

AI and You

Transcript

Guest: Chanuki Seresinhe

Episode 115

First Aired: Monday, August 29, 2022

Hello, and welcome to episode 115! My guest today is Chanuki Seresinhe, head of Data Science at Zoopla. She began working as a commercial data scientist at Channel 4 and, more recently, was the Director of Data Science for Culture Trip. Chanuki has a PhD in Data Science from the University of Warwick and it's her research there and at the Alan Turing Institute that caught my attention. She used big online datasets and deep learning to understand how the *aesthetics of the environment* influences human wellbeing. For example, how might we design our future cities to be conducive to our wellbeing? She's been featured in *The Economist*, *Wired*, *The Times*, *BBC*, *The Guardian*, *The Telegraph*, *Scientific American*, *Newsweek* and *MIT Technology Review*.

She kept her research going long after she got her PhD because she was so enthusiastic about it, and you'll hear that come out in the interview. This is a fascinating added dimension to the application of AI, and if you're wondering how AI can help judge how beautiful a place is, well sit back and listen to the interview with Chanuki Seresinhe.

Welcome to AI and You.

Hi, it's really great to be here.

So, you're involved in this fascinating field, it seems, I've never encountered anyone doing this before. So, I think you're not only leading the field, but defining it, of AI in the aesthetics of the environmental, perhaps you'd like to describe for us what that is, and then I'll ask you to talk about how you got into that.

Yeah, absolutely. So, I just wanted to understand the connection between our environment and well-being, and specifically focus on beautiful environments, because it's something that we haven't really been able to look at on a large scale and I think it's something very intuitive that we would think is connected. So, when you're when you visit a more beautiful place, we just automatically feel so much better. So, it all came about because I wanted to quantify that connection and that's when I started to also then quantify, what is a beautiful place composed of? Because that's something as well, we haven't done been able to, in the past, do on an enormous scale.

And the word that stuck out there for me is *quantify*, because it's not a word that people who are usually into aesthetics employ a lot, what is your background that led up to this?

So, I've been working in design for a very long time, even though now I'm working as a data scientist. But my past career in design, informed me that there are some aspects of aesthetics or beauty that actually can be quantified. So, when you want to design something beautiful, there

are some principles you follow. So, it's not completely out there that we can follow certain types of principles or certain types of patterns in order to create something beautiful. So, and of course, beauty is subjective to a certain degree. But there is a component of it that you can also quantify because there are commonalities between humans basically. So, that's kind of what I was getting at.

Okay, so what I want to do is pull this out of, or away from, a trope that people might be thinking right now, if anyone has seen *Dead Poets Society*, and is thinking about Robin Williams, the teacher at the beginning, going over the textbook of poetry with the kids in his class, which is written by, I forget the name, but it's a fictional author, and who says, "here's how you analyze poetry, you calculate something like the ratio of syllables to lines, and that tells you it's a good poem," and he tells them to rip the pages out, of course, it's a straw man argument. I don't think anyone ever actually did write that. But that's sort of where my mind goes when we're thinking about quantifying art and for the benefit of artists listening who are thinking, no - I'm sort of making a cross, fingers crossed sign there - this is anathema, describe how AI has benefited this field. What sort of results have you seen from applying AI to this that have advanced the appreciation of the aesthetics of the environment?

That's a very good question. So, I mean, I think there are some things that are possibly difficult to quantify. So, the beauty of art, the beauty of poetry, I could see where that's quite tricky. Now, the way you kind of define beauty of an environment is you can break it down into certain components. So, what I did was I taught an algorithm how to predict beauty - and we can discuss a little bit more about those details on how I did that - and then I also then tried to analyze statistically what did the algorithm find beautiful and then it's broken down to components, right? So, it's obviously it's going to find lots of natural features has been beautiful there's like trees, valleys, coastal areas, that you know - it's all intuitive like things that humans find beautiful there are going to find beautiful but I thought it was really interesting the algorithm also found characterful buildings beautiful like lighthouses, bridge-like structures, as being beautiful as well. So, again, it says this is broken down into certain components and then saying these are beautiful components if you combine this in a scene, this is makes it beautiful seeing on the other hand, it also recognizes things are not so beautiful and obviously if it's like a very industrial scene, if it's got a bunch of factories or so it's basically just looking at the scene what the components is made out of, and then making aesthetic judgement as to how beautiful that scene is?

