Sound Expertise

Season 2, Episode 11 – Cracking Algorithmic Recommendation with Nick Seaver

TRANSCRIPT prepared by Andrew Dell'Antonio

SPEAKERS
Nick Seaver, Will Robin

Nick Seaver 00:00
If you're an anthropologist, and you're not interested in -- how do they do the weighting in their algorithm, but you're interested in -- how does this class of company understand what musical taste is, and the relationship between humans and machines and media technologies and expressive culture? How do they understand that, this general class of companies, that's actually something you can study without having to get into Mark Zuckerberg's locked file cabinet.

00:29
[intro music]

Will Robin 00:50
Welcome back to Sound Expertise. I'm your host, Will Robin. And this is a podcast where I talk to other music scholars about their research, and why it matters. You've probably heard on NPR or cable news, or some other podcast, a pundit proclaim that algorithms govern our lives. It's pretty much true, from Netflix telling you what to binge to the much more disturbing effects of predictive policing. Now, algorithmic music recommendation, when Spotify and Pandora suggest new music you might like based off of what music you've listened to in the past, might not sound like the worst part of our current tech driven music economy, compared to dismal royalties and the inability for most musicians to make any income off recording anymore. But these recommendation systems do deeply shape how we consume music today. And mainstream critics of algorithms rightfully argue that we have surrendered our ability to make decisions, a crucial aspect of our humanity, over to machines. But what's often overlooked in that conversation is that algorithms themselves are not exclusively machine generated. They're made, interpreted, and endlessly tweaked by people who work at tech companies and have their own ideas and values about what they want their platforms to do. And so to understand what algorithms really mean for the world of music today, we need to look at who those people are, and how they think about shaping our musical tastes. That is the expertise of today's guest, Nick Seaver, an assistant professor of anthropology at Tufts. Professor Seaver has spent many years embedding himself enough into graphic fieldwork and conversations with the creators of algorithmic music recommenders. And we'll hear from him now about what he learned, and importantly, how he learned it.

02:39
[intro music]

Will Robin 02:50

Let's talk a little bit about algorithmic recommendation systems, which is at the center of your research and which shapes how we consume media today, and how we consume music on services like Spotify and Pandora. What is an algorithmic recommendation system?

Nick Seaver 03:06
That is a good and surprisingly hard question. In general, this is funny because as someone who studies algorithmic recommender systems, you think I'd have a really good ready to hand definition. But as an anthropologist, I'm sort of of the camp that you don't need to know what something is in order to study it. And it's sometimes better to not know exactly what it is in advance, because then you can find out and you're not necessarily going to determine what everything that you find is, but in general, and not to be too obtus, an algorithm of recommender system is a system for connecting listeners to music, and vice versa, usually on the basis of some sort of model of those listeners or music. So Pandora radio in the US is is a common reference, Spotify, Discover Weekly, all of these things. It used to be the case that I had to tell people what a recommender system was, because you couldn't be guaranteed that people would know. But nowadays, I mostly take it for granted that people have used one of these things before, and they probably have strong opinions about them.

Will Robin 04:11
What is the recent history of, mergers and collaborative filtering, I think is the process, how does that… What's the history of that? How does it relate to how Pandora and Spotify work now?

Nick Seaver 04:23
Sure. So the the potted history that you'll get if you talk to someone who works in this field, so the kind of people that I study as an anthropologist, is that modern recommender systems sort of get going in the mid 1990s, with a technology called collaborative filtering, which is more or less a spreadsheet. Imagine users along one side of the spreadsheet and music or movies or whatever along the other, and then in the cells of your spreadsheet will be ratings. So if you liked a movie five stars, it's a five and so on. Now, most of the cells in that spreadsheet are gonna be empty, if it's all the users and all of the songs or all the movies because people haven't listened to or watched everything. And the goal of the recommender system in here is to guess what's going to show up in those empty spots. So the system says, Oh, we think that you are going to like this four stars, because based on your other ratings and other people's other ratings … there's a lot of ways to do this, but that's the basic gist of it. And so once I get those, I find the things that I think you're going to like, that you're going to rate highly. And I show them to you in some way. And that's what a recommendation is. That's the canonical collaborative filter. And what people thought was really cool about that, and sort of still sometimes think is cool, is that it is domain independent, which is, like I've been saying, it doesn't matter if those things are music, or movies, or recipes, or jokes, or hotels or academic research papers, or potential partners, they could be anything. And that kind of system will still be used and is still used today.

