Your Inside Track Newsletter

Description

Seventh Edition

Interview with Manaswi Mishra

Part III

t watermark By Don Franzen and Judith Finell

To view Parts I & II of Manaswi Mishra's interview, published in our fifth edition, click here.

Don Franzen

So far, the Copyright Office and the courts are taking the position that purely AI generated art, including music, is not protected by copyright if there's no human involvement. So, how do you see that rolling out? Is that a real concern or is there always going to be some human involvement in the creation of music so that there can be a workaround to still get copyright protection?

Manaswi Mishra

In the Copyright Office guidelines from March 2023, they have mentioned that if there is sufficient human expression involved in addition to using the AI tools, the creative work is certainly considered copyrightable.1

If the early examples of AI image generators (Midjourney and Stable Diffusion) are examples to go by, the first thing that anybody does to get AI generated creative work is to enter a text prompt and hit the generate button. Now, when the cost of creating an idea is so cheap, one is able to create thousands and thousands of ideas and then select the ones that they want to represent their idea, right? Arguably, this is similar to hitting a slot machine. Would the human act of selecting and the original text prompt be sufficient creative expression? While this is just the first example of how one might use generative AI, if you are an artist, you want to uniquely identify your style, right? You want control over

the output to realize your artistic vision.

You want to develop an identity aesthetic, you want to make that context be the reason your art has value, and why your fans respect and love your work. I already see artists trying to use it as a tool within their arsenal of music making or image making skills. Therefore, the AI parts blend in with many other aspects of their workflow. We see that blending because AI tools are also getting deployed within existing software that people use for creative work (Adobe tools like photoshop, Music recording software plugins etc.).

The onus is currently on the artist, because an artist is going to go to the Copyright Office and say, "I want my work protected." They are going to have to describe which parts were AI-generated and which parts were human made. It would be great if technology, as well as the law, supported making these kinds of claims: like, if there was an obvious way of recording provenance where you can get a receipt along with your piece of music when you export it. If one could indicate in this receipt: "these were the Al models used, these were the non-Al parts used" it would make it simpler to go to the Copyright Office and say, "Here's my music. Here's the standard form of representing all the edits I have done. Here's how much AI had a role to play."2

The Copyright Office would also be happy with that. So, technologies that would allow a recording of what parts were AI, what parts were human, will benefit both artists and the Copyright Office. The initial ruling that purely AI generated works are not copyrightable may not have a huge role to play.

Don Franzen

ault wate That could be the exception, not the rule. Is that what you're saying?

Manaswi Mishra

Right.

Don Franzen

Let's say a musician comes up with a little melody, like a one measure melody, and then asks AI to generate some chords or harmonies to go with it. Now he's got a melody, he's got the harmonies. Does he have to say to the Copyright Office, "I wrote the melody, AI wrote the harmony"? What if he adjusts the harmony a little bit? What if he changes a note or two in the chord structure? Now, is it his, and not Al's?

Judith Finell

I was thinking about this. Say if you did submit it to the Copyright Office with sort of a breakdown of the elements that were AI created. It could have three columns: "AI created," "human created" and "combination, extended from the AI."

Don Franzen

Hybrid.

Manaswi Mishra

AI assisted?

Judith Finell

Yes, hybrid. But the Copyright Office is fairly agnostic in the way it looks at this, and it doesn't delve into the originality of materials; they've never seen that as their role. So, I'm wondering instead if they will start to see it as simply a tool, the same way that you would have used a digital workstation as a production tool. It could be seen as really enhancing what you've created by the various sound effects and timbres and orchestration selections, and even the note generation after your first measure as Don was suggesting.

So how do you think the Copyright Office will look at it? Will they require more documentation so you must prove the human contribution, or will they instead see this as a tool like any instrument or recording production possibility?

Manaswi Mishra

I know it's super important to future proof ideas, especially when you're talking about rapidly developing technologies. So just referring to them as a tool, rather than making an exception to what an AI tool is, might be in their interest. Perhaps if a certain creative work makes a lot of money and somebody comes along later and says their software was used in it, a case-by-case judgment might make more sense. There are already so many pieces of music which have used the preset of a software or a keyboard that has a demo button. You press the demo button, and it produces a sample and people use that directly in their music.<u>3</u>

Judith Finell

You can already command it to create permutations and variations of the melody you've created, right?

Manaswi Mishra

Right

Judith Finell

You might say, give me five permutations of this one bar melody that I created. Right?

Manaswi Mishra

Right, right.

Judith Finell

So, who wrote those permutations?

