Allison Jerzak: Hi, my name is Allison Jerzak. I'm currently a PhD student in Musicology at the University of California, Berkeley, where I study the history of music recommendation. In particular, I'm interested in the digital infrastructure underlying the earliest music recommendation systems. I recently spoke with Nick Seaver, a cultural anthropologist from Tufts University. As he puts it, he studies how people use technology to make sense of cultural things. His recent book, Computing Taste: Algorithms and the Makers of Music Recommendation, explores music recommendation companies in the United States in the 2010s. I wanted to speak with him to gain a greater understanding of how developers of music recommendation systems were thinking about music, and how their knowledge and beliefs about music shaped how they developed these systems. This is our conversation.

I'm talking to Nick Seaver and he is an Assistant Professor of Anthropology at Tufts University. Broadly, I would say Nick studies the anthropology of data. His first book, *Computing Taste: Algorithms and Algorithms and the Makers of Music Recommendation* is really about the people behind music recommendations and how they are thinking about music and actually programing it into these systems. His more recent work is on attention and how attention works in socio-technical systems. To start, I'm curious about how you became interested in this topic generally — either music recommendation specifically or the broader idea of the anthropology of data.

Nick Seaver: Well, first, thank you for having me. The story about how I got into music recommendation as an object of study is, I don't know, maybe boring, but it's my story.

I've always been, even as an undergraduate, interested in the relationship between music and technology, and how the way that technology works in music undercut a lot of ideas that people frequently have about how music is uniquely human and how things that are uniquely human are not technological. We think of music as being expressive, subjective, all of these things. And yet music is also extremely technological in all its forms.

Before I started working on music recommendation, I was doing a master's program in media studies and working on the history of the player piano, where I was also interested in how people thought about the collision of technology and musical stuff, but in an earlier historical moment. I got to music recommendation by thinking — you know, I want to do a project that's about something that's happening now. What something that's happening now? Music recommendation is happening now. And that was it. So I got I got in there fairly arbitrarily, and I turned into an anthropologist of data and algorithmic systems, but I certainly did not set out to become one. It just sort of happened along the way.

Allison Jerzak: Did you notice any historical differences between how people were thinking about the player piano and how the relationship between music and technology shows up in algorithmic systems?

Nick Seaver: Oh, that's a great question. One of the big things that I was interested in during my player piano work was trying to recover some of the historical plausibility of the player piano. I think it's easy for people nowadays to think about player pianos as sort of always-already obsolete, like we knew that they were never going to work, and that people back then must have known that they were going to not be the thing. I like to ask people: if you imagine in your head the sound of a player piano, you probably imagine a piano that's a little bit out of tune because that is how player pianos exist in the cultural imaginary nowadays.

In any case, I was really interested in how people at that moment might have thought of them as being viable or interesting. And one of the things that I wrote about when I was working on the player pianos (and again this is ages ago now), was how they were trying to think about what was distinctively human or distinctively musical in relation to these technologies.

Many player pianos had an interface where you could control some things — you couldn't press the keys, but you could control how fast things went; you couldn't control the dynamics. There were all these like advertisements and pamphlets where people are saying, well, you know, the true seat of musical expressiveness is actually in — surprise — tempo variation and dynamic variation. The notes themselves, eh, not so important. I think what you see today, sort of unsurprisingly, are similar efforts to think about what's human in and around these systems. Especially, I would say, with more recent stuff like generative AI, to start to imagine what the human is again in relation to what these systems can and can't do.

Allison Jerzak: Moving in the direction of your book a little bit — where and how did you do your ethnographic fieldwork for the book? It sounds like you interned at a couple of different music recommendation companies. Can you outline what that was like?

Nick Seaver: One of the challenges of anyone who wants to study these kinds of systems is that it's very hard to get access to companies in particular. It was true then; it's probably even more true now. Companies don't have a lot of incentive to let outside researchers in.

And so what I did to start was I started going to conferences. I did a lot of field work at conferences where people were working on recommender systems or working on music informatics. And I would go to those regularly thanks to a little bit of funding that I had got from my university, and I would start to see similar people at all the conferences I was going to. And eventually one of those people said, “Hey, do you want to study us?” And I was like, “Oh, thank God.”

