[music]

00:06 Sarah Crespi: Welcome to The Science Podcast for May 31st, 2019, I'm Sarah Crespi. In this week's show, we start with staff writer, Katie Langin. She wrote a piece about how science graduate programs are dropping the GRE requirement like a hot potato and I also talk with Max Jaderberg of DeepMind about training an AI to compete and collaborate using a multi-player video game from 1999.

00:33 SC: We also have a book segment this week on an artificial intelligence and more importantly, we have a new book reviewer, Hickey Stanford. You may know her from the This Week in Science Podcasts or many other science communication efforts. For this month's book segment, she interviews Marcus du Sautoy about his book, The Creativity Code: Art and Innovation in the Age of AI.

00:58 SC: Now we have Katie Langin, a staff writer at Science. She's been trying to quantify the perception that graduate programs have begun the great GREexit or the GRE Exodus. Hi Katie.

01:10 Katie Langin: Hi.

01:11 SC: Okay, just a background question here. Did you take the GRE?

01:15 KL: I did take the GRE way back in 2005. I took the GRE and went through the process of studying for this standardized test to get into grad school.

01:27 SC: Me too and... You took the general or was it the biology one?

01:32 KL: I just took the general test.

01:33 SC: Me too. If you know people in the sciences these days, you've probably heard at least anecdotally that grad schools aren't really asking for the GRE as much anymore but that's kind of been something a school will announce or people will start talking about it but how did you go about finding out that this was more than a case here or there, how that was more than an anecdote?

01:57 KL: Yeah, so there's been this chatter on Twitter in particular, that the so-called GREexit movement has been gaining steam and a number of programs have been announcing that they've dropped the GRE. But I was really curious to quantify how many programs have actually dropped it and how that differs between scientific disciplines and so what I did is I sampled PhD programs at 50 of the top-ranked research universities in the US and looked at whether they required the General GRE test in eight different disciplines. So Molecular Biology, Neuroscience, Ecology, Chemistry, Computer Science, Psychology, Physics and Geology.

02:40 SC: And when you say top-ranked, where do those rankings come from?

02:44 KL: They came from The Times Higher Education Ranking System.

02:49 SC: Okay. When you looked at all these schools and all these different disciplines, did you... Were you able to kinda say yeah, something's happening here. The GRE is starting to go away in a big way?

03:00 KL: Yeah, so it was interesting going into this. I knew that there had been a shift in recent years but I didn't realize how dramatic it had been and what I found is that if I had collected these data just a few years ago, the vast majority of PhD programs would have required the GRE general test but the change has been really dramatic in the life sciences in particular. So in 2018, 44% of Molecular Biology PhD program stopped requiring GREs scores and just the year before, 100% of programs, at least in my sample required the GRE.

03:36 SC: That is a big change and you saw some decrease across the board.

03:42 KL: Yeah, so Neuroscience and Ecology, about a third of programs, did not require general GRE scores in 2018 and a lot of them are actually going to move to not require the GRE in 2019. So these numbers are on the rise for sure.

03:58 SC: In 2019, yeah. Were there any of these categories that were not doing this?

04:05 KL: When I looked at Chemistry, Computer Science, Psychology, Physics and Geology, more than 90% of programs required the GRE. But in a number of those disciplines, there does seem to be a movement as well. So Chemistry, there's a push to drop the GRE. Geology actually 100% of programs in 2018 that I surveyed required the GRE but there are a couple of programs already this year that have announced that they're gonna drop the GRE and so there seems to be movement in those areas as well. It just seems to be a little bit slower or not as dramatic as what I found in the life sciences.

04:42 SC: Well, I'm not gonna make a Brexit joke because I don't really understand what is happening there but what do we know about why these institutions are stopping requiring the GRE? Did you talk to people and ask them what was the main driver of this?

04:57 KL: Yeah, so there are a number of different reasons. One reason is that a number of studies have come out in the last few years questioning whether GRE scores are predictive of success in grad school, so admissions officers look at the GRE scores and a lot of people think that they are indicative of someone's intelligence, their ability to succeed but when you actually look at the data, in a number of programs, studies have found that they aren't predictive of things like the number of first author publications, the number of fellowships that students receive, their time in grad school, whether they graduate.

05:34 KL: Some of them have found that the GRE scores are predictive of first semester grades but for research-based PhD programs, I think a lot of people view that as not a particularly important thing to predict.

05:45 SC: Yeah.