Manaswi Mishra

In many ways, it's similar to if you start writing a story and you use ChatGPT to write your story. I bet there's a big difference in telling ChatGPT to just write the story and then it writes the story, versus using it in a creative way where ChatGPT can play a character, and I'll write down the roles of other characters, and they kind of interact. Some of the dialogue is written by AI, but it's a tool within my creative expression. That could be one of many completely unique ways of generating permutations.

What we are talking about here is asking ChatGPT to give me five romantic ideas, then I pick the one I like and ask ChatGPT to make variations of it. In all these different ways of expressing in 2023 and beyond, I think transparency is the key. I think artists would want it so that they get credit where it's due whether they use AI or not. Artists might also want to show to their audience that there was no AI involvement in their work.

The Copyright Office's work will be easy if there's more transparency, and if they can really compare works on a case-by-case basis in the future. It'll be interesting to see what they do, especially since this is happening at the same time as the technology is evolving, so it's not clear which are the best atermark techniques and standard practices around it will evolve to be.

Judith Finell

Manaswi, you're a creative artist and you're a scientist. You're a blend actually of these different worlds we're discussing. As you describe the present protocol of obtaining copyright protection and getting permission to use something from the big library of music that was written before today, hundreds of years back, do you see copyright as becoming obsolete because it interferes with creativity? Or do you see it as necessary as the only way that will incentivize creative artists to in fact create?

Manaswi Mishra

I am not sure if I have all the information to answer this question. I have desires as an artist and as a researcher. Even if AI musicians of the future use many AI tools to create an idea, they would want their idea to be copyrightable. Music copyright has many different layers to it, right? The sync licensing and the mechanical licensing are just some variations of use cases. If I'm an artist and I have made a piece of music, I should benefit from that. If that music is heard a million times, that should proportionately benefit me somehow. If that music is used in training of several AI models, that should also benefit me. It makes sense to keep a record of the music authorship, so it incentivizes an artist to create because they know that they can benefit from their work. We live in an age where the boundary between artist and audience has almost disappeared. You can just post music directly to TikTok and have millions of followers that are directly listening to your music. So, I'm not sure if copyright ideas are obsolete. Arguments could be made that we don't need copyright. But like I said, even the AI artist wants their ideas protected.

Judith Finell

That's a very good summary.

Don Franzen

Well, that leads into an area that we were hoping you could discuss. There are at least three class actions pending right now in the United States. One is against GitHub, Microsoft, and OpenAI. Another one is against Stability AI, Midjourney, and DeviantArt. And a third is against Stability AI, mostly having to do with the use of copyrighted visual works as training data. So that's a very, very complicated area. I wonder if you could comment on that. In other words, both from a technical standpoint, what does it mean to use material like this for training data? And what are your thoughts on whether that might in fact be an infringement of the owner's copyright?

Manaswi Mishra

These lawsuits have the potential to guide the regulatory frameworks for what happens regarding generative AI, so I've also been following them with a lot of interest. The ones against GitHub, Microsoft, OpenAI are for coding and the others are against Midjourney, DeviantArt by Getty for images, so they are asking slightly different questions. While the former is more about the copyright of the "outputs" generated by the system, the latter is more about the copyright of the "inputs" to training the system.

Programmers have been using Github as a service to share, store, update and keep track of code for many years. The GitHub co-pilot is a generative AI model trained on this large collection of software code and can generate new code. But Copilot generates fragments of code that are very similar and in some cases even copies of code from its training dataset. While some code is generic and fundamental to programming, at other times it can be more sensitive and private

As you can see, this is about outputs and how similar are ideas in the outputs to examples in the input data sets, which is going to be crucial for any new generated media. If there is music generated from AI and then there is a certain phrase, which is exactly the same as something in the data set, then how does current copyright law deal with it? It has always been a case by case with the jury historically.

It is important to note that sometimes a tiny phrase of just three notes might be the core idea that has to be protected, like the opening notes to Beethoven's *Fifth Symphony*. That's just three notes, but it's a very recognizable idea. Sometimes longer chord progressions might be copied, but those are the fundamental building blocks of music, and therefore arguably shouldn't be copyrighted from being reused: that aspect of copyright protection is super important for music when evaluating what ideas in the output are similar to the input. Already, I see that the Co-Pilot and GitHub platforms are continuing their services. They haven't stopped their business. Their services still exist, and we can use them.

A novel way forward could be if generated outputs also include citations to examples in their input data that was used for the current generation. Some language models are already citing sources within the training dataset along with generating an output. I think that's a great technological solution where, along with producing the idea, the model also quotes citations that indicate which pieces of music were closest and were used in coming up with those ideas.

Judith Finell

That's fascinating.