Will Robin 05:57
And so, as you mentioned, you used to have to explain this to people. And now everyone kind of knows what it is. How did you — you started your fieldwork at a time where these were not necessarily a dominant force in cultural life, how did you come to want to study algorithmic recommendation as an anthropologist?

Nick Seaver 06:15
Oh, well, so I started wanting to do it before I was an anthropologist, my little academic bio is that I was a literature major as an undergraduate who wrote about music technology, I wrote a weird little thesis about noise, music and recording media, as a total outsider, didn't know what I was doing, thing. But it was super fun. And I did a master's degree in media studies where I studied the history of the player piano. And I was sort of a fake historian for a brief period of time, I went to one archive. And that experience was so great that It turned me off of archival research and got me wanting to study something where you could talk to people who could answer your questions. And I was always interested in this relationship between automation and music and the way that music, this is an idea that I've learned a lot about in reading the work of, say, Georgina Born in particular, the way that music is this thing that destabilizes a lot of ideas that we have about categories, cultural categories, ontological categories, humans, machines, expression, math, feeling, all those things get mucked up in music as you and I'm sure that people listening know. And so I thought, okay, I want to study automation and music. What's the thing where that is happening right now, okay, recommender systems exist. I know about those, it's 2009. And I'm thinking about going to get a PhD, that's what I'll do. And so that's what I proposed in applying to graduate school, I applied to a bunch of different kinds of programs. One of them was anthropology, I'd never taken a real anthropology class before. And long story short, I went to anthropology grad school and became an anthropologist who studied recommender systems on the basis of that.

Will Robin 08:00
Interesting, okay, and so, when you made the decision, you want to study this. And then presumably, you had a couple years of coursework, where you figured out how to do anthropology, what was it like going into the field in terms of, what was the field to go into to do your research? And what were the assumptions or things that you wanted to figure out in going to tech companies and hanging out with these young dudes making algorithmic recommendations?

Nick Seaver 08:28
Yeah. So first off, you're right is mostly young dudes. And it was certainly mostly young dudes when I was doing field work, although they are making slow progress diversifying, and we can talk more about that later. But I'd say there's two halves of this: one is the theory half, what kinds of things might you want to know? Or what could you study? And the other is the practical half — how do you do this? How are you supposed to study these things? So maybe we can take them one at a time. Let's start with the practical stuff, because that's something that people are always interested in. You want to study an algorithmic system, say, pretend you want to study Facebook, or something, you want to study how they do the newsfeed, how are you going to do that, if you don't work at Facebook? They're not going to let you into Facebook, they're not going to let you go see the way that the sausage is made. That's impossible. And in the literature around algorithms that has blown up since I started doing this research, in the humanities and social sciences, this kind of access question, this framing of the black box that's the defining quality of algorithms …That's been the big thing. The problem with algorithms, in some cases is Oh, because they're too rational, they're too, whatever. That usually pops up in the music context. But in general, the problem is that they're secret, people don't like the fact that they can't have access, and for good reason in many of these situations. So just deciding that you want to study

them is not quite good enough to get you access, as we'll say, as ethnographic fields. You want to go study someone, you got to get there, you got to be able to go. So that was a really challenging problem for me that I dealt with in a very classically academic way, by theorizing my way out of it and trying to say that access didn't matter as much as you might think —which is, I think, partly true. So we have this idea that you really get knowledge by going somewhere and gaining access to the site. And we built up this sort of mythos of the black box, this impenetrable fortress of rationality, that if you really wanted to understand you've got to get into. And the problem with that is that once you do get access, so this is something I can say, because I did get to inside of a company to do field work. So this, whatever, don't listen to me. But once you get in there, you realize that there's not an algorithm there. There's not — oh, there it is, I open the door, and there's the algorithm. There's a bunch of people, and there's more doors, and there's more secrets, and there's more opacity. And you know what, that's true in everyday life, too — even in a field site that you might imagine was totally open to you, as an anthropologist, historically, pretend you go to some small-scale society and you're living in a village and you're like, I see everything. No, you don't, you don't see half the things that are going on. Historically, you really don't see half the things because you're probably a man. And you probably don't realize that women do their whole own thing in the place that you're studying. But even setting that aside, you just don't see everything. That's what it's like to be a person anywhere. And so one of the things I tried to do was to really dig into that idea and to say, Okay, what do I actually want to know? And where do I actually need to be in order to know that? We can really fetishize the fly on the wall and the meeting where the big decision is getting made. But that's not the only thing to know about algorithmic systems. And this idea that once you got in there, people ask me, how does … they'll ask me of companies where they think I did field work, how did this company's recommender system work, as though my job as an anthropologist is to go in there to get it, and to bring it back out. If that were true, then you could just ask the engineers who work there — they could tell you and my job, I'd like to think, as an anthropologist has to do with something a little bit more than to collect secret facts from engineers and just tell them to someone else. There's something else that my job is. So that brings us to the theory side. I'm sorry, I'm sort of rambling, but…

