundation model, perhaps it's going to result in new models of authorship and sampling.

It will be important to make a distinction between such a foundation model, which can only be built by a large corporation or community with access to a large amount of data centers and resources, versus individuals who may make their own smaller AI models.

Judith Finell

Say a big company like Universal or Warner Brothers that owns a large music catalog, if they wanted to control it, maybe they're already working on this, couldn't they create a blocking mechanism that doesn't allow it to be fed into AI without permission?

Manaswi Mishra

This would be great actually, because it doesn't exist right now, this kind of technology to fingerprint, watermark, control and restrict the flow of data. These have to be innovative solutions because today it's very easy to make digital copies of data and extract out watermarks. We are actually seeing the beginnings of this with images. There are these open-source methods that are being developed where you can take a piece of an image and apply an algorithm to it so that

it continues to look pretty much the same to the human eye. But when the algorithm looks at it, it messes up the information a little bit in order to protect your media and piece of art when you don't want it to be ingested, at the point of creation. Which is super powerful and maybe similar ideas could be explored with music as well. Such authentication techniques don't exist as a standard right now, but it's a great way to encode the opting in and opting out that we will definitely need in the not-so-distant future.



Every piece of media that is created would need such a mechanism for authenticity and consent, encoded within the media, moving forward.

Judith Finell

It does seem inevitable, right? If they want to be considered a financial partner in the use of it, then they'll have to control it.

Manaswi Mishra

Yeah, that would be wonderful actually, if we could do that.

Judith Finell

Well, maybe not everybody, but all of our readers would probably agree with you that it would be wonderful.

Don Franzen

All the readers who own copyrights in songs or recordings would agree that it's wonderful anyway.

Manaswi Mishra

That's right. Having an extra layer of opting in and opting out, which is in control at the source of the origin, where the artists, as they create, are able to encode their consent and identity is a great incentive. This is crucial for the protection of both the individual artist and catalog owners. Whether you decide to opt out or you are like me, I might want my music to be used in the next training of models. If it offers a new revenue stream, whether I use AI in my artwork or not, I would want to protect my bespoke style and ideas. I also imagine this opt-in/opt-out might look much more interesting, where one might say that I want my music to be used in the training of the next big foundation models, but I refuse to use it in these specific contexts like a political rally. One possible future is where we build techniques to encode consent and fingerprints into pieces of music.

This would be an incentive for not only catalog owners but also individual artists and musicians to future proof their creations¹.

Judith Finell

Right, in terms of moral or affiliation considerations that represent philosophical differences between some parties.

Manaswi Mishra

Right.

Right now, how catalog managers find and compare samples is, they have a range of signal processing and information retrieval methods to look at the sample and compare it against their collection of samples. That's how the royalties get collected. That technology for fingerprinting when extended to AI-derived content will look different now in the coming years.

Don Franzen

This leads to the next subject we wanted to talk to you about, which is, in what ways do you think AI offers opportunities for enhancing creativity among musicians?

Manaswi Mishra

This is something I'm really passionate about because I think it is not obvious what new ideas look like before you create them. The first results of any technology are to mimic and copy what already exists, and that's what we're seeing now. Everybody is trying to mimic a style, copy a canon, but this pastiche will soon give way to making completely new ideas in music. I can imagine it spawning in many different ways.

Music could today be thought of differently for the medium we consume music through, in the same way, we have access right now to musical ideas from any time period in the world and from any country, through our streaming services. Our future musical AI tools could act as portals for quoting ideas and mixing sonic concepts together from different time periods, different parts of the world, and doing it in totally novel ways. We are showing some examples of this already in our Operas where we have incredible musicians react and respond to different genres and musical worlds, allowing one to collaborate and perform with ideas from a certain era or decade².

We've been testing it out in different contexts for non-professional musicians as well, through our communal AI sonic sculptures that allow anyone to compose and mold music by listening and manipulating knobs and levers. We have also been experimenting with newer ideas of shared authorship in a fictional AI radio context where an individual is allowed to collaborate and reconstruct a constant stream of generated music³.

It seems if music has to exist in a digital world, it doesn't have to be exactly three minutes long, the size of a CD, which I don't know why we need to stick to today, maybe it doesn't have to be sampled at 44.1 or 48 kilohertz or doesn't have to be a fixed form of instrumentation.

I think it is important to show many such examples because in talking about it, it's not clear what a new sound or music is, or what it means to mix country music and Indian classical music, ideas about authorship, who gets to do it and also why! These are emergent properties of core musical ideas, the communities that create and grow from and the fundamentals of now.

Judith Finell

I would appreciate your explaining how is that different from just exposing the musicians to Indian classical music or country western? Are you saying that the AI is going to be exposed to it and blend that in? How does that apply here to the connection with AI that you foresee?

Manaswi Mishra

It is similar in many ways to exposing oneself to different identities and different styles. It is similar in concept but not in practice. For example, one might make music that has key ideas and elements that have been composed, and then based on the person who is listening to it, the ideas take different forms in every instance of the playback. A future artist might release an album that produces an infinite amount of generated music from a fixed set of composed sonic ideas and sounds different every day.