Manaswi Mishra

But it's not so obvious because sometimes the idea that a scale degree 2 and 5 is followed by a scale degree 1 has happened in so many pieces of music. You can't really point to a certain music and say, "This was the music that contributed to this idea." So, we are going to see a hierarchical breakdown of units of ideas, whether it's a single word or phrase or substantial parts of a work. Similarly, in music, it would be notes or chord progressions and transparently citing sources that is going to be crucial. The people who build the technology would also want to be able to provide for such an ability. Companies that are building foundation models for music also need to protect themselves from this copyright infringement.

When someone can say "rock music" and write a prompt, and then press a button and get some rock music, the generated music and the service has to make sure that it is not incorporating phrases and elements that are identifiably copyrighted.

Judith Finell

....ean?san? Yeah, there have to be safeguards auth Watermark Don Franzen

I'm glad you differentiated between the GitHub case which has primarily to do with copying code, using code, whereas the other two cases have to do with using the actual visual works as training data. So, what about that? That's a different set of issues.

Manaswi Mishra

Absolutely, and that one deals more with input. Previously, we talked about more of the output: is the output most like something in the dataset? The case of Stability AI, Midjourney, and DeviantArt talks more about input, which essentially asks: is it legal for a company to ingest copyrighted art that is just out there for public consumption? Do I have the right, at the input level, where I can opt in, opt out, even get credit for this foundation model training on my work? This one's going to be interesting on how it develops. We are already seeing that different countries are treating this issue differently. Japan put out a statement about allowing any kind of input, copyrighted or not, to train such models. 4They said that they are going to focus on the outputs and look case-by-case to see if a certain output infringes on copyright, but they're not going to make restrictions on input. Europe and the U.S. are trying to make sure that inputs are also licensed, which makes sense because places that own large data sets, like Getty Images, want some say in their entire catalog being used.

Don Franzen

Right.

Manaswi Mishra

And I think with music, that part is also crucial.

Don Franzen

Sure. Let's say that music publishers are the analog of Getty Images, right?

Manaswi Mishra

Right.

Don Franzen

All the music publishers hold vast amounts of copyrighted material. They don't want to see it used as inputs for training models, necessarily.

Manaswi Mishra

termark Well, right now they don't I think because they are obviously seeing it as a thing that they don't have control over. Different companies are building models with different financial motivations and priorities in mind, but I can imagine that in the near future each of these catalog owners will want to train their own models, and release it in some form for people to use.

So, it seems to me that regardless of what decision is made, we are going to have foundation models trained on catalogs of music. Whether each catalog owner decides to do it in their own way and say: "we are not going to share our foundation models with you" or, if people decide to join forces together and make a giant foundation model, it's going to look similar. I think the first reaction has been to not allow the training of foundation models, but I can see the next steps would be for the catalog owners to figure out how it could become a revenue stream.

Don Franzen

Well, that's really interesting. It's kind of like the labels' first reaction to digital music was to try to make it go away, through lawsuits and what-not. But ultimately, it became the savior of the recording industry when they finally embraced first downloads and then streaming. And now the recording business is back up to its old levels, or even beyond.

Judith Finell

They had to create a licensing model as sampling became popular in hip hop and they built that whole catalog of music. They had to create a licensing model because the behavior wasn't going away.

Don Franzen

That's right.

Judith Finell

Do you think that they will have to license the material that they're training from?

Manaswi Mishra

Well, it makes sense that catalog owners will figure out some revenue mechanism. Either they build it themselves and collect revenue every time their models are used, or if they rent it out to Google, OpenAl or a new company, they produce a deal. There's going to be new revenue streams for catalog owners. I think they have the power and the legal teams to figure that out. Artists as individuals do not have the legal teams and the power to create new revenue streams like the publishers. Hopefully publishers will work with artists in laying this groundwork.

It's going to be interesting to see how it develops, but we can't wish it away. Because, right now, the catalog owners are trying to take the sounds off or are threatening Spotify and Apple Music to not share their catalogs, which totally makes sense to me. The world of open source works in a decentralized fashion where people are just excited for these models. They will just stream it in different places, and if they don't have legal ways of doing it, it's just so easy now to make copies digitally. I can screen record something, I can record audio on my microphone. lefault

Don Franzen

Well, that leads us to the last thing we wanted to ask you about after talking about all these concerns. You are a musician, an artist yourself. How are you planning to use the tools of AI in the work that you do?

Manaswi Mishra

I love this guestion [chuckle], because it helps me talk about the things that I'm actually working on in my day-to-day life. What I'm interested in is what musical instruments will look like. Every generation has its own kind of music because music has many functions, but one of them is to connect us to ourselves, to each other, to the world around us.