This is actually a very standard anthropology fieldwork story — you find one person who thinks that you're worthwhile, and they can invite you in to whatever it is that you want to study. And everyone else who used to, you know, not answer your emails will talk to you now because you're associated with this guy. So that's what happened to me. I was able to do an internship at this music recommendation company that I call Whisper in the book. It was that, and it was supplemented with a bunch of fieldwork at events, like music hackathons. I lived in in New York and San Francisco for a little bit, during the course of this as well, to see what was going on in those local scenes.

But it's a very challenging thing to sort of get into one specific spot. One way to deal with that is to think about what exactly it is that you want to study. You might say, I want to get into a company because I want to know how they do things. But, if you don't want to know about the details of one particular company, but you might want to know about sort of patterns of thought across a whole sector, you can actually find out a lot of that kind of stuff without getting inside of one particular company. You don't see everything you might want to see, but you can actually find useful stuff if you don't imagine that everything that's useful is by definition secret.

Allison Jerzak: Was there anything that surprised you about how people who work at these companies think about music or talk about music?

Nick Seaver: Yeah, there's always weird stuff going on. I mean, some things you could probably imagine from the outset. One thing that I talk about in the opening of the book is that my original idea that I was really interested was in the relationship between theories of taste and technical infrastructures.

You can imagine kind of a simple story, right? Which is that somebody’s got a theory about how taste works and they build a system to implement that theory, and then that's how that exists, right?

If I think that people like music because of how it sounds, I might develop a system that somehow analyzes the sound of music and then uses that to make recommendations. And then you can imagine this sort of second order story, which is like — sometimes people have tools and then they have theories. So, maybe if I have a system for analyzing musical sound, I might be encouraged to think that musical sounds are very important and then [create] a theory as a result. So, I was sort of curious about that dynamic between tools and theories.

And what I found in the field was that people didn't really have theories about taste. I expected them to have pretty strong theories. Instead, they had this real openness. They really wanted to be open to any of the ways that taste might work. And so most of the technologies that they were building were aimed at trying to be ready for whatever might happen, in terms of taste. People might like music because of how it sounds; but they might like music because of what their friends like; or they might like music because of where they are or what they're doing. And there are all sorts of influences that they wanted to try to be ready for, or to develop systems that were kind of generic and open enough to capture it. And that was surprising to me, actually.

Allison Jerzak: And I can imagine it's not just all these different contexts or ways people might have tastes behind music, but the idea that all music can be encompassed in a system as opposed to just little clusters of music.

Nick Seaver: Yeah — there's a kind of spaciousness that some of these systems aim for. It makes sense that they would try to be open-ended because by the time I'm doing field work on this stuff in the early-2010s, you're not developing systems for each group of users, really. Most of these systems are aimed at this everybody kind of audience. Whereas you might imagine if you had to build a recommender system for some very well-defined musical subculture or something — first off, they might not want it; they might not need it because they know what they like. But second, they might have some very specific thing that they're looking for and you might build for that. So these people are definitely trying to build for a kind generic audience.

Allison Jerzak: In your book, one of the things you draw attention to is the idea of listening. And listening comes up a couple of different times in the book. In one of the chapters, you talk about how listeners really break into two categories: lean-forward and lean-back. And a lot of times there can be tension between developers when they're a different type of listener than many of their users. I'm wondering if you could describe these two kinds of listeners, and the relationships or tensions the developers sometimes encountered.

Nick Seaver: Yes; I should say that the important thing here is this is how the people that I talk to in the field really imagined the variety of listeners. One thing I'm interested in in the book is trying to catalog a bunch of different things about how this group of people think.

And one chapter, the one you're referring to, asks how do they think about listeners? A very dominant model of listener variability in the field is usually glossed as lean-forward or lean-back. It's a little bit of a mishmash of categories. But, in general, the idea is that a lean-forward listener is someone who is actively looking for things — they're searching for stuff, they're clicking on things, they're interacting with the interface. They're willing to expend effort to find new music. The lean-back listener, by contrast, is listening to music in the background, and maybe does not want to put in so much effort to click and find things; maybe doesn't even care as much about music, and so on.

Now, this is obviously like a very sort of idealized dichotomy, and I don't want to pretend that people in the field like weren't aware of that. But it did pop up a lot; it was something that that that came out very, very often. And I think one thing that's interesting about it is that it really aimed to stand for — in the minds of a lot of the people I talked to — the major difference between the people who built these systems and the people who used them.