Judith Finell

Is that what you mean by changing, having a flexible format length, it would change each time?

Manaswi Mishra

That's just one way of imagining it, definitely. It does enable a bunch of novel ideas because you can now suddenly get portals to different periods of time, different musical languages, and also to your own material. One specific way in which the AI makes it different from a human scholar who exposes oneself to many pieces of music, is that we can now quote an idea without actually exactly copying it. Many different musical contexts have this kind of recontextualization – like sampling in hip-hop, or quoting phrases and licks in jazz, reharmonizing standards, etc. Music is all about the human context around expressing ideas – new and old. I think AI tools of the future will have a big role to play in reinventing this form of quoting and contextualizing ideas. AI might unlock some new ways of implementing it, of course, by making sure that authorship and credit distribution are done the right way, which we talked about earlier, and that's a big hurdle.



I also want to mention another powerful idea I foresee. One more viral and contentious form of AI music spreading around the world is voice mimicry. When done right, with due credit systems in place, I see this as a powerful new form of fan fiction – a kind of intricate audience and artist interaction. The music that they make, with the raw ingredients from an album or AI models that the artist has provided, might unlock completely new revenue streams for the artist, as well as new ways fans and artists can interact. This could be very powerful in the way it manifests.

Judith Finell

That's a really interesting idea. So you see that this gives greater fluidity and flexibility to music than ever existed before. Is that what you believe?

Manaswi Mishra

I hope so. Before there were streaming services, Napster was a big controversial leap that caused a massive disruption, but it did allow people to now listen to music from around the world. Looking forward, you can imagine AI will disrupt the world of who creates music, what we value in the arts in an abundance of media, and our creative processes.

Don Franzen

You've explained some ways that AI could kind of transform both the creation of music and the enjoyment of music. So we see positive things there. But what are your concerns? What is there in the development of this technology that also concerns you and that we should be thinking about, perhaps even apprehensive about?

Manaswi Mishra

The first one, of course, is what we discussed earlier, a need for new forms of authorship. That one is key, on who's creating the piece of art, AI models were used, and a kind of provenance, that's a challenge. How do we know if my style and my music were used by someone else in another context? That one's a challenge that everybody is interested in addressing, so I'm looking forward to ingenious solutions. It's clear why that's the first thing that people are addressing. But another challenge that I see that is not being addressed and might go slightly under the radar is appropriation. These AI models I mentioned can model a certain instrument, a certain culture, a certain genre, and a certain slice of reality based on the data that it is trained on. What then, is our responsibility as we use it?

In many ways, these AI models are blocks of culture whether it's for a community or genre, or whether it's the identity of an individual. It is not much different, in my opinion, from the national anthem or the national flag of a country where you have this representation of a historical and ethnographic context. At an individual level, this representation is being fought with the right of publicity, especially with the example of *fake Drake* songs. And we observe UMG also trying to take down unlicensed AI generated music, rightly so, because identities and IP are being misrepresented.

This might happen not just at an individual scale. I think it can happen even with a culture where a lot of people might be upset that I used indigenous music in a way that was not respectful of it.

Judith Finell

Have you observed, or can you predict there could be abuses of it in terms of impersonating an artist or a community's indigenous music?

Manaswi Mishra

Absolutely. I can imagine it being used in ways that certain people who have a context and respect for that culture will not be happy about. At an individual level, we can already see it manifesting where people's voices are being copied and therefore you can make them sing things that they may not sing otherwise. Grimes, for example, had her AI voice model that people were allowed to use, and she made sure to filter out the ones that use obscenity or talk about things that she would disagree with. It might be harder to do that when it's more for a culture where there is no one spokesperson to represent and protect.

Don Franzen

That's an interesting issue that has come up. The whole issue of cultural appropriation and whether indigenous music or traditional music and culture, or traditional visual art for that matter, should enjoy some kind of protection. Not so much in the United States, but some Latin American countries have actually enacted laws trying to protect indigenous art and music. But you bring up both dimensions, both the traditional Western copyright concerns, and then also a non-traditional cultural appropriation concern. I think both of those will have to be grappled with as AI develops and embraces more styles of music.

Manaswi Mishra

Absolutely. We absolutely need a diverse set of voices, cultures, communities and individuals to have access to the mathematical ideas of AI and help shape the tools and sounds of this new creative age.

Part III of the interview continued in the next issue.

default watermark



Manaswi Mishra is a current graduate researcher and LEGO Pappert fellow at the Opera of the Future research group, MIT Media Labs. His research explores strategies and frameworks for a new age of composing, performing, and learning music using AI. He joined the MIT Media Lab in 2019 and completed his MS in Media Arts and Science, developing his work “Living, Singing AI,” to democratize the potential of AI music making with just the human voice. Prior to joining MIT, he has received a master’s in Music Technology at UPF, Barcelona and bachelor’s in Technology at the Indian Institute of Technology Madras. He is passionate about a creative future where every individual can express, reflect, create, and connect through music. Manaswi is also a founding instigator of the Music Tech Community in India and has organized workshops, hackathons, and community events to foster a future of music and technology in his home country.

Footnotes

Date Created

September 13, 2023

Author

don-franzen

default watermark