Will Robin 12:21
No, it's OK!

Nick Seaver 12:22
The theory question. So what do we want to know? And I don't think that what we want to know is, how does it work? I mean, in broad strokes, you can get how it works, from reading academic publications, sometimes we're looking at patents, although patents can be unreliable, those are a little fantasy documents. But in general, you know a lot of these things, or it's not too hard to know these things in broad strokes. And if you wanted to know the details, the problem is that the more detailed you get, the more transient the things that you would know are. So say you use Facebook, and you want to know how they weighted something, in the algorithm that decides to share something on your feed, that might be different tomorrow than it is today. Because there is constantly... machine learning is happening, the models are adjusting themselves, but also people who work at Facebook are changing things constantly. It might be different for you than it is for someone else, because you're in different test groups, or you've been personalized to; it would be different in different countries. There's all sorts of differences, especially once you get down to these details that you know you might be interested in.

And if you're interested in those things, you got trouble, you have a lot of stuff you need to do. But if you're not, if you're an anthropologist, and you're not interested in -- how do they do the weighting in their algorithm, but you're interested in -- how does this class of company understand what musical taste is, and the relationship between humans and machines and media technologies and expressive culture? How do they understand that, this general class of companies, that's actually something you can study without having to get into Mark Zuckerberg's locked file cabinet. You don't have to go and crack his computer in order to ...

Will Robin 12:41
It's not about proprietary secrets anymore. It's about what are the values of the people.

Nick Seaver 14:09
Right. And the argument that I'll make then is -- what is an algorithm? Well, people like to say an algorithm is this series of steps. It's a recipe, it just does the rote thing that it's going to do. That's not really true of the algorithms out there in the world. The algorithm isn't really an algorithm in a technical sense, but has all these people in it making decisions, and if you want to know how an algorithmic system is going to change, or what it's going to do later, then you can get really far by knowing how those people in the system, how they think about what they're doing, because how they think is going to shape the decisions that they make. So if I work at Facebook, and I think, Well, my job is to connect the world and to make us more open or whatever it was that they used to say, then you can look at their decisions and say, well, they're probably going to try to do things that they can claim connect the world or you can be more economic reductionist and say they're going to do things that make money. But then you actually have to have a good theory of how they make money. And it's also true that a lot of critics of these companies don't have good ideas of how they make money, either. Which is, that actually gets a little harder, because once you ... companies get very sensitive about that particular detail.

Will Robin 15:17
So when you're figuring ... once you decide or theorize your way to understand what you want to know, and how you're going to figure it out, you do end up... Willow is one of the "companies"... tell me a little bit about ... you looked at a couple of different actual companies making this software which you then anonymized and renamed for anthropological purposes, as an anthropologist would do when you're doing fieldwork in another country and have to rename all your interlocutors or whatever, but what was that experience actually like, of hanging out with the people in the settings of their workplaces, but also at the tech conferences you were going to?

Nick Seaver 15:58
Yeah, so great. So that's good question. So what does the field actually look like? For me, in particular, what it looked like, was a lot of anxious interviewing of people, I had to... I ended up doing a lot of conversations with people. One of the things that characterizes anthropological research, and we're always a little bit salty about how we are not like journalists, because, oh, I don't know, if you're just interviewing people, then why don't you just do it faster, then. But one of the things that we say is that we do have this sort of time duration. So I've been interviewing people and knowing people working in this field, since about 2011. And so there's a difference, something happens, when you start talking to

