Entrevista a Geraint Wiggins[1]

Sección: Entrevista
Recibido: 12/12/2024
Aceptado: 30/12/2024

Geraint Wiggins Interview

Santiago Negrete Yankelevich
Universidad Autónoma Metropolitana-Cuajimalpa, México
ORCID: 0000-0002-6558-6538
correo: snegrete@cua.uam.mx
DOI: 10.46530/virtualis.v15i28.468

25 de febrero, 2025

Santiago Negrete (SN). This number of Virtualis is about Computational Creativity and the Arts, creativity in Design too, and I wanted to interview you because of your musical as well as theoretical experience in Computational Creativity. This is what we want to talk about. So to start off, I would like to ask you what you think of the new technologies regarding Large Language Model (LLM) technology and its use in creative work. Is it useful? Do you think it's interesting? What is your general opinion?

Geraint Wiggins (GW). It's very difficult to answer that question in just a few words. I think it depends a lot on what one is doing and why one is doing it. So, I mean, maybe it's worthwhile starting off talking about what a LLM does. And essentially what it does is it learns, broadly speaking, an approximation of the statistics of language so that, in the event that it finds a question that is sort of inferable from its learned information, it's capable of guessing what a human might respond to that question. And I think it's really important to put an emphasis there on *guessing*, and what a human might answer because, in general terms, in terms of these LLMs, it is absolutely clear that very many people do not understand that there is no notion of truth in them. So they make up

[1] Geraint Wiggins is a professor specializing in computational creativity and artificial intelligence. He is affiliated with the Vrije Universiteit Brussel and Queen Mary University of London. His work focuses on AI's ability to simulate intelligent and creative human activities, and he has played a key role in developing the field of computational creativity. He also leads the EUTOPIA AI Learning Community, which fosters collaboration among universities in AI research and education. This initiative includes specialized courses, seminars, and even AI programming tournaments. https://ai.vub.ac.be/team/geraint-wiggins/. [Geraint Wiggins es profesor especializado en creatividad computacional e inteligencia artificial. Está afiliado a la Universidad Libre de Bruselas y a la Universidad Queen Mary de Londres. Su trabajo se centra en la capacidad de la IA para simular actividades humanas inteligentes y creativas, y ha desempeñado un papel clave en el desarrollo del campo de la creatividad computacional. También dirige la EUTOPIA AI Learning Community, que fomenta la colaboración entre universidades en la investigación y la enseñanza de la IA. Esta iniciativa incluye cursos especializados, seminarios e incluso torneos de programación de IA. https://ai.vub.ac.be/team/geraint-wiggins/].

sentences, but those sentences are not necessarily true. And when the LLM does make up such a sentence, there is no claim that it is true.

We see it on the screen, we see it written down, and we know from psychological studies that people are more likely to believe something written down than something that's said to them. So, you know, we are gullible animals. We choose to believe things when perhaps we should be more careful. For example, if you hear something which rhymes, you're more likely to believe it than if it does not rhyme. So, these are ways in which people are gullible.

Now then, when we're talking about using LLMs for creativity, we could be doing that in different ways. I mean, if we're specifically meaning LLMs, then it needs to be some kind of linguistic creativity, presumably. In which case we are generating new sentences from old in some sense. We might be generating them according to some predefined structure like a sonnet or something like that. And it's extremely impressive, I think, from the engineering perspective, that models such as ChatGPT can do this to order. When they're instructed to write something which has the structure of a sonnet, they can more or less do so. The amount of reasoning that has to go on to do that is extraordinary. I think we need to be very careful when we see that going on because, effectively, all that's going on is an extraordinarily complicated statistical pattern match and projection inference. And humans have a tendency to anthropomorphize machines. All the time we talk about our cars as though they're people. We talk about our computers as though they're people. We generally say the computer *thinks,* something like that. No, it doesn't. The computer doesn't think. It calculates, that's all it does. And when we have systems that are capable of generating human-like language, we are really inclined to believe that they are living entities. And that's dangerous because they're not. They're just not. We hear people from Google even saying: “oh yeah, we've built a machine which is conscious”. No, you haven't. So far there is only one way that we know of making conscious human-like entities and that's the good old-fashioned way that we all enjoy doing so much.