05:46 KL: Another reason is that, especially in the current climate of wanting to increase diversity in science, there's concern that the use of the GRE actually disadvantages under-represented groups. So women and underrepresented racial and ethnic minorities tend to score lower on the GRE and it's kind of an expensive test for a lot of people that puts people from lower socioeconomic backgrounds at a disadvantage.

06:15 KL: So the test cost $205, it costs extra money if you want to retake it to up your scores and that's a fairly common thing for grad school applicants to take the test multiple times to get a better score. It also costs money to travel to take the test and it costs extra money to suppress GRE scores so you can take the test multiple times and then suppress the scores that you don't want schools to see.

06:41 KL: For an undergraduate student who's maybe working one or two side jobs to pay for their education, it can be a financial burden to pay for this test but it also can be a time burden. So they have to spend a lot of time studying. In some cases, people who have the means to do so actually pay upwards of $1,000 or more for a GRE prep course. And so, it really advantage those people who have the time and the money to study and pay for this test.

07:12 SC: It sounds like there's kinda a few categories of reasons that the universities are thinking about dropping the GRE as a requirement. But let's talk about the testing company that runs this test, what do they say to the fact that the studies are saying, this isn't correlating with things that departments particularly care about and there might be a barrier. What is their argument for continuing use of this test?

07:37 KL: From their perspective the GREs are another piece of data and they're a rare part of the application that's actually standardized across applicants. They also say that the other parts of the application, reference letters, educational background, they're also subject potentially to bias and so it's hard to drop the GRE and think that there is no bias remaining in the admissions procedures and so they argue that it's another piece of information, it shouldn't be used to make admissions decisions alone but it should be part of the package.

08:14 KL: They also take issue with some of the recent studies. So one of the drawbacks of these recent studies is that they're looking at admitted students and for the most part, those students have fairly high GRE scores to begin with and the students who have lower GRE scores probably really stand out in other parts of their application and so maybe they just had a bad day that day and they happened to get a lower GRE score. They didn't quite study enough.

08:39 KL: And so what they say is that we really don't have the best experimental design in those studies and that they're not really designed to find a correlation between GRE scores and these measures of success in grad school and so ideally, what you would have is basically a study that randomly admit students across the range of GRE scores and then look at success in grad school and it's hard to imagine a study like that ever happening.

09:08 SC: So the testing company feels that they provide enough of a service that institutions should continue to use this but it does look like the trend in the science and the science is that they're saying, there's not enough science there for this to be a test that we use and it also seems to create a barrier. So if this test is dropped more broadly, do you think it will change not only who's accepted but also who applies to graduate school?

09:33 KL: Yeah, that's a great question. So there are actually some people looking at this right now and looking to take advantage of the so called GREexit Movement. Some researchers at the University of Minnesota sent out a survey last month to Biomedical PhD programs that have dropped the GRE and also ones that have maintained the GRE requirement and what they're interested in trying to figure out is whether dropping the GRE requirement will diversify applicant pools and there's some argument that, that might be the case. Some students might do poorly on the GRE and maybe they can't afford to retake the test and up their GRE scores and so they just decide not to apply to grad school. Others might not be able to afford the test to begin with.

10:17 SC: What about the idea of making the GRE optional? So not making a requirement but adding it as an extra data point to your application.

10:26 KL: There are proponents of that and so one reason that might be a good thing is because there can be some difficulty for applicants that come from less well-known schools or maybe had some sort of disadvantage in their past that hinders their application, who could benefit from having a high GRE scores so it could help them get noticed in the application process. So people say well, you should give them the option of submitting their GRE score.

10:52 KL: There are other people who I spoke to, who think that that's not a very good idea because basically, when you give applicants the option of submitting GRE scores, they actually look at that and think that that is a requirement and one person I spoke to said that in his experience working with under-represented minorities, they in particular see that as a requirement because they think that that's the only way they can really compete with other applicants, is if they have a strong GRE score.

11:21 KL: Another drawback is that faculty members may look at applicants who don't submit GRE scores in a bad light, they might be biased against them because the applicants who do submit GRE scores are more likely than not gonna have high GRE scores. And so, it could lead to a question in the faculty member's mind, even if it's not conscious, even if it's at an unconscious level, they could be thinking, what about the student who didn't submit GRE scores? Is it because their scores were not very strong.

11:52 SC: Thank you so much, Katie.

11:53 KL: Thank you.

11:55 SC: Katie Langin is a staff writer for science, you can find a link to her story at sciencemag.org/podcast. Stay tuned for an interview with Max Jaderberg about training AIs to play capture the flag, the video game version.