someone regularly over the course of many years, you learn things, you start to understand unspoken things, even if you're not working in a company for three years or whatever, you get to sort of know things about them. And they get comfortable with you, start to talk to you in a way that's not -- you don't just get the spiel that they give to every journalist, which once you start studying this stuff, you'll interview someone, you'll look them up online, and you'll find interviews they've done with journalists, you see every article says exactly the same things that you heard in your first interview, because that's what it looked like. In any case. So the field work was, I went to a lot of conferences at first. So academic and industry conferences, on recommendation on music informatics, I got some funding to be able to go to those fairly regularly, which was really useful, because eventually ... academics listening to this will know the weirdness of conference dynamics, it's much worse when you're trying to do field work with people, because then you're really anxious about whether people are talking to you or trying to hide or not. But eventually, you're one of this sort of core group of people who are showing up at these different conferences, and you start to get to know people. And this was actually how, after trying on purpose to get into companies, I just got invited into one where I ended up doing more fieldwork by someone who knew me because we had been to these conferences, and at some point, I probably -- this is a sign I probably should have asked earlier. But he was like, Do you want to come to our company and study us just where we are. And I said, I'm like a vampire, I can only come in if you invite me. But it was... so what I did was I ended up spending a few months as an intern at a music recommendation company that I call Whisper. And that involves sitting in on meetings, I had fairly free range inside of the company. So sitting in the company chat channels, talking to people, interviewing as many people as I could drag into offices to interview, and sort of seeing how they worked over the course of basically a three month period. Most of those people... sorry, not most of the people I'd interviewed, but most of the people that I already knew there I had been interviewing for a few years at that time. So I was a known quantity for at least a handful of folks in the office. So I went to conferences, I did this bit of field work there. And I lived in a few different cities in the US so that I could try to go to meetups and talk to people locally, and so on. And I should say that this kind of access issue is not unique, as I mentioned, to studying tech. It's arguably there in any sort of anthropology. But one of the things I really like about anthropology as a discipline is that because anthropologists study so many different things, I can draw on really weird counterparts to suggest method stuff to me. So some of my key methodological touchstones were a book by Lilith Mahmud, which is called the Brotherhood of Freemason Sisters, where she did an ethnography of Freemasons. So secret societies, literally, a book by Graham Jones called Trade of the Tricks which is about magicians, which is another set of people who have secrets and don't want to tell people what they are. And then a bunch of work by Hugh Gusterson, who is an anthropologist of primarily the nuclear ... the American nuclear complex. And so studying nuclear scientists, who also have a lot of secrets. And what's nice about that is you see, okay, my problems aren't my own. There are things that are in common here. So Gusterson can't get in... even harder than my case, he can't get into the nuclear research labs. But he can go to the restaurants where people hang out in town, he can talk to people, go to their birthday parties, he can be all over the place. around there. There's a kind of, he calls it... He calls it polymorphous engagement, which is maybe too many syllables. But there's this basic idea that you scavenge, you can get access to stuff, all sorts of stuff, if you don't think of yourself as pursuing access, just straight away getting in the door, and then knowledge will happen.

Will Robin 18:18

And so as you are immersing yourself in these different experiences in building this knowledge, how did your views of what an algorithm is and, specifically, what the role of music is too in all this, begin to change or adapt? What are the conclusions you started to be able to pull away from how these companies are shaping our musical lives? How they are ...

Nick Seaver 21:11
Well, so I want to pull apart two things there. One is ...

Will Robin 21:14
Yes, please deconstruct my question! I'm talking to an anthropologist,

Nick Seaver 21:20
This is the way it happens. So one big important thing, I think, is that the way these systems change our musical lives is not a question you can answer from the kind of fieldwork that I did. So there are people who are ... I'm interested in this question, obviously, I think it's an important one, as anyone who listens to music today, I think should be interested in it. And I think that my research is only one part of this. I can tell you what people who work in these companies think and I think fairly reliably, although I think as my fieldwork recedes into the past, things might be changing. So I'm working on the book, it'll come out eventually. But there are people interested specifically in this question of how a recommender system shapes what it is to listen to music, how people develop their tastes, and all of that. And I think that's really important. And I think a lot of people have this problem, where they get really technologically determinist about it. They say, well, this system works in this way, and therefore people must respond to it in that way. And that is not how technology actually works. So it's easy for me to say, Oh, well, so and so, this person thinks that musical taste is like this, that idea informs their system, that system informs how people listen, therefore, their theory becomes true. Voila, that's what [unintelligible] means.