So, is this supervised learning? Have you fed it a bunch of input data that has been classified and ranked by people as beautiful? Is that how it establishes its standard?

Absolutely. So, I use a website called ScenicOrNot and it's collects data from people. I mean, it's just a game people play, they read this in, which is between one and 10 and in the background, you have an image per square kilometers of Great Britain. So, by the time I came across the data, there was about 1.5 million votes with 200,000 images and I only chose the images that were rated at least three times so there's some kind of consensus that's there. Then I took those images, put it into a deep learning algorithm and it learned to tell what is beautiful or

not based on human input. So, if that, so and then I was kind of surprised, because I, initially when I did this, I was skeptical, I didn't think it was going to work. Because of course, you think beauty is subjective; and yet it did a good job predicting what humans thought of as being beautiful and then you look at the results. They're pretty intuitive. I mean, we wouldn't argue against it. So, that I said, I think it's about the components that are creating a scene and it's recognizing to understand what's beautiful.

Can you feed art into this now can you give it a Turner and other - I forget the landscape artists, - but those paintings, those semi realistic paintings, and let it loose on those and see what it comes up with?

It would be really interesting. I mean, normally, our algorithms are very fine tuned to whatever they were trained on. So, this algorithm is trained on images, they're from a website called geograph. So, they're documentary images. Now paintings are a slightly different image. But saying that I accidentally had images in my data set when I was basically predicting how beautiful London was and in this data set, I accidentally had indoor images and some of the most beautiful images were actually sculptures and I thought, this is really interesting and fascinating, that there was something not only about the composition, but about the form and the shape that it was also picking up on.

Will it identify features in an image as being particularly contributory to the ranking?

So, that's the one tricky thing about deep learning is, I don't know 100% why it's choosing an image as being beautiful. The only thing I can do is I can kind of post analyze it, take whatever it did find beautiful, and then extract the components out of it, to understand what was in the image and then kind of infer because this scene has mainly trees, then the algorithm found it beautiful, but that's that one only drawback when you use deep learning is that we don't fully understand why the algorithm picked it as being beautiful.

Would it be possible to essentially run this backwards and have it generally generated the idealized, most beautiful scene that it can think of?

Well, ideally, that would be amazing and I yeah, I really wish it was possible. I know, there was a team that did attempt to do this and there's a team behind Good City Life and they created an algorithm called facelift, but nothing to do with faces, but they were, it was about taking an image of a street and trying to see if you can generate the more beautiful version of that street. And what you can do is you can generate a template, so it's not necessarily something you would like, understand what that looked, as a human, but then you can match that template against another image of a street that's similar and you can see, ah, okay, maybe if add a tree to the street it'll become more beautiful. So, it's kind of, you can do it in a bit of a roundabout way. But at the moment, the tech is not there to fully understand what's actually in the setting to take that leap to generate something that that a human would go, Oh, wow. I mean, the machine I think is absolutely beautiful. But it doesn't necessarily make sense.

If I have my machine learning algorithms right, a generative adversarial network, a GAN, could be paired up with what you're doing to do that, perhaps?

Yeah, so I think I believe they did use a GAN. So, you use a GAN to generate and I would call it a template because the image itself is not a human understandable image and then you take that image and you kind of do a similarity analysis with actual images, and then you can find what the closest like, human readable image is to so that is an interesting technique that they use to kind of get around the fact that, unfortunately GANs, they don't fully quite understand what we perceive. So, it's hard for them to fully create a scene. I think a scene is quite complex. I think we'll get there someday, but we're not quite there yet.

The dataset that you were talking about; what are the images in it? You said, is it Britain, London? What does it cover?

So, yeah, it's Great Britain tells me, it's all UK-based. So, it's just a sample of an image per square kilometer. And we did another study on Rio, because we just want to just curious as how it might look like in a completely different setting and, again, quite similar themes. But obviously different things. Compared to Rio, though you have your promenades, which you don't have necessarily in England. It'll pick kind of aspects that are more common in that area as being beautiful. So, it can be trained for different countries, we just don't have data on hand.