I think these myths and legends that are growing up around these admittedly extremely clever systems are potentially extraordinarily dangerous for us as humans. When we believe things that are not true, we can be misled. So, overall, I think that's an important thing to say at the beginning and I find it deeply disturbing that the companies that are producing these proprietary entities are showing absolutely no concern at all about these dangers. I think it's irresponsible and I think there needs to be regulation of some kind. And that's particularly true because in a fact that not very many people know and even fewer people care about –distressingly– is that the companies that make these

things don't actually know how they work. They might know what the computing architecture is, but once the system has learned what it has learned and starts to generate new stuff, it is not possible, in general, to ascertain why it has done so. Even by asking it. Particularly by asking it, perhaps. And this means that these systems are complete black boxes, which may well be doing things that we simply do not understand. And which may be extremely bad for us, possibly. We can't tell. It's unlikely, but it's possible.

Now, there's a tension here because if the system is a black box system and it is not open to scrutiny, we can't tell what it does, but it seems to be doing cool things. So, normally speaking, when we're engineering, when we build a bridge, for example, it's the engineer's job to show that that bridge is safe. You can't be an engineer unless you have this notion of producing safe artifacts. And in computer science that pressure is very weak; and in AI, in this kind of AI, it is non-existent. The people who are making these things genuinely do not care that they do not know how these things work. In fact, it's cool –they think– that they don't know how they work. It's like a child with their bike doing wheelies and saying, look ma, no hands. This is dangerous for all of us. It's irresponsible.

Then moving on to creativity. Generating new sentences is, of course, a creative act. It's what I'm doing now. I know what I want to say, but the exact way I'm saying it to you here is something I'm improvising as I go along. That's a kind of creativity, it's a kind of everyday small kind of creativity, people do it all the time. In fact, some of them never stopped doing it. And so, you can say the models are in some kind of sense creative. And there's increasing evidence that the kind of creativity that humans do when they're talking in an everyday way is doing something rather like what the models are doing at the level of statistical generation.

We tend to speak in terms of expressions that we've heard. We tend to repeat expressions that we've heard and so on and so forth. So, you know, that's the kind of statistical learning that these models do in a way. But there's a difference here, which is that we understand what the words mean. And when somebody says to me the word *blue*, I have an experience in my head which is connected to my experience of seeing the color blue. I have that meaning. If I cut my hand and someone says: does that hurt? I know that they can share the feeling of pain, even though they're not feeling it right then. They know what it feels like to have pain. A computer does not know that. And this is the real difference that, it seems to me, very many people do not fully understand or, if they do understand, they don't care and they probably should because one of the things that allows humans to care about each other is the ability to share feelings. So

something, an entity, an active autonomous entity, which has no ability to share feelings is kind of like a psychopath. And generally speaking, psychopaths are viewed as dangerous things.

So then, coming back to the creativity –I just got a little bit into my old rant there–, one of the features of this kind of learning is that the results it produces tend to be within distribution. That is to say: they tend to produce structures, whether they be sentences or musical pieces or whatever it is, which are rather like the things they were trained with. And human creators, the ones that are considered to be great creators anyway, tend not to do that. So, I'm going to quote an example that was given in a talk I saw by Brian Ferneyhough, who is a modernist composer. He pointed out that the Sydney Opera House was a very great piece of art. Which it certainly is. And his argument was that it was a very great piece of art, not because it started a new trend, but because it ended it, because it made it completely impossible for anyone to do anything similar without it being pastiche. It is a unique entity, and therefore a very great creation because it is so distinctively what it is that you can't copy it. So there's no style as such, there just is that thing. That's an out of distribution creativity. So, you can attribute the same kind of creativity that we have in these models, they're not literally statistical models most of them, but they are, effectively, statistical models because it's the statistics of the data that is being modelled by the neural networks, fundamentally.