Will Robin 22:38
So if this software engineer has some idea about how the algorithm will recommend the music, and that's how the algorithm maybe works, and then people have a musical experience with Spotify, then that's all this one-to-one thing, but that's not actually ... it's more complicated.

Nick Seaver 22:53
Yeah. I certainly wrote grants where I suggested that that was the case. And that was why I want to ... we want to study these set of people. And I think there is ... obviously, there's some aspect of it that has to do with that. But we might want to think about the mechanisms. Because the other thing that comes up, actually, and this is -- well, we can talk more about this, because it's interesting. Artists, people who make music, whoever they are, fall out of this a lot. So when I was doing this sort of research, there's a lot of talk about listeners, and a lot of talk about how people build these systems. And I would find myself realizing that I had not talked about musicians at all for a long time because music in these systems equals recorded music. So one person memorably put it to me, he said -- music exists for me when it's digitized, which is not quite a claim of, there's no such thing as music outside of here, but more of a claim of -- well, the music that I can do anything about is recorded music, and not only recorded music, but digitally... music that exists in digital form, and then more narrowly the specific

music in digital form that exists in my library. And there's a whole world of music and musicking that's out there that is not at all contained in that definition. So I try to be wary of replicating the thing that many of the people in the field that I was studying were doing, which was equating listening on their platforms to music in general. What's music today? Spotify tries to do this a lot, they put out those year and things to be like, what did we listen to? And it's like, excuse me, not everybody listened to music on Spotify, but you're -- I see what you're doing. I see this little trick. Now I don't remember the question I was supposed to be answering. [both laugh]

Will Robin 24:37
So what were your takeaways from fieldwork -- obviously that's a big question, your whole book is about that! But what are the questions that you began to hone in on at least the areas where you realized Okay, this is interesting. talking to these people is going to start to ... the things that maybe changed from when you went into your field work that you realized would be different?

Nick Seaver 25:03
Yeah, that's it. Okay. Great, yes, there's absolutely a big thing here. So one idea was that I had a theory going in, that there is this relationship between ideas about musical taste, and techniques for doing recommendations. So we've talked about collaborative filtering, this idea that there's patterns in preferences as expressed through ratings. We didn't talk about how those ratings -- so five stars, or thumbs up, or whatever, are usually replaced in collaborative filtering by implicit ratings. So if you listen all the way through a song, that's implicitly a positive rating, not as good data as a thumbs up or whatever, but enough times, we start to guess that you like this, and so on. But there are other ways too. So my theory going in was that techniques and theories of taste had this kind of one to one mapping, where if you thought that people liked the music that they liked because of how it sounded, then you might try to build a system that analyzed musical sound, and then see what songs are alike, right. So Pandora famously has this music Genome Project, which is -- they've got people, they call them musicologists, but I believe that they're people who have undergraduate music degrees, I don't want to offend real musicologists. And they listen to a song, they fill out basically a several-hundred question survey about each song that scores it... if people are familiar with Alan Lomax's Cantometrics project, it's kind of a weirdo offshoot thing of that, although they claim to not have known what Cantometrics was, but it's a thing that you answer questions like, is there a saxophone in this song, zero to five amount of saxophone... I don't know if that's literally one, they don't actually release them. But that's the idea. So you would have this data for every song. And then you could map those things to each other. And they say, Okay, well, you liked this song. And then as you listen to Pandora, you may have clicked on the link that explains why they played something for you. And it's always like, we played this for you, because you like music that has.... And it's always super generic. It's rhythms, or [unintelligible], but it's always groove -- a groove, four on the floor beat, female vocals, whatever, that kind of thing. So that's a theory, and it's a theory that people like music because of how it sounds, so the good way to recommend things to you is to give you stuff on the basis of the sound. And that the best way to know how something sounds isn't to do some sort of fancy computer analysis of the audio, but to ask experts to code it. Now, the trick with that, though, is that most systems today -- so this is what I found, so I go in thinking that's what it's going to be; you think you like how music sounds, or sorry, you like music because of how it sounds, you do something like this, you like music because that's what your friends like, let's try to get your social network data. You like music because of -- you're

at the gym, or you're at the office, let's try to figure out where you are. There's theories. But it turns out that all these systems now -- and this became really clear to me when I started asking people why they thought people ... sorry, this became clear to me when I started asking my interviewees, what is your theory of taste? Why do people like the music they like? And they would all look at me, like, what? I don't know, that is an impossible question. And I'm thinking, don't you spend...