That was a great opportunity that you had that dataset sitting there that was big enough to do this, which was already labelled and are there other datasets that are equally big enough and unlabeled for the aesthetics of other countries?

Not that's readily available. I know MIT were doing some work and they were creating datasets they just - is this is this image more beautiful than another image. Looking at T Street; there's a few research. Again, the team behind Good City Life also did the same. So, they've created datasets, but they're not necessarily open datasets, researchers can access them if you want to collaborate. But we don't really have large scale open data on aesthetics that can be used.

Of course, there are datasets on the subject of aesthetics of other things besides landscapes; I'm thinking people - swipe left swipe right - kind of data, not sure how available it is, because it's obviously proprietary to some companies. But have you thought about going there? Or would you rather stay with landscapes?

I would rather stay with landscapes and like I said, the original reason I went into this, is I really wanted to understand, what is it about certain environments that make people feel happier, and I'm much more interested in that question. I'm less interested in what makes a person beautiful? Because I don't know, that is not something I really don't want to go down that rabbit hole.

Fraught with risk for misinterpretation. Now, do your images contain a lot of buildings enough to be able to say this kind of architectural pattern is more beautiful than this one?

It's a good question and I don't have enough data. So, because it's England, and it is one image per square kilometer, when you sample that data, you're going to get a lot of historical buildings and, that's why is that the characterful buildings come out. But where you have more interesting architecture is in the city, and its high-density areas, and I just didn't have enough images in

these high-density areas to be able to do that. So, at this moment, I really don't know like, which styles of modern architecture people find more beautiful or not and it's something I've always wanted to do and I might still do it these days, because I'm quite curious if there are any, do we have any collective patterns of friends for modern architecture?

Well, and as it's also Britain, there are also miles upon miles of ghastly office buildings that would be rated zero or negative by most people; and did those make it into the set or were people only looking at the images that they actually liked?

No, they're looking at any they're looking at any image that was there in that square kilometer and yes, there were, ugly grey tower blocks, and that those did get rated down. So they're picking up on certain types of architecture. I couldn't go on style, but yes, those ugly office buildings did get rid of down people don't like them.

I'm reminded of this wonderful book that actually came to my attention because it was recommended in computer software development, but it has nothing to do with that. It's about architecture by Christopher Alexander called *Towards a Pattern Language* and it develops his whole taxonomy of Architecture starting from how to make a window all the way up to how to design a city. It's very appealing for geeks. But when you look at its conclusions about how to, for instance, create some small aggregations of buildings, call it a village if you will, or courtyard, then you have now a way of explaining what's wrong with so many places and what's appealing about things, for instance, in rural France. And I wonder whether you are familiar with that, or whether you have any ideas about developing this to the point where you could draw, say, some of the same conclusions, or say, yes, this maps on to these parts of that work?

Yes, I'm absolutely familiar with the book I've got the book is so fascinating. It's just amazing how somebody went to that depth to find patterns, not just for how to build your home, but how to build a whole town how to build a whole city out. And it's just fascinating, the kind of the amount of work and thought that went into it. Now, it is also a little bit too prescriptive in some ways. But maybe it's almost, perhaps it's about how you interpret the instruction so that you could have that. But some of the things, they, for example, they'll say, don't build buildings higher than four stories, supposedly high rises, make people crazy and there's abundant evidence that supports that, I don't know if that's actually true. But this was in the days before big data. Now, we do have, and actually, we can probably, we can actually start quantifying that. So, there are probably a lot of things in our book that actually can now be tested, and with large scale evidence, and we can actually see, is that true, is that not true? There's some of the things they say intuitively make sense. Some of the things seem outrageous, but what would the data actually reveal about what's actually good for humans? It's a question I'm absolutely fascinated by, like how should we build our future cities to be places that people really love to live in?

That fascinates me as well, I? And on the height of the buildings, I don't think he would say don't build any cathedrals. I think that would probably still work. Anyway, I'll leave that question aside for now. So, is anyone interested in in this, I can imagine tourist boards or, say,

municipalities in the Cotswolds wanting to slap a *Rated 9.9 by AI*, as most beautiful place in England, that sort of label. Has anyone been sniffing around like that?