These things are like rolling dice to find out what next word you're going to use or play. You know, that idea was invented in the 1700s, attributed to Mozart, not so clear that that is actually correct, but it was invented around that time –the dice game indeed– in order to allow people to create new pieces of music from little snippets, which were usually a couple of measures or something like this, that you could assemble in any order. So it took some considerable skill of the composer to write them in such a way that they could be reassembled in any order. And of course that skill was a skill that was in that case done by the human. Now it can be done by a much more complicated statistical model than a pair of dice. But it's still generating within a particular framework which is inferred from the data it's been given.

So let me give you another example of true human creativity which is very close to my heart. This is in music. In the early 1960s, Dr. Bob Moog invented the synthesizer, or in fact popularized the synthesizer. Similar ideas had been around for a while, but Moog invented an idea called voltage control, which meant that the synthesizer could become a really dynamic and exciting instrument in a way that it hadn't been before. This idea of voltage control meant that you could use simple voltages to control different aspects of a sound like its pitch or its tone or its vibrato, things like this. And so Moog thought,

well, wouldn't it be useful if I built into my synthesizers a thing that I'm going to call a *sequencer*? That would create voltages to order in a sequence, a repeating sequence, so that then the person who's playing the synthesizer can use that to control stuff. So that's what he did.

These synthesizers were crazily expensive and not very many people owned them in the early days. But one person who did was a member of a German rock band called Tangerine Dream. And he bought one of the early Moogs, which is the size of a large wardrobe, and used it in Tangerine Dream's work. But he didn't actually play it. The person who did play it was a guy called Christopher Franke, who saw this sequencer thing and thought, aha, this is a way that I can make new kinds of sounds by using it to drive bass –well, actually not only bass, but quite often bass–, rhythms, and then by changing the sequence live as it was being played, turn it into a performance instrument. So we started off with a technological invention, which was then subverted by a musician. It wasn't meant for this live usage, but that's what he thought. So there's a kind of creativity. This guy saw something he had never seen before and thought, I can use it like this, which was completely novel. In Tangerine Dream's album from, I think, 1974, *Phaedra*, this technique is used incredibly well. And then we have a sequence of events here.

A bit later, along came a songwriter called Giorgio Moroder, who is still one of the world's greatest songwriters. And in the early 1970s, he heard Tangerine Dream's music and realized that he could do the same thing in a pop context. And he wrote the song *I Feel Love* for Donna Summer, which uses exactly the same idea. And when you listen to that song in context of the earlier Tangerine Dream, you realise that basically what he's done is take the essential idea of the sequencer and repurpose it for dance music. Because Tangerine Dream's early music certainly was not dance music. And that led ultimately to what we have now in EDM as well –electronic dance music– but that's a longer trajectory.

Okay, what's the point here? My point here is that no LLM could ever reach this level of creativity. Because each of these different things, each of these different, um, processes is saying that I'm going to take an idea and use it in a different context. And that's a thing that, certainly at the moment, AI is not very good at doing. It's only good at sampling from its own distribution, so its creativity is restricted.

The question about whether these things can be creative is a nuanced one, and the answer is: “yeah, they can a little bit”. But not very. And when they are, they do not know what they are doing. They are not setting out to make a new symphony or even

to surprise people. They're merely generating stuff because that's what they've been designed to do. And I find it very hard to attribute creativity to something which is not making a decision to create something. I think teleology is rather important.

So that's a good introduction for you.

SN. Yes. Thank you.

You were saying that LLM are used for text, but there are many other applications for video and for image that are being used for Design and other applications.

GW. Right, but then they're not just LLMs, right? They're linked to other models of different modalities.