[music]

12:14 SC: This week's episode is brought to you in part by KiwiCo. KiwiCo creates super cool hands-on projects for kids. They make learning about Science, Technology, Art and Math fun for everyone. With a KiwiCo subscription each month, the kid in your life will receive a fun, engaging new project which will help develop their creativity and confidence. The projects are designed to spark tinkering and learning in kids of all ages.

12:41 SC: All projects, inspiration and activities are created by a team of product designers in-house in Mountain View, California and rigorously tested by kids. Every crate includes all the supplies needed for that month's project. Detailed easy-to-follow instructions and an educational magazine to learn even more about the crate's theme.

13:00 SC: KiwiCo inspires kids to see themselves as makers and is on a mission to empower kids, not just to make a project but to make a difference. KiwiCo is offering Science Magazine Podcast listeners the chance to try them for free. To redeem this offer and to learn more about their projects for kids of all ages visit kiwico.com/magazine. That's kiwico.com/magazine.

[music]

13:32 SC: Now we have Max Jaderberg. He's here to talk about training an AI to play Capture the Flag. Well, video game, not in real life. Okay, hi Max.

13:41 Max Jaderberg: Hi. How's it going?

13:43 SC: Good. So this paper is about using reinforcement learning to train this artificial intelligence to play a cooperative first person point of view video game. This is a lot of things for an AI to do all of a sudden. Can you talk about what kinds of larger picture questions you're trying to address with this kind of learning or trying to teach an AI this type of game?

14:07 MJ: We chose this type of game for a few reasons. The first one is that it's inherently multi-agent. You don't just have one agent running about in the world trying to solve puzzles. [14:15] ____ is interacting with three other agents and other teammates and then two opposing agents. And so as soon as you have a multi-agent game, you introduce a huge amount of complexity into the game, which you don't find in single player games and this is a whole topic of AI research, multi-agent research. How to train agents to play in these multi-agent environments.

14:38 SC: You're not doing a two-player game, like Go or Chess?

14:41 MJ: The key difference between Go, Poker, Chess, StarCraft is, these are all one versus one games and you don't have that teammate or collaborative elements to the multi-agent problem which you have here. The other dimension that this environment, this Quake III Capture the Flag environment goes in, is that it's a full first person video game so each agent only sees it's view of the game from its first person view as opposed to seeing the full board in a game of Chess, the agent only sees a very small portion of the world, the pixels from its own point of view.

15:15 SC: Right. So there's no overhead view showing everybody on the board or any kind of locator map. There's just, whatever you see those are the only clues you have about what's going on.

15:25 MJ: Exactly, exactly.

15:26 SC: So, you mentioned the name of the game which is long, but it's Quake III Capture the Flag and it was first published in 1999. Why did you pick a game that came out so long ago?

15:37 MJ: This game came out a long time ago but really is one of the canonical first person multiplayer video games. It defined a lot of the mechanics of how to play these games and it was really popular back in the day and it had a big professional scene which is still alive and it actually served as really the inspiration behind many of the modern first person video games that you see these days.

16:04 SC: You talked about it being professional, what does that mean?

16:07 MJ: So that means it was so popular that it supported a professional scene where you would have players and teams of players competing in tournaments for cash prizes and people would have livelihoods based on playing this game.

16:21 SC: Wow!

16:22 MJ: Yeah.

16:23 SC: Let's talk about the training here. Can we start with what reinforcement learning is. That's kind of a big part of how you got these agents to play and play well.

16:32 MJ: So reinforcement learning is... It's all about training agents which take actions in our world. The key difference from something like supervised learning where you have human annotations of what you should do at every point and time. Here you don't have human annotations and so the agents learn by trial and error, by interacting with the world and then receiving pieces of reinforcement signal so a plus one, you did good at the end of a game for winning and you did bad for losing.

16:58 SC: Right. There is a visualization. I saw it on your website, I think in a video and it just shows all these agents in different little boxes. So they're all going up against a different team and over and over all at the same time, so how many play-throughs did your agents, were they subjected to?

17:17 MJ: Each agent through its lifetime for these paper results, played about 450,000 games of Capture the Flag. So that's, for each agent that's I think just over four years of experience of playing this game.

17:30 SC: Wow!

17:31 MJ: It's not just one agent playing this by the way, it's also 30 agents are all training in parallel. So that's what that visualization tries to get out is that there's a lot of stuff going on at the same time.