Not necessarily, tourist places, but councils have been interested in understanding. I mean, I think people want to understand, when you change an area, what kind of impact you're going to have. So, if you can start to look at if you want to change something, is that going to make a place more beautiful, more appealing to people or not? So, in that sense, I've had government organizations actually interested in finding out what you can do with these approaches.

And your input data, you said it was one image per square kilometer, is that right? Of Britain? And it was covering the whole country, and does that mean that you could now produce a heat map of Britain with one kilometers resolution showing the more or less beautiful areas of the whole country?

Absolutely and I could do it beyond one kilometer. Because I then because I used the images, which I had at a one-kilometer resolution to train an algorithm, as I was speaking before, to then predict the beauty of new images. So, I can basically take any image, for example, from a website like geograph, which collects multiple documentary images of all around the country, and predicted at a minute scale where it's beautiful, where it's not beautiful. So I actually did it for London and that's in now, one of my papers is a map about in London and you can really see where the beautiful spots are and where the beautiful spots are not. And of course it'll pick the parks as being beautiful. That's not a surprise. But what a neat interesting surprise was, there was a little path and I was like, it's not on the road. Like what is that path? This is really interesting and turns out was a canal. So, of course the river and canal trust really love that. Because this algorithm is smart enough to understand that we love, this canal, it's really beautiful to walk down. And the algorithm picked that out as being beautiful. So, it works really well. It's so versatile. I'm still amazed about how AI can work so well.

May I ask, which organization underwrote or sponsored or grant provided the grant for doing this?

So, this is all my PhD research. So, it was at the University of Warwick, as well as the Alan Turing Institute. So, yep, so they were the two organizations behind this research.

Okay, have you got your PhD now?

Oh, yeah, several years ago, I think I got PHD.

You're still working on developing this further? Where do you want to take it?

Yeah, that's the thing is I love to doing this research. So, I work commercially as a data scientist now, I'm head of data science at Zoopla. But I love doing this research so much that I just couldn't stop working on it. So, I am still dedicated time every week where I purely dedicated myself to keep doing something with the beautiful places and I'm just really just trying to understand, where I can use it to do good things for humanity. So one of the projects I'm working on at the moment is with an organization called GoJanuntly and it's a walking app, and they suggest walking routes where people can get more exposure to green spaces and really

appreciate nature and connect with nature and I think it's such a wonderful service they do. So, we're working on a pilot, where we can also suggest beautiful routes. Because in my research, I also quantify the connection between beauty and wellbeing. People that visit more beautiful places to report more happiness. So, we're also going to do a pilot project to create beautiful routes and see, what kind of feedback we can get from people, what kind of products we can create to make people happier.

What about urban planning, or municipal planning? If, say, municipality is working on revamping, it's inner core, wants to redesign the layout; might they be interested in running this on it to see which of the proposals are rated higher?

Yeah, I mean, that would definitely be a future application. I don't think the algorithm's ready for that level of detail because planning permission's mainly on in urban areas and as I said, I don't have enough crowdsourced data in urban areas to be 100% confident that is going to be able to pick up the nuances. So, I would like to first do a larger scale collection of crowdsourced data in an urban area before I would create an algorithm that could be used for planning purposes. I think right now it's okay for walking on the street and seeing different settings. But I don't think - it's just not nuanced enough to comment on planning applications. But that is definitely a place we can go and be really useful in the future.

Are there any ways that you could crowdsource more data like this? Any way of getting people to do those ratings short of paying them on Mechanical Turk?

Nobody was paid to give the readings, and it was just a game, it will just be a matter of - actually wouldn't take that much effort - if there's any volunteers out there that want to help me out. All you have to do is basically create a website like scenic or not, but just feed it with urban images, and then send it out there. I'm just surprised how, like people just love contributing to projects like this. I mean, we've seen it with like, Galaxy Z, for example and I think for the love of it, people will contribute ratings. I guess the tricky thing with planning is when we're doing anonymous ratings, we don't know who's contributing the ratings. So, there will have to be some thought done to make sure we're capturing the opinions of a wide demographic of people and that we're not accidentally introducing bias because only certain types of people want to rate the images for free. So, there has to be some thought of making sure that we are capturing kind of a wide range of people's preferences.