The idea is that there are a lot more lean-back listeners in the world than there are lean-forward listeners. And the kind of person who ends up working professionally and music recommendation is probably a lean-forward listener, right? There's someone who's really interested in music. And those people (according to this theory) want different things out of a recommender. If I'm a lean-forward listener, I want a recommender system that helps me find new and obscure music. If I get a bad recommendation in the in the stream, it's not that bad because I'm really I'm looking for stuff and I'm willing to deal with it. Whereas lean-back listener would have different goals, right? I don't want to be bothered. Don't give me weird stuff. Give me things that I'm don't have to get up and change the music and so on.

What's funny about it is that, you know, people recognized that that was sort of, you know, boxing people in. And these are systems that imagine themselves to be really unboxing people, to help people find music that they never would have been dealt — music that never would've been served to them before. And so they've tried to sort of relativize it and they'll say, well, can maybe there's not lean-forward and lean-back *listeners*, but really lean-forward and lean-back *situations*. You might be lean-forward at one point in the day, but lean-back at another, right? I might be happy to hop on the computer and actively search for music in the evening, but if I'm at the gym, I am not doing that. I am effectively lean back (ironically enough) because if I'm exercising, I don't want to interact with the thing. I want the music to just play. And that leads to a whole other set of issues, which is, you know, how do you tell what situation someone is in? How do you know whether someone's at the gym or not? And there are a lot of different ways that you can try to do that. I end up talking in the book about how the rise of context as a thing that people in music recommendation care about really comes about as a response to this concern that they are unfairly boxing in listeners, despite their ostensible goals of introducing people to music that they've never heard before.

Allison Jerzak: Did you come across the idea of filter bubbles in your work? I can think of an issue with filter bubbles where you were recommended a kind of music in one context, but you actually might be interested in many other clusters of music. And if you're filtered into one space, maybe you will never be recommended music in the other space. Did you come across that at all?

Nick Seaver: Yeah, I think filter bubbles are, in some ways, another name for the concern about the positive feedback loops that might happen from recommendation, right? If you get recommended one thing, and you listen to that thing, and the recommendations are based on what you listen to, then you're still listening to that thing, and you get recommended more that thing and what's going to happen.

I think the idea of the filter bubbles are really interesting. One, because it really emerges originally in the context of social media recommendations, which are kind of a different category of thing than music recommendations, weirdly enough, because there's user generated content, there's community, there's communities communicating with each other. Whereas most of the people who are receiving music recommendations are not also making the music that's going on to the onto the platform. So dynamics are a little bit different in the context of social media.

There have been a lot of critiques of the filter bubble idea — empirical critiques — that indicate that it maybe doesn't happen as often as people think it does. It seems like it logically makes sense that that's what might happen. But, it turns out that, you know, people who are interested in political news online get all sorts of political news online and they are not, practically speaking, filter bubbled. That doesn't mean that we shouldn't have concerns about the influence of these systems or the possibility of positive feedback loops. But, the sort of "they're stuck in place-ness" really could only happen if you are getting everything exclusively through this one system.

A lot of these systems are predicated on the idea that there is some kind of outside input. If you use Spotify, for instance, and you keep hearing the same stuff over and over again, I bet that the people who work at Spotify would hope that you, at some point, are going to hear about some other artist and search for that because otherwise it's hard for them to know how to find some other thing to recommend to you.

One thing I think that came up for me when I was writing the book was realizing (and I talk about this in the in the epilogue of the book) that a lot of the ideas behind these systems were formed at a moment when music recommendation was not the only way that you found new music. We live in a world now where if you open Spotify, for example — basically everything that you see is going to be algorithmically recommended to you in some way. That didn't use to be the case. It used to be that a recommender system was a weird little feature on top of, or in the corner, of some other way of listening to music. And so sure, if it's a little thing that you can dip into occasionally be like, "Let's see what the weird computer that's always wrong says about what I want to listen to," that could be that could be useful. But now, we're in a moment where lots of stuff gets run through the recommender. It's very easy to only listen to stuff that's been recommended to you.

And the assumption now, which is very different than it was even ten years ago, is that recommenders *work* because a filter bubble only happens if a recommender system works, right? In order for you to get filter bubbled, the recommender system has to be correctly identifying things that are the same as each other and showing them to you. If it doesn't work, that's actually kind of a plus. Little mistakes are (in some sense) a way to explore the broader world of music.