Will Robin 28:31
Isn't that what you're supposed to be doing?

Nick Seaver 28:33
Isn't that all day? Isn't that what you're doing? I don't understand. I thought this was the whole point of this. And it turns out that the way that most of these systems work now, Pandora included, is that they are ensembles, which is to say that they coordinate a lot of different kinds of data, and a lot of different kinds of signals into their model. So they don't... no one uses only the musical sound, including Pandora, they might use some data about the musical sound, there's lots of ways to get that data, they might use ... so we can take Pandora as a good example. Because I think people will be familiar. So it's personalized radio, that you can get things thumbs up or thumbs down. There's this music Genome Project thing. So if you say, make me a Britney Spears radio station, that sort of canonical Pandora move will say, Okay, well here are other artists that sound like the sort of prototypical Britney Spears song. And so we're going to give you those. And that'll be great for you. Because it's not only going to give you popular stuff, it's going to give you anything that sounds like Britney, regardless of how popular it is. And that sort of makes sense to a certain extent. But then you start to wonder, Okay, well, to get more songs into Pandora's ecosystem, we need to have people score them. And that's a bottleneck, that's not easy to do. And so you start to get computer audition, machine listening systems in play. And so I saw someone from Pandora give a talk at these conferences where they're like, well, we have all of this data of experts who have coded songs with qualities. So we can learn, a machine learning model — which of these genomes, which of these genes as they call them, map to acoustic features that we can analyze computationally? Which ones can we reliably predict from that data? And so now we can just do that.

Will Robin 28:55
So they replaced the “musicologist” with AI “musicologist,” or something.

Nick Seaver 30:31
Basically. My understanding, and this is probably a few years old, but my understanding was that it was a sort of stopgap thing. You can put things in using this AI system, and you would still use the human experts to make a good version of this…

Will Robin 30:46
… check it…

Nick Seaver 30:47
… but at least for the stuff that was easily … there are certain things that are not that hard to figure out, like what key a song is in, what the tempo is, there are details that are tricky about it, but they're not as

confusing as something like, is this a happy song? Or is this a song about sports or whatever, there's lots of songs about that. So in any case, so they have that, but they also have all this thumbs up and thumbs down data. And as people I talked to who worked at other companies, and they had no doubt in their mind, even when this wasn't public. And I don't know if it even is a public thing. And I don't actually know if it's true, I should say, but it's a rumor, that Pandora, holding all of this data about skipping songs or thumbs up, it would be ridiculous for them to not do collaborative filtering with that, to not use that a little bit to be, well, other people "thumbed up" this, it's not only the audio. And so my understanding is that now they use all these things together. And that's true of all these companies.

Will Robin 31:44
So it's not just, the next song you listen to is one that sounds like the previous song. And it's not just the next song you listen to is one that your friend liked. It's all of these different things working together.

Nick Seaver 31:55
And it's always more. It's like anything. So when I'm doing fieldwork, this is primarily in 2014 is the bulk of my in depth fieldwork time, which now feels like a billion years ago. But when I'm there, you're seeing people constantly trying new sources of data out. They're saying -- Well, what if we used this? What if we used that? This could make sense. And so what you had wasn't really a theory of taste that was really specific. That was — people like music because of how it sounds, which is what I thought you might have, but you had instead a kind of open attitude toward what might constitute or affect taste. So people when I asked them, why do people like the music? They look at me, like, Are you stupid? Because they were thinking, well, I don't know, it might be different for everyone a little bit, it might be different at different moments in time or in different situations. And so the goal of my system, and my goal is not to figure out the real reason why people like what they like, but to build a system that was open, to build a system that was able to accommodate whatever the hell taste might be. And to support people in finding things that fit their taste, however that taste might come about. And that's a very standard machine learning kind of take, people always like to cite the Wired article about how big data means the end of theory, that we don't need to have a theory about why people do anything, we just look at the data. And that's it. That's kind of a lousy take, because of course, there are theories, there are implicit theories, as Marilyn Strathern, a sort of arch anthropological theorist, says, one of the things anthropologists do is study how our generic terms are actually specific. So that set me up for the real anthropology, which is -- okay, you want to just let taste be whatever it can be, you're just gonna open it up. Okay, what's the structure of that? What does openness look like? What does — taste could be anything, we don't know, what's the shape of that? And it's not nothing so narrow as — it must be the sound, or it must be the friends. It's shaped in all sorts of weird and subtle ways that are harder to specify.