SN. Yes, but they have this generative structure where you write a prompt and they generate what you ask for. Also I find interesting that, in creativity, you introduced the notion of truth. If you're using these systems for law, certainly truth is important, but if you're using it for the Arts, then maybe something that is not true or that is not accurate can bring about ideas like those you're saying that people in Tangerine Dream and Giorgio Moroder had, something unexpected, something that is not true, and they find it interesting. They say: “oh, this is really crazy, but we can use it”. It can bring something new to the process of creation. So I find that interesting in the use of these technologies for the creative purposes.

GW. Can I just interrupt you there and disagree with you very slightly? Because truth does matter in creative contexts. Science is creative. Engineering is creative. And those two things both rely on truth. So in artistic cases, maybe it doesn't matter so much. Maybe the notion of truth is different in artistic cases. And in the case of music, you can't even really say what truth is, right?

But I think that's really important to remind ourselves of that because it's certainly in my experience of working in this area for thirty odd years now, as soon as you start talking about creativity, people do assume that you mean the arts and, actually, creativity is in pretty much anything interesting that humans do.

SN. You mentioned some examples of prototypical creativity, like the Sydney Opera House, a brilliant case, and there are many other activities that have to do with arts that maybe the computer wouldn't be able to do because they fall into this category. But there are other situations where they can come into play and prove valuable as

creative partners. We are now talking about co-creativity. That is, how machines can be used by people to create and whether they could be creative in this context. For context, imagine a Hollywood film or any film. In order to produce it, you have people doing different things and they're all creative in their own disciplines. Although they say that the director is the star, the rest of the people working there are also creative and very talented. And if one of these roles is played by a computer, maybe the computer can be deemed creative in this context.

GW. Well, of course, that's already happening, in computer animation, a lot of in-betweening is done automatically, for example; and this used to be a highly skilled animator's job. So yeah, absolutely. People who are worrying about their creative jobs being removed by computers have something indeed to be worried about. Because computers generally, in the long term, are cheaper than humans. You don't have to keep on paying them. You just have to buy them once.

So, there are some consequences there, which we can certainly be talking about, but to come back to the point of your question about co-creativity, I completely agree with what I take to be your premise that this is actually a situation which is more likely to arise and which is more likely to give rise to exciting creativity, let's say. Because for the moment, the computers are indeed restricted to the content of what they learn. It's very difficult for them to look beyond that and to see something completely unexpected and to understand why it would be good. Whereas humans can do that. Humans have the capacity to spot serendipity. In science and in arts and in lots of other contexts and say, “yeah, I'm going to grab that before I lose the opportunity”. And the co-creativity is a place where this can really happen. So even if a computer is capable of generating stuff which is really unexpected and out of distribution, it still needs some way of noticing that it's good. And or that it's worth keeping or whatever you want to say. Good is a difficult word to define. And so having a human there who can do that is perhaps a great way to work. I think in those circumstances it's rather difficult to draw the line between where the creativity of the human and the creativity of the computer is. So I would prefer to call such a system co-creative and not to attribute the creativity to either one or the other. What is certainly the case is that we can use computers or any other kind of generative artefact such as tossing a coin or rolling a dice to stimulate the creativity of humans. By just helping humans think outside the box a bit. So co-creativity is interesting and there's a lot of scientific work on co-creativity looking at the interaction between humans and machines and the emergent results of such interactions.

SN. How do you think this works or would work in music? How does a composer use co-creativity?

GW. It could be all sorts of different ways. To borrow one example, I have a former student of mine, Lizzie Wilson, who now has her doctorate, is a very accomplished, internationally accomplished, live coder who performs music by writing code that plays the music as it's being performed. And she did work looking at how co-creativity can be monitored, be understood, and how a system can generate suggestions for a live coder to use and then either reject or accept as they're working live. And she looks at the effects of these on the audience and whether they increased the audience's perception of creativity. And if I remember correctly, it's a while ago now, they did indeed increase the audience's perception of that.