And I will say that, if a system knows enough to filter bubble, to give you all the same stuff, then by virtue of that fact, it knows enough to break you out, right? If it knows what similar, it knows what's dissimilar. So it could be tuned to give you things that are outside of your bubble. It's hard to say for any given service that's out there in the world, but plenty of them are built this way. So, you know, if you're listening to an algorithmic radio stream, there might be specific songs that are testing you. That are like — maybe you like this? And those might be test. Those might only happen if the model thinks that you're in a lean-forward mood, so to speak. Like, let's not do that while we're playing the romantic date night playlist. But absolutely, these systems try to build in a little bit of that. It's called an exploration-exploitation trade-off in the literature. They can exploit what they know about you and get you the same things or they can explore. And maybe it's a little bit riskier — you may not like it — but, they might find that other thing that it turns out that you like.

Allison Jerzak: That's fascinating. I haven't really thought about how in some instances it's beneficial to push someone out of their bubble. The other way you really talk about listening in your book has to do with I might call machine listening (although that's taking on a very specific term). But you talk about how developers essentially incorporate algorithms to listen in specific ways, and I'm wondering if you could walk us through how that process works.

Nick Seaver: Yeah. There's a there's a chapter in the book called Hearing and Counting, which is about listening in a few different ways. One is the kind of listening we would call machine listening, right? This is when computers analyzing audio data and doing things about it. I think of that, in some ways, as being a special case of a broader practice of listening that I witnessed all the time.

Which is, if you're a developer and you're building a music recommender system. You want to make a similarity system and say, okay, give me an artist or a song, show me other songs that are like it. What are you going to do as you're building that system? I'm sorry; this sounds speculative. It's not speculative; I've seen it happen. You put in something that you know and see what the system returns to you and literally listen to it. Like literally play it in your headphones and say this, this, this, this, this.

I did this as a participant in some of these spaces, to help work on evaluating these systems because there is not an objective, say, auditory similarity measure. For a while, I was doing this sort of scoring work during music informatics conferences where they have algorithms to measure musical or audio similarity, and they would have people go in, listen, and rate how similar these two clips are on a scale of one to whatever. And so I did that for a while.

It's so interesting because how does that work? Well, if I think that things are similar, I am drawing on my own in enculturated listening practices. One time when I did this, I ended up scoring two things. You know, they play a maybe ten second clip of audio and then give you 50 results out of these various algorithms that people have submitted to this contest. One of the examples I got was a classic rock song or something like that. So, I could go through the things that were said to be similar and it raises funny questions, right? There's a rock band…is it classic rock? I don't know. Here's an orchestral recording. This is clearly not similar in my mind. Okay. Well, the other one I got was a rap track. And I knew from being an enculturated person that there are regional differences within rap that I am not intimately familiar with, but that are very significant to other people. I was not equipped to say like, "Oh, this is really different. This is like a particular regional style," or something like that.

And so you can bring what you know into these systems. And because this happens all the time while people are developing these systems, there's this kind of pervasive enculturation of the overall system that happens in a very casual way, just by the people who happen to be there. Because you're going to be able to test it on the stuff that you know. And yes, they're going to test it in other ways, too. But you can't beat this constant, informal testing. And it's something that's kind of unique to music. You can't do this if you work in like a company that does credit scoring or something like that. You might develop like an intuition for credit scores, but you're never going to be able to be like, "Oh, I remember the credit scores. My parents used to play on the radio when I was a kid," and therefore know. There's this deep way that people can draw on their own experiences as listening subjects to feed into these systems and change the way that they work.

Allison Jerzak: This fascinating. Was there any kind of self-reflexivity on behalf of the developers for these systems? Were they aware of how much their own cultural positions could potentially bias these systems?

Nick Seaver: For sure. Yes. I think that this is another important thing to keep in mind. It's easy for people who come from my corner of the university to imagine that like everybody else who works on this stuff is some sort of idiot and that they don't know. They never thought about the fact that there's like a cultural bias based on their own experiences.

But, you know, they work on this stuff every day. They read things. I often found that when I was in the field, people were reading stuff that I thought of as being from my corner of academia. I remember people sending me stuff, like Kate Crawford, who is a prominent critic of Big Data practices from academia and beyond, and saying, "hey, did you read have you read this article?" It's like — oh, you read this too. Great. So they know.