Will Robin 34:07
And so this is maybe a good way to talk a little bit about — one of the central discussions around algorithms is now algorithmic bias, which at its most basic level is, white guys make facial recognition software that only recognizes other white guys. That's not the case in how algorithms work with music, and one of the claims that these companies make is that these recommendation systems allow them to be post identity or post demographic. But you found other ways in which you see a kind of algorithmic

bias emerging, or at least a way in which the people who work at the company are encoding something into these algorithms that has to do with who they are. Can you talk a little bit about that?

Nick Seaver 34:53
Yeah. So that's a great question. So one of the huge questions on algorithms in general, are these questions about sort of disparate impact biases. And it's usually in domains that people understand to be much more significant than music recommendation: predictive policing, drone targeting, approvals for mortgages, all sorts of very, very important things that affect people's life chances. Now, I don't need to make the argument here for like why music is also important. But clearly, a decision that gets made around a music recommender is not going to be as momentous, that one decision, in someone's life as one of these other ones. So one of the things that was useful for was that I could see dynamics of decision making around algorithmic systems that might be a little bit more obviously problematic in other domains, people would be a little bit more open about them in in this one. So people would come to me and say things like, so you know how people like different music because of their race? And I could say, well, there is a sociological fact here about racial differences in music listening and music industry, and the racial structure of the genre system, okay. And then they said, there's not a way we could use that? And I'm thinking, they're afraid of racial profiling as a named thing. They're concerned about bias — the people I was working with didn't have a super elaborate language for talking about this. They wanted… they understood themselves to be the good guys, and they wanted to make good decisions. And so this doesn't mean that they did. And it doesn't mean that what they did was always good. But what was useful about this was that you got to see, okay, what do people do when they're trying to do the right thing, when they're trying to make systems that they think are going to not unfairly bias in one way or the other, are going to let people explore, yada, yada, yada. So, all around collaborative filtering, which grows up alongside the web, you have that classic internet -- on the internet, nobody knows you're a dog post identity discourse, a collaborative filter only knows your ratings. It doesn't know anything about your demographic qualities. And all the papers and ads and puff pieces in Wired Magazine, are all like, demographics are out. We just know what you actually like now, and isn't that great, you're free. And so you have a lot of important and already early on work. People like Alondra Nelson writing about how online race is always seen as either a kind of … if you can see race online, that's imagined to be a failure of the system. Or race is either a baggage that you should be freed from, or it's evidence that something has gone wrong. And that is absolutely the theory about race that you see in a lot of these conversations, including around bias, actually. The people who are involved in algorithmic bias conversations are not immune to this, this idea that a system should be race blind, that race should disappear. That comes up surprisingly often. But obviously, in a domain like musical taste, and like musical production, race is all over the place, in the way that music gets made, and listened to. And so what happens is what people in this field call a proxy problem, which is that, okay, we only are looking at what people thumbs up and thumbs down or whatever. In that data, you're going to have patterns that are racialized, because the people are coming to you from a society in which race exists. And they're listening to music that's made in a world where race exists. And so race is in there. And so you have this problem, where if your system pretends to be post-demographic, you are going to find in your "natural, true, real" data, you're going to find race again. And worse, you're going to find a version of race that you think is authentic and true, and the way things really are. So inadvertently, you've thought that you've freed yourself from demographic profiling. But you can potentially re-naturalize some of these demographic categories. Race is the most obvious one. But of

course, there's gender, and all sorts of other categories in here as well. So what's the point of this? The point of this is that people working in these systems were very insistent that demographics were not the way to go. So when someone comes up to me and says, Is there a way we could use race to make recommendations work, they are wary because they don't want to do a thing that they understand as racial profiling. But they also understand that using race goes against that initial idea of what a recommender system is good for, that a recommender system is about sort of liberating people from boxes. It's about letting you find whatever you want to find. And so this idea that I mentioned earlier of like, oh, we're going to use any data we can get, who knows why people like what they like, we're omnivorous about this stuff. That starts to come into tension with this original idea of Well, we're gonna let taste do its own thing, unencumbered by something like demographic marketing categories. Because all of a sudden demographic data becomes something you might get. And something you might use. Something that seems to these people plausible. And not unreasonably so. There is, of course, a way that if you want to know what music someone likes, we say this now, if I tell you ... this is becoming less true, it feels like, but if you say what your favorite song is, that's a very classic way to subculturally identify someone. Like, what bands do you listen to? And so we might say, oh, that shouldn't be built into the very infrastructure of music distribution in the way that it is, but of course, I hate to break it to you, that's how the music industry has worked for a long time, you know that. So it's not that it's new with these systems, it is taking on a slightly different shape.