And there will be the opportunity to do all kinds of data correlating with the demographics there like, young people prefer this, old people prefer that; we all have some stereotypes that we might want to see if they are proven. Maybe you could get into that, does that interest you?

Absolutely, I'd be really interested. I'm also really interested in cultural differences, because I am curious as to what we're picking up on, it's just kind of the European ideal of aesthetics. You know, what is it? What is aesthetics means for people coming from different countries where they grew up with, landscapes that look completely different, what is what is feeding and, remember, there was a time when I was in Australia, and I went there with this European ideal of "wow, this is nature, it's so beautiful, and I just really want to go on a hike," and then

everywhere I went, it was like dangerous. And it made me realize that there's a European idea is completely not true in other countries and we have to kind of take that into account that people have different experiences in the environment and how does that impact their aesthetic preference?

That's Australia, it's like Pandora, half the wildlife is trying to kill you. And people outside Australia find kangaroos cute and cuddly and inside Australia, they're pests. So, what sort of surprises that you haven't mentioned? Were there any other surprises that came to you, as you were doing this that you were able to put into your thesis?

Yeah, absolutely. I think the biggest surprise that was that people read flat areas of green space as being not beautiful and I thought that was absolutely fascinating. So, if you have an image, and it's completely green, and yet it's flat, there's no trees, no hills, no contours, it's going to be rated down. And that says a lot about green space policy, because green space policy at least in the UK, it's about let's add some green space. But nobody talks about the quality of green space. And that seems to have an impact on people's preferences, it didn't matter if it was just green. That was surprising, because I didn't expect that I just thought it looks lush and green and even if it was flat, people were really high, but they did not rate it as high.

So, I wonder what we should do with the flat space, either make hills there or put something that's not green on the space or plant a lot of trees, any thoughts?

Yeah, I mean, plant trees, flowers, it's just I think there just needs to be more detail. So, we're talking about really flat spaces, like almost looking like an athletic field. But you see that in some neighborhoods like that is their green space is flat piece of grass that looks like an athletic field and that's supposed to be someplace that people are supposed to go and spend their time in, and it's not appealing. There needs to be trees there. If you can't create the contours. If that's a bit more difficult, then you just yeah, this odd flower is just it needs something a little bit more interesting than just the green space and I think we can't underestimate that the quality is also really important.

Does this have any crossover with your work for Zoopla?

I guess a little bit. I think it's an interesting crossover, that really is definitely a correlation with house prices and beautiful places. People do pay more for homes and beautiful areas compared to non-beautiful areas. So, I mean, I don't work with the beauty specifically there. But I know that beauty is a component of predicting house prices. So it's an interesting kind of area. Because my ideal place would be to get where, there isn't such a big difference in people areas, because what we're talking about is that, there's privileged people, including me, I'm working in tech, I'm privileged in that sense, I have access to beautiful green space around my corner and that's not that's not the case for everybody else, because they're going to be priced out of a certain area and, their green space is not going to be that beautiful. And I would like, governments at least have a kind of a basic quality that they're giving to humans, to give them access to certain respects. I mean, we knew the impact of that, during COVID, when all you can do is actually go for a walk and people that lived in privileged areas probably didn't have that horrendous impact

that other people did when they can only walk out of the street into a concrete jungle, or a park that's completely unappealing and that must have been terrible for people. So I just really wish that governments at least think that this might be just a basic human need to create green space that's beautiful for everybody.

There's something foundational about this that just seems to appeal to something in the subconscious, something fundamental to us that I imagine people wanting to dig into this a lot more, like maybe, is the data available for people to create bucket lists? So, they can say I visited the top 10 most beautiful square kilometer's rated by AI in Britain?

It's not there yet. But I think I can actually make that available in the future. So, it is something I've always also wanted to do is create a map of all the beautiful places around the country, and then just give it away for free so people can go visit these places. I think that's kind of what we're trying to do with GoJauntly by creating these beautiful walks, because I'm giving them the AI data. And then people will know, like, hey, there's a beautiful spot just around your corner. It's like, there's so many beautiful spots that are just around the corner that we don't even know about. So that's what the algorithm can actually review.