What's interesting about this is that music recommendation is a domain where nobody can pretend that there's not something cultural going on. They know that what they're doing is cultural. We have a kind of a script in Science and Technology Studies where we find a domain that is technical and say, “hey, this is secretly cultural.” You can't do that in recommender systems because people know that's what it is.

However, it's not obvious what to do about it. And so when we say, “are they being reflexive?,” I can say —for sure, they're being reflexive. Does that mean they're making the right decisions? Absolutely not. There's nothing about being reflexive that guarantees you're going to do the right thing. And we have a tendency to equate reflexivity with (ironically enough) just being right or doing things correctly. One of the things I try to do in the book is to talk about that, to talk about how these people were reflexive, because like lots of people, they are thinking about what they're doing. It's just that there's more than one way to think about it. They're going to think about, you know, how are we different from our users? Well, you know, we're really into music, right? Like I was saying earlier, there are other ways that they're not like their users, but they're going to foreground to that one, in particular. Is that being reflexive or is that not being reflexive? Is it both being reflexive in one way and not being reflexive in a way that that some people might wish that they were?

Allison Jerzak: In the book, you talk about a process that you call acousmatic listening. I was wondering if you could just explain that concept a little bit.

Nick Seaver: Acousmatic listening is a term that I picked up from the music literature on a particular history of experimental music, music concrete — music that was made largely from recordings of “nonmusical” sounds. And the idea in a lot of the acousmatic listening literature is that when you listen to a recording of something presented to you as a recording, without all the cues about what it actually is, you might start to do a different kind of listening. Acousmatic listening refers to the kind of listening that you do when you don’t know what the source is, or where the sound is coming from.

There’s one school of thought that suggests that when you do that, you're going to listen for timbre or features of the sound, rather than being like, that's a train, which is what we do in our normal listening. However, there are plenty of people who say, well, what was actually happening when you listen acousmatically is you're just constantly guessing about what the thing is. And so acousmatic listening is this listening for something, and trying to imagine what's on the other side.

One person I've used to think through this is Nina Sun Eidsheim, who talks about this in her book *The Race of Sound*, which is largely interested in how people listen for race through recordings of singers. They say, here's a recording of a black opera singer, for instance, and people will listen through the recording and try to hear race. When you listen to music in a recommender system, there's a different order of acousmatic listening happening. I'm referring to that experience — which is a peculiarly contemporary experience — of hearing a song and thinking not, “what's the source of this song?,” or what instruments made it. But rather, “why is this song being played for me right now? What is the logic that delivered this song to me?” If you open up your music recommendation platform of choice now and hit play on something, you might have a theory about why that thing is being played for you. That's a very common feature of listening in these in these systems. One of the arguments that I usually associate with Eidsheim is this idea that all listening, to a certain extent, is acousmatic like that. We're always listening for, we're always adding meaning in to the signals that are coming to us. And so it's not surprising that we do this also with the outputs of recommender systems.

I have a story about this, which I always think is sort of funny. While I was doing research for this, I was visiting my family and having dinner in this steakhouse that was pretending to be a hip, modern thing. And the music was horrible. It was the worst playlist. It was a bizarre mix of corny crooner music, like Michael Bublé, mixed in with a very strange mix of things. My sister was there and she was like, I think this is a Frank Sinatra Pandora station, because a Frank Sinatra Pandora station would make sense if you were running a steakhouse and you want to have this kind of cool vibe. And so we went to the bar and found the place where the music was playing, and that is what it was. It was incredible. It was actually a Frank Sinatra Pandora station, and they never pressed the buttons on it. They never pressed like thumbs up or thumbs down. So there's this poor Pandora station running in the background. It's frantically guessing at what you meant when you said, Frank Sinatra. Did you mean crooner? Did you mean a rat pack? What are you looking for? And there was a way to sort of listen for that, and to imagine what's going on behind that. That is a very weird kind of listening that that we can do nowadays. I mean, you could imagine before, like why did this deejay pick this or that? But it's taken on a special a special form when we have all of this weirdly algorithmically mediated listening.

Allison Jerzak: Yeah. And I can imagine you might create or give the system more credit or more intelligence when you listen in this kind of way. I know from reading about Ringo, which is the first music recommendation service, that even those developers were saying that people ascribed Ringo an intelligence that it clearly didn't have. But there is some deeply humanistic instinct to try to ascribe an explanation to why something was recommended to you.