Will Robin 40:56
And so, with all of these public debate style issues around algorithmic bias, or just around ... are these systems, these algorithmic systems, or even these music platforms the death of art or whatever, how do you conceptualize your role as someone who can contribute to these conversations? Is it just, we need to complicate, complicate, complicate, which I'm sure it is, but how do you ... do you have useful takeaways? Or is the takeaway, you have to just think about this for a really long time?

Nick Seaver 41:38
Okay, that's a fair question. Actually, there's sort of two things. One is that as an anthropologist, I am, of course, biased toward this, let me try to understand how this group of people make sense of the world, and try to find the reason in it. Now, that's not the same thing as saying it's good or that everything that people do is the right thing to do. I do think that it's important to do because we have this idea that we all know, we understand how these people think, we understand how these systems work. And I don't think that's true. I think that people don't necessarily appreciate the weirdness that goes on in these places. And I do think that some of the weirdness is an opening for critics. I think that some of the ways that these systems work that aren't exactly what you might expect, those might be spots where we could push, I should say that the other thing that's true here is that the effects of ... when I say these systems, we're thinking of different things. Most people are thinking of Spotify, in general. And I'm thinking of recommender systems, which are part of Spotify and other companies, but aren't the entirety of them. And very, very conspicuously do not include things like the licensing agreements and payment structure of the systems. So you had music recommended, this is something actually worth pointing out, because it was hard for me to remember. When I started doing this work in 2010, on-demand, large catalog, music streaming services, did not really exist in the United States, they were not obviously going to work. People had tried, and they had failed. And so music recommendation research for most of its history to date has happened in a world where that meant recommending you things in

your own iTunes mp3 library. The first music recommender system, by most accounts, is one called Ringgo, developed at the MIT Media Lab, and I think the paper's 94, early 90s. They're recommending CDs that you can buy, which is not remotely the same thing as recommending what song in the Spotify catalogue should you listen to now. And so I don't think you can talk about these systems independently of the circulatory infrastructures that they're part of. But there are very good reasons to have concerns about the centralization of music circulation in these big platforms, and all sorts of things that get sort of metonymically reduced to the algorithm, I think, as the algorithm stands for tech companies in general. But which are actually much bigger than a recommender system, per se. So that's the first answer, which is that the recommender system is, for better or worse, often isolated from some of the questions that we might really want to know things about, in terms of what's going to happen to music and to the music industry. There are specific things that do matter, so people are interested in things we haven't talked about yet but fairness for artists. So recommending songs, are recommendations going to be systematically biased against female artists, or against musicians of color. And those are empirical questions to a certain extent that can be hard to answer for people on the outsides of these systems, but are important and get framed very differently by people who are inside these systems and people who are outside, but there's a lot of work on algorithmic fairness now that I've been tracking. Not -- I don't know if specifically on music, but they always say it's generalizable. But book recommendation, what kind of metrics can we come up with for understanding whether our recommender system is unfairly biased towards one group of authors over another? And how would you know? It's not an easy question, actually.

Will Robin 45:29
Great. Well, thank you so much. This was really helpful and I learned a lot. I appreciate you talking to me.

Nick Seaver 45:33
Good. I'm glad.

45:39
[Outro music]

Will Robin 45:41
I'm very grateful to Nick Seaver, who is assistant professor of anthropology at Tufts, for that fascinating conversation. You can check out links to his scholarship over at our website sound expertise.org. And you can follow me on twitter @seatedovation. If you like the sound of Sound Expertise, check out the music of our producer D Edward Davis on Soundcloud at warmsilence. Many thanks to Andrew Dell'Antonio for transcribing our episodes to make them more accessible. Those transcripts are up on our website. I'm very excited for our episode next week, a discussion of intimacy tourism and sexuality and electronic dance music with ethnomusicologist Luis Manuel Garcia.

46:26
[Outro music]