So yeah, I mean, there are some interesting things that can be done. When you want to have really tightly linked co-creativity, where there's a real bond between the machine and the human, there are all sorts of other questions beyond the creativity that become deeply difficult. Like how do the two things communicate? So as a composer, on my own composition work, I tend to use a sequencer, that's to say, a piece of software which enables me to write notes down but not in standard musical format. And then I listen back to them and manipulate them and so on. Now, personally, I don't use AI for that. I use my ears and my brain and I create what I want to hear using that method.

But you can well imagine that perhaps the system could make suggestions for me or whatever. And then how would I communicate with that? How would I work with that? Would I, I don't know, if I've got my score in front of me, would I like draw a circle and say, suggest something to go in here? And then maybe listen to it back? Or would it be better if the system was noticing what I was doing and then perhaps making suggestions alongside or marking places where it had a suggestion? Personally, I would find all of these things annoying. I don't want the computer interfering with my work. So, my point is, it's not clear what the right way to do this kind of interaction is, and there's a whole bunch of research that needs to be done. I suspect it's probably deeply personal for most composers how they would want things. I suspect there's probably quite a lot of research to do there to find all the different possibilities.

SN. So are there any risks in using this for music? Do you think people would be disappointed if they know that algorithms have been used for the music?

GW. Some people would be and some people would think it was really cool. I think this is an entirely subjective thing. I, as someone who doesn't choose to use AI tools in my own creative practice, have indeed been deeply impressed by some of the things that have been achieved using generative systems in music. In 2014, I think it was. Maybe

in 2016, we ran a concert in London at the end of an EUfunded research project called Learn to Create. In this concert, apart from one piece by John Dowland, which was there as a sort of reference point, an hour and a half's music all was generated by computer, from nothing, not by taking existing work and repurposing it. Admittedly, it was out of learning from data sets of publicly available music, mostly by people who are long dead –so there are no rights issues there– and generating new things in their style. Quite a lot of what we did in the concert was about a particular style of music and generating that style convincingly. That was extremely successful and some of the music was really good; some of it was less good, but it was still quite interesting in the way that it was not so good sometimes. So, I think there's a lot to be learned and, of course, this is a classic from the jazz world. In jazz, if you make a mistake while you're playing, improvising, then standard practice is you repeat it. Because then people think you meant it. And what happens is that that can change the nature of a piece of music, right? Because you play a wrong note, that changes the harmony and then you think, oh, gosh, well, that's an interesting thing to do. Maybe we could go there. So one of the things I found in listening to our concert was precisely that I was hearing things allegedly in the style of such and such a composer, which that composer would never have done, but which were nevertheless interesting and inspiring to perhaps do other things.

SN. So it can be enriching to use the computers to get some new ways of composing.

GW. Absolutely. Potentially. So long as it's done in an ethical way, which does not infringe anybody's rights.

SN. That brings me to the beginning, the discussion in the beginning about these, the dangers. What do you think can be done to avoid these dangers?

GW. Well, there's a whole range of different dangers. There's, for example, the issue of control. if we have an autonomous AI system, if we give it too much power, too much control where we can't stop it doing things. We don't know what it is possibly going to do. And you end up in the science fiction scenarios where, you know, it decides that we should extinguish mankind, because of the fact that mankind is screwing up the planet or something like that. Okay, now, this is science fiction. But, there is no way that you can look at the code of ChatGPT and say that it is not planning to do that. It's not possible.

On that level, this sort of global threat level, I think there is some serious thinking that needs to be done by the politicians and the EU is now doing it to prevent misuse of AI.

It's pretty hard to prevent the misuse of something when the systems are not open to scrutiny. So you can't tell whether the system is doing the thing that is illegal or not. So what the EU is very sensibly legislating is what humans are allowed to do with the AI. And that way you get around that problem. So that's the big scary stuff.