Nick Seaver: Yeah. So this was in the early days of collaborative filtering, right? This kind of technique that exists in the RINGO system and in a bunch of others in these early days are the classic, “people like you liked things like” this kind of recommender.

These systems know nothing about audio data. They only know about patterns in ratings or maybe patterns in listening behavior that have been interpreted as ratings. What's funny about them is that you get these results out of collaborative filters that seem like they know something about the audio data. Because people often will listen to songs that sound similar to each other. So if you go based on listening patterns, you'll get something that seems like it knows about the audio data. Or like I talk about in the book, a movie recommender that appears to know things about genres, that puts a bunch of movies that are in the same genre next to each other in this space that it creates. People will call this the magic of collaborative filtering. And it's really surprising when it happens, and it leads people to think that the systems are using data that they aren't using.

One way to describe it is as you just did — to say that it leads people to imagine that they know things that they don't really know. On the other side, people might say, well, what does it mean to know the thing, then? Maybe you can do it with this kind of data, maybe that's enough. But I was always very struck by how much the magic of collaborative filtering, so to speak, resided, in the end, in the enculturated minds of the listeners, to return back to that listening question. The only way that you could look at the graphical output of one of these clustering algorithms and say, “Oh my gosh, look, all those movies are similar to each other,” is because you know what similarity means in that context. If you didn't know anything about those movies, you couldn't do that. So the magic does rely on you; relies on the listener.

This is a very similar thing to ChatGPT nowadays. A lot of what looks remarkable in those systems only looks remarkable if you know enough about it to be able to fill in the blanks. Like it knows that if you don't know anything, there's a lot of problems, right? A lot of these things don't seem to work as well as well as you might imagine, if you don't know enough to cover over their mistakes.

Allison Jerzak: Getting back to the idea of clusters, I want to talk a little bit about musical space, which is another chapter in your book. You talk about how data is imagined in very spatialized ways, and musical space is no different. It seems like genre, at least in some cases, is one of the ways people try to cluster this music. And I'm wondering if you found anything besides genre people use to cluster music, or how are people thinking about musical similarity?

Nick Seaver: Yeah. The relationship between genres and clusters is an interesting one, and it's one that a lot of musicologists have started to take on as well, which I appreciate because I'm not one. I guess what I could say is that clusters sort of take over, maybe, the role of genres in these systems.

If you, like many people, at the end of the year get your Spotify Unwrapped. They'll say, "Here's your favorite genres" or whatever. And what are those? It will sometimes be a list of styles of music that you've never heard of before. Like, the names don't mean anything to you. The names might not mean anything to the artists. Because what they are names for clusters of listening behavior. And this is really striking because it is a very different thing than what genre has been historically.

In the book I draw heavily on a couple of musicologists who have written about this — Eric Drott, who has written about genre in the age of algorithms, and Tom Johnson, who has also written about something he calls genrecraft in this context. The idea is that the reason that music in the same genre sounds similar used to be because someone would make a bluegrass record on purpose. Whereas today, you might discover after making something that you have made a Float House album, and you did not know that Float House was a thing. But that is what Spotify calls it.

You know, you certainly are still influenced by other people but what it is changes when it's treated like a cluster. And into that mix, you get a lot of other things that are not obviously genres in a traditional sense. If you look at the of map of genre names used by Spotify, for instance, as I do in one of the chapters in the book, you'll see genres with names like Cello — which is not a genre. Or deep cello, which is deeply not a genre. Or comedy, which is arguably not a genre either. Or focus, which is more like a context than a genre. But what it does name is some kind of coherent listening pattern that's out there.

And it's a weird thing because as the people who work on this stuff will tell you, those clusters are real. Which is to say, they do reflect like actual listening behavior. That doesn't mean that they are genres in any absolute sense, because what is a genre anyway? But it does raise some interesting questions about what to do about them. What are these things? What does it mean to have our musical experience organized in this way that it didn't used to be organized?

I think genre theorists have been interested in this for a long time. Our ideas of genre are always historically specific. What is a particular genre means varies over time — not even what genres exist, but what the structure of genres is varies over time. And a lot of our ideas about musical genre are rooted in this idea that the mid-20th century recording industries are the natural state of music or something like that, right? I don't think we would actually want to argue that explicitly. I think a lot of arguments about these contemporary systems imagine that what they've done is screw up the way that things are organically, but that the way things are organically is the United States circa 1960, which I'm not sure we would want to do.