Every generation has its own kind of music, and the musical instruments are also changing every generation. There's a very famous piece that goes, "Roll Over Beethoven, because it's not your time anymore." The same thing applies to every generation. The musical instruments of the past were made from the technologies and resources that we had access to. We had bone flutes, goat skin percussion, and then soon we had electronics, and now we have apps on our phone and digital workstations. Almost all music today is made at some stage on a computer.

It is only natural for me to think that musical instruments are going to change in the generation to come when people are going to have access to large datasets and AI models. My work involves thinking about AI as a musical instrument. An instrument is very personal.5 It allows you to express yourself. The early examples of AI media that we are seeing, where you enter a text or press a button to get art, do not really feel personal. I might have this black box and say, "this is an AI machine that makes

beautiful music" and press the button, or you might press the button. I do not think there'll be much difference. You would not feel like you created the music. I think the musical instruments of the future might have a different set of buttons and knobs and strings and ways for you to express your own ideas.

What I'm doing is building prototypes of these musical instruments, thinking about which ideas stick, which don't. I've been able to make these musical instruments for myself as an individual, and I do performances with these instruments. We have also been able to make experiments of new AI musical instruments with large orchestras, through Tod Machover's operas 6 at the MIT Media Lab, and test out what a musical instrument might look like when an expert musician improvises, presses the buttons and knobs, creates music that they really feel like they own.

We have also been experimenting with musical instruments that can be performed by hundreds of people (like a sculpture) or just single, average people who have never thought of themselves as musicians. Maybe, these tools can allow them to make music, too.

Judith Finell

Do you think that the performing virtuosos of the past, the opera singers and the instrumentalists, are going to fade into the past because the sounds will be controlled by the composer, who will create all of the sounds within the studio environment and with AI enhancements? So they don't need a violinist to read a score and determine the bowings, or anything like that. Where's it going?

Manaswi Mishra

lefault Well, I think that's an interesting way to look at it. How is music made and who are the contributors? There's the composer, there's the session musicians who are doing the recordings, there are mixing and mastering engineers who come later, and then there are distribution channels. There are AI tools involved in each of these steps. There's AI in the composition, there's AI for recording, and like you just said, there's even an AI violinist. The most successful early examples of AI have been in mixing and mastering because that's an expensive process that not too many people have access to. I don't have access to a string orchestra, but I would love to record with them.

Al will augment each of these individual steps in the process. If you are somebody who has a skillset, you will fill the holes where you don't have the skills and resources with AI, perhaps. I don't think any individual portion will disappear completely. Composing will exist. Recording will exist, just as guitars and pianos still exist, even when we have digital sampling methods for them. I don't think AI replaces any part of the composition and recording process, it just augments it. If someone doesn't have a certain skill, they might rely on AI because it's a cheaper tool that they can access.

Judith Finell

But you won't be giving a score to a string quartet and saying, "play this piece." How are you going to communicate to a third party, meaning the instrumentalist, when you've created all the sounds already? I am trying to understand what that lifespan would look like in a performance.

Manaswi Mishra

Right. In some ways generative AI could play the role of translating ideas from one expert to another.

It could be exactly translating a composer's ideas into a score notation for the violinist to play. In performance, this could mean, composing and improvising in real-time ways not possible before. It could also be an inexact translation where musicians interpret the composition in their own way. I use algorithms to give score and performance instructions to many live musicians. They perform it, they play it in real time interpreting the instructions, but I'm composing on the fly. These are just some examples of how creativity might look different in the age of AI.

Each musician could use AI to assist in expanding their creative imagination and vision. Music is the best thing because it only exists in context. There must be somebody who wrote it; some human context around the stories that were told, why the notes were chosen. There must be humans who listen to it and agree or disagree to hate or love it. They may become fans of it. The more we have this AI talk, the more the value of live performance is going to go up. I would love to see a live musician play a musical instrument much more than I did before perhaps, because I'm now exposed to so much AI generated media. I really think popular music is called "pop" only because you're able to share that context with other humans. It wouldn't make sense if all the music was just an AI generated hyper personalized solipsistic feed. I think as individuals and as a society we will use AI in processes, situations, and performances that expand the definition of art. In this pursuit, we must cultivate and vigorously protect our ability to create and imagine beyond our tools.



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

© 2024 Your Inside Track[™] LLC

Your Inside Track[™] reports on developments in the field of music and copyright, but it does not provide legal advice or opinions. Every case discussed depends on its particular facts and circumstances. Readers should always consult legal counsel and forensic experts as to any issue or matter of concern to them and not rely on the contents of this newsletter.

Date Created June 5, 2024 Author yourinsidetrack

default waterm