What about viewing angle? Have you got data on the same place viewed from different angles?

So, what I then also did was I created an algorithm that can also work on Google Street View images. Because, when we're using the images from Geograph, they're basically uploaded by people that want to document a certain area. So, there's just, you know, one photograph of that space. Which is different. If you use Google Streetview images, it's kind of it's more, you have to capture every single angle of that space. So, then you have much more comprehensive data and there's definitely nothing like not missing anything subjective there. Because it has to collect the data in every single point from every single angle for the Street View to work properly. And so yeah I did actually train an algorithm and I had to feed it. Google 3D images that were rated by people, and I was able to then again, still create the algorithm that can rate how beautiful places by taking the beauty of every single angle of that place into consideration and yeah, I did it for London, and you see, again the same patterns, as you'd expect.

Just so I'm clear on your background, you're data scientist now; you started out in which field?

In digital, so I started in digital, since 1998, I was always interested, I'm always been interested in tech and pushing the boundaries of tech and at that moment, it was the web. That was absolutely fascinating. So, that's where I worked in. And I ended up moving a bit more into design. So I did some design, I was teaching designers and developers to create websites. But our main clients tended to be designers and architects. So, this is where the conversation started, about aesthetics and design. So, when I later studied economics, and then did my PhD in data science, it was kind of there always in the back of my mind and I think that's what really inspired me to kind of let's quantify this beauty. We see the environment and let's understand, how does it impact our well-being?

I think it's been fascinating to just see how someone explores this new space using AI to stretch into an area that we wouldn't normally associate with AI, and you've got such an evident

passion for it that is really infectious. I can't wait to see what you do with it next. Do you have any ideas of where you would like this to be 10 years from now, if your wildest dreams were realized?

If my wildest dreams are realized, I mean, I'd love to have beauty ratings for all over the world, mainly focus on the major cities, because I'd like to monitor how beauty changes over time, especially in the cities, and how that actually impacts human wellbeing in the city. Because I really, truly want to understand how we design cities to be better for people to live in and it'll be such an interesting measurement to have that we've never really thought of measuring before. So, that would be my ideal in 10 years' time.

How should people find out more about what you're doing, get the app, and follow you?

So, you can find out more about what I'm doing on beautifulplaces.ai and you can download the GoJauntly app. Just yeah, just Google GoJauntly. It's on the Apple store as well. That's where you can find the app as well. So, and then it's the beauty walks are not out yet, but they will be in the next few months. So, keep an eye out for that.

Well, thank you so much, Chanuki Seresinhe for coming on the show. You've provided a wonderful, beautiful new dimension to AI for our listeners.

Thank you. It was a wonderful to chat with you about this as is definitely one of my favorite subjects.

That's the end of the interview. You can find out more about this research at beautifulplaces.ai. I wanted to ask Chanuki where the ugliest place in Britain was, because I have my suspicions, but I decided that would be... unnecessarily divisive. Didn't that discussion just make you think of all kinds of potential uses of AI in evaluating aesthetics? With it now presenting the possibility for objective, consistent ranking? It's just not something that had been within my horizons until I learned of her work.

In today's news ripped from the headlines about AI, a paper out of MIT with the frankly irresistible title of "Nanosecond protonic programmable resistors for analog deep learning" was published by researchers who created hardware equivalents of analog components of the human brain, but because they were semiconductors without the limitations of biological cells that have to live in a saline solution, they could be made much smaller and faster than the biological equivalents – in this case, a million times faster, but with heat generated that was comparable to that of a human synapse. This kind of analog device has potential within deep learning. If you listened to episode 95 with George Dyson, you heard him predict that today's researchers would make significant advances if they pursued analog strategies, and here we have an example of that appearing to pay off.

Next week, I'll be talking with Chris Summerfield of DeepMind and also Professor of Cognitive Neuroscience at Wadham College Oxford, who lives at the intersection of neuroscience and AI, studying the lessons that each has for the other. That's next week on *AI and You*.

Until then, remember: no matter how much computers learn how to do, it's how we come together as *humans* that matters.

<http://aiandyou.net>