Allison Jerzak: Yeah, that is very true. And I know in music a lot of times genre is thought of as this highly contingent social phenomenon. There's musicologist Matthew Gelbart and he basically argues that genre is a social contract and that it's really predicated on groups of people and a certain kind of sociality.

I want to go a little bit further with the clustering idea and that there are these clusters that reflect real phenomenon —the systems are finding real things either in the audio data or in the listening data — but it is not immediately evident to the developers of these systems exactly what the computer is finding is similar. I'm wondering about that gap and if there is some amount of unknowability and guessing or intuition that goes into these systems.

Nick Seaver: Yeah. I mean, I think what you're pointing out is a basic question in any of these "data-driven" enterprises, right? Because data don't tell you anything on their own. There's always interpretation of what the data mean. And there's always this question about whether you should believe it; whether it's pointing you to something that's salient for what you're trying to do or whether there's a mistake somewhere. There are all these levels of rejection that could happen with the data. I always like to say there's really no such thing as data-driven decision-making. Because someone who wants to make a different decision can always find some reason to say, "Well, that data is actually not relevant to the specific issue that we have right now." And so, yeah, when people encounter these sort of outputs, there's always this interpretation that has to happen. What does this mean for me? Is this relevant? Is this the kind of thing that we're that we care about right now in this in this setting?

One spot where that sort of changes is when things are algorithmic to the extent that, you know, sure, Spotify can spit out these names of genres to you and say, well, there's this Float House or whatever, and that's one kind of thing, right? We can imagine all sorts of funky feedback loops where an artist discovers that they're in this kind of genre, and so they start to make music differently, or a listener discovers that they like this kind of genre, and they start to listen to a different set of artists than they might have done previously — performative feedback loop kinds of things.

But in the insides of recommender system, these categories can still be functional without ever being named. If you're listening to the daily mixes on Spotify. These are a bunch of different personalized playlists that sort of replicate the experience of listening to radio presets in your car, Ostensibly, they're supposed to get a bunch of different kinds of genres that you're interested in. They don't usually tell you (it's varied over time), but they don't tell you what those genres are. They'll tell you the names of artists and they rely on you as a listener to be like, "Oh yeah, artists like that person." And one of the reasons they'll do that is because, like we were talking about, the users will fill in the blanks. If you tell me "artists like whoever," I'm going to think, "oh yeah, I like that." Whereas if you give it a name, I might be like, "No, that's not the thing. That's not the thing I like." But under the hood, those same categories might be what's making things happen, without telling me necessarily. And so it's very possible that I have a Float House radio station on Spotify that is not called that, or one that is partly that and partly something else — I would imagine that is how they're probably doing it.

It's a funny experience where those labels — not the labels themselves, but those categories — can start to inform the recommendations without ever being named as such. I think a lot of people resist the names and say like, "Well, that name thing is wrong." Or have you ever done the thing where you find on Google, the advertising list, where it's like, "here's the list of things that we've profiled you and we think you're interested in?" It's always a very goofy list, and people are like, "this is ridiculous. These systems are so bad." But companies don't use that internally; they don't need those names. They can give it a name to show it to you, but they never have to name it for themselves. They can say like, Cluster Five is the thing that you're interested in. It doesn't matter what the name is, they will always be incidental to what's going on inside. They'll use it to show it to users. Sometimes they might use it to show advertisers. If you want to say to an advertiser, "oh, you want to advertise to people in cluster five," they'll probably try to give it a name. But yeah, it's weird that there's this level of functionality in these systems that doesn't require naming in that way.

Allison Jerzak:. And it seems like the opaqueness happens on several different levels and for several different reasons.

Nick Seaver: Yeah, you're right. This is exactly an example of algorithmic opacity. And there's a lot of them, it happens at many different levels. One of the things the acousmatic listening idea gets you to is that this happens under conditions of opacity. It happens under conditions where you don't know what's on the other side. But if we really we take that acousmatic argument seriously — the fully-fledged when we find with someone like Eidsheim, we see that everything is opaque at some level. And this is not an excuse for the algorithm companies. But it is to say that we should not imagine that the alternative to this situation is one in which we know why everything happens. That everything is clear to us. I don't want to hold an algorithmic system up to some standard that lets us know why everything occurs, because that's just not a normal state of affairs. But again, my goal is not, strictly speaking, normative in these contexts. I'm really more interested in describing how these systems work empirically.