On a much more small scale, personal level, what is going to happen here is that artists, craftspeople, musicians, whatever, are going to be put out of work. They are. They have already been put out of work. So when the synthesizer came along in the 1970s, it put session musicians out of work, especially when the sampler came along. Things like that. And that has now happened and there is nothing that can be done about it.

This happens every time there is a new technology. AI is not new in that respect. What is different about AI is that it has the capacity to learn, it has the capacity to adapt itself, and it has the capacity to do things which previously were the reserve of highly skilled people. And perhaps those people have a louder voice than, for example, the weavers who were put out of work by Hargreave's spinning jenny in the Industrial Revolution.

I think there's some sociological stuff that's going on here which is interesting because the people who are currently at risk are altogether more influential than the people were previously when they were put out of work. People who are put out of work by, for example, the combine harvester didn't have much of a political voice. But only today or yesterday we had authors or songwriters like Elton John and Michael Rosen in The Guardian arguing that AI needs to be controlled. And actually, I don't disagree with their principle, but they're making the wrong argument. AI does not need to be controlled. It's the people who use it that need to be controlled. AI does not pick up a huge database of music that is copyright and then learn from it. AI doesn't do that. People do that with AI. So there's the issue.

So I, speaking as someone who is a composer who has indeed had their work stolen in the past, I fully understand the concerns and the hurt it causes as well, the emotional hurt. That does need to be controlled. Unfortunately, the cat is out of the bag. And now that the cat is out of the bag, getting it back in I think is almost impossible. So that's a tricky one.

Another thing that's perhaps worth mentioning here is the following: I've heard a couple of times people asking, from the perspective of a particular songwriter, let's say Joe Bloggs. The songwriter, decides, realizes, or has a contention that his song has been stolen by a company and then used for training in AI system. Why can't we just make that claim and then have his song removed from the AI system? Because you can't. You

would have to retrain the whole thing and that would then mean using vast amounts of energy, and so on, to retrain it without that one song, which is just not an economically viable option. So, the issue here is a very complicated one in that individual's ownership is, in a certain sense, subjugated because their contribution, unwilling or otherwise, is not distinguishable from all of the other data in the system.

And that's another way that there is danger. In standard copyright law, if I listen to somebody else's melody and I repeat five notes of it in one of my melodies, then I can be sued for, copyright theft. However, in standard pop music style, it is not actually so unlikely that you might produce five notes that are exactly the same as another piece of music because there aren't twelve notes in the scale, there are only eight and, actually, only five of those are in common use. So, you know, the probabilities are not that small. There's a problem even with the current copyright law. Now, if you have a computer that is doing this, it's using statistical models, so it is somehow randomizing to generate this stuff somewhere along the line. That randomness might come from your original prompt or it might come from some deliberately random procedure or whatever. And if that's the case, then okay, so the system produces five notes which are from my song, but it may well not have done so by looking at my song. There's no way of telling statistically. And then because the systems are black box systems, you can't even go back and look at what it actually did because that is not available as a construct within the system.

SN. How can you tell whether we do the same when we compose something? We've been listening to music all our lives and maybe some ideas you may not even be conscious of where they comes from. We all have a culture and we constantly borrow from it.

GW. This is absolutely right. There's a couple of points there. So I do like to throw that point in when I'm talking to artists about this issue, because it is indeed a point that very many artists haven't thought about. Can any single one musician say they've never been influenced by anything? Of course they can't. That would be ridiculous. The issue though is that as someone who is supposed to be a creator in this area, you should reflect on what you're doing and make sure it isn't a copy of somebody else. You really should say, for instance: “that sounds too much like whoever it is”. So there again, we come back to this issue of reflection and of intention in creativity, which our generative models do not have, they're simply programs which do whatever they've been programmed to do, generating sentences, generating pictures or whatever. They don't decide to be creative. It's their human users who decide to be creative.

SN. That's right. Okay, well, I think this is a very good point to stop. Thank you very much for this interview, it's been most interesting. I hope we can collaborate again in some other time, in some other interview.

GW. That would be lovely. Thank you very much for listening.