Allison Jerzak: The last thing about your book is more of a comment than a question. It was striking that your last chapter takes on pastoral metaphors, talking about parks and this idea that people in these systems are (in some sense) caretakers of this thing that they can't quite control. In musicology, there's Guido Adler, who we consider the founder of the discipline. In his first essay that lays out the aims discipline, he also invokes a gardening metaphor. He thinks of it more as a garden, and the task of musicologists are to preserve and cultivate the garden. So his conception had more control. But it was really striking reading that in your book.

Nick Seaver: I didn't know about that! That's so interesting. This chapter in the book is about how a lot of the people I spoke with when I was doing fieldwork would use these funky, natural metaphors. Like, I'm a data gardener; machine learning is like farming; I'm a park ranger. And so I analyze some of these in the book and try to figure out what's going on there.

There is in the world of data studies, a well-filled critical niche where people talk about these naturalizing metaphors as the "naturalistic fallacy." They say that the problem with them is that they say that data is just objective. It's out there in the world. And these nature metaphors are a way to disavow responsibility, to say, "oh, you know, that's just the way things are. We're just like farming."

What I try to do in the in the book is to offer a different reading. I don't think that what people are doing with these metaphors is, strictly speaking, naturalizing, in part because a garden is not natural. A garden is the intersection of nature, culture, and technology. When people use these metaphors, they're often talking about the amount of control they have, which they experience as real. They know that they can change things. The guy who designs the genre system at Spotify — he can change how that works. And he knows. He's not like this is only objective. He knows that there are arbitrary choices he's made, but he also knows that he can't do anything he wants. He can't make the data say *anything*. He can't produce literally any goal. The experience of working in these systems is one of being constantly stymied and confused by like surprising stuff from outside of your control. I think those metaphors usefully talk about the kind of confusion and intersection of control, arbitrariness, chaos, instrumentation, all of these nature, culture and technology things that these people encounter it in their work. And it's so interesting to hear that it's also used in the context of music. It makes sense, though.

Allison Jerzak:,Yeah, it does. And I think there's a resurgence of people getting interested in issues of style, with gardening metaphors and botanical metaphors. So yeah, it'll be interesting to see where that goes in our field. So then, the last question. It seems like your newest work is on attention. You're really thinking about the different ways in which humans are thinking about attention and how we have certain values around it. But then, there is also attention in these computer systems. And similarly to what we've spoken about before, those are not always aligned in the way that we think that it is. Where do you think your work is going to go next?

Nick Seaver: That was that was a great description of where it is now. The idea for me is that I became interested, through working on recommender systems, in the value that people place on attention. I noticed that people put a lot of value on it, and they put different values on it in different domains. People who are critical of social media will talk about how Facebook erodes someone's ability to pay attention, to control their own attention. Someone who is working on a system like a music recommender will think about their goal or the way that they sort of measure success, as how much user attention they are able to command, which is related to that first problem.

But you'll also find in machine learning nowadays — in most of these generative models that are getting a lot of press — have at their core what's called an attention mechanism. Which is a technical bit of machine learning tomfoolery that has a very loose relationship to the thing that it's named after (just like a neural net also has a loose relationship), but is also treated in terms of attention. The paper that introduces this mechanism is called "Attention Is All You Need" because these are systems that are built largely out of these so-called attention mechanisms.

And so the project now is basically a straight-ahead STS project: here is this thing, attention. It's on the one hand, like a scientific object, a kind of psychology thing. It is also clearly a cultural object. It's a thing that people use to make sense of their lives and talk about what they value. Attention is a deeply used and valued thing in our cultural moment. And it's operationalized in computers, among other places, and it's measured in all these different ways.

One of the things I really liked doing in my first book was digging into really specific metrics and close reading them for how they implemented certain ideas. How they embody certain ideas about what people are like, about what they're trying to measure? And so I'm trying to collect as many different ways of measuring attention as I can find, which range from things like dwell time — how long are you on a website, which gets used for advertising revenue, for instance, or online survey design. I've been interviewing the people who design trick questions in online surveys to see whether you're actually paying attention to the survey or not. There's a whole body of theory and research about how to design those questions. So interviewing people who work on that and trying to pull together these really fragmentary moments of attention, to give an anthropological picture — as opposed to the psychological or economic pictures we usually get when we talk about attention these days — about what attention means to people, and about the many meanings that attention has.