17:41 SC: It's the montage of the team all doing their trainings. [laughter]

17:44 MJ: Yeah, exactly. Exactly.

17:46 SC: That was one thing that really caught my attention and I don't really know the answer to is, we talk about these as agents and they all have different training, they're not all identical at the end. You kind of show it as an evolutionary tree.

18:00 MJ: Exactly.

18:01 SC: Do you have different agents at the end of this and they all have different strengths and weaknesses?

18:06 MJ: Yeah, absolutely. So at the end of training when we decided to name... Actually, we need to write some results down. We stop training and we have this whole population of agents, in our case it's 30 and they all have different strengths and we run a tournament and we pick the best ones to really evaluate. But there's a whole selection in that mix and as you've gone through trainings, you'll have a whole variety through training.

18:31 SC: Huh! Once you pick your gladiator, you face them off against a human champion and you see what happens. So when you ran these agents against humans as competitors or as collaborators, did you find that the computers were consistently better?

18:47 MJ: Yeah. We found the competitors were consistently better both as teammates and as opponents as full teams and interestingly, we ran these tournaments in a setting where humans and agents will be matched up but they wouldn't know who they're playing with or against. So they just play a series of games. Sometimes the teammates would be humans, agents. It's all mixed up. So that was a really interesting experience.

19:09 SC: Were these agents passing the turing test then? Can the humans not tell that they were agents.

19:14 MJ: I think at the beginning there was a lot of confusion and then by the end... We actually had this questionnaire. At the end of each game, we asked each player to try and identify whether their teammates and opponents were humans or agents and at the end of the day, I think humans could tell fairly reliably which plays were which but that was mainly based on... From questioning them it was based on how good they thought the other player was. So if they're a really good player, that they were playing with, they thought it was an agent. If they were bad they thought it was a human.

19:44 SC: And you had people who were new to the game, naive to the game but also people who were expert players.

19:50 MJ: We asked everyone who was participating in this to have had some first person video game experience and then, we also post hoc stratified the human samples into different skill levels. So based on their results, we had a section of strong humans, average humans and then we dropped the really bad humans.

[laughter]

20:12 SC: Was that someone on your team, perhaps? [laughter]

20:15 MJ: I can't comment on that.

20:16 SC: One thing I want to know is that if you saw any free loaders, anybody who had typical video game experience where someone just kind of hangs out and relies on their partner or camps out.

20:27 MJ: Yeah and actually some of the early days of this project, we had what we called the lazy agent phenomenon, which is where one of the agents in the team and this is in the setting where we always train the same two agents together so they're always in the same team and one of the agents could learn to just sit back and relax and stare at the sky, while the other agent did all the work.

[laughter]

20:50 MJ: And actually, by making the problem harder and not fixing the teams to the same two players and making it more this ad-hoc training scenario, where every game you're with a different agent as a teammate, you remove this lazy agent phenomena.

21:05 SC: Did you find that the agents use the same strategies as the human players?

21:11 MJ: To some degree, I think we can say that there are motifs that we definitely saw human players and agents do alike. So there were things like base camping, opponent base camping where you would sit in your opponent base. No one's there, the flag isn't there but you're sitting there waiting for the flag to reappear so you can immediately pick it up again and run home and we try to automatically identify these things by actually drawing inspiration from some people who are trying to model, for example, rat behavior and you can create these behavioral models and you can evaluate agents under these models and humans under these models from our recorded games as we try to relate some of the behaviors that we saw in both.

21:49 SC: So you extracted these patterns and then look for them in your agents?

21:52 MJ: Exactly.

21:53 SC: I kind of off-handedly said well, computers are gonna be better at computer games but I mean that's not true but they could be faster, right? They could react faster than a person. Was that something that you had to take into consideration?

22:06 MJ: Yes. This was a big part of the analysis and the process in developing this paper was yes, agents could react very fast so they had a hard limit of 67-millisecond reaction time. This is quite a bit faster than reported human reaction times. However, we first of all did some post-hoc analysis where we hindered the agents by reducing their tagging accuracy, so their ability to tag other players and we found that even reducing that to human level, many agents were still stronger than humans.

22:39 MJ: And then we did a further study where we actually trained agents, which had an in-built reaction time of 267 milliseconds and this is in line and comparable to human reaction times of video game players in simple psychometric studies and so then we ran a new tournament with these delayed reaction time agents and we found that even these agents were still stronger than humans.

23:02 SC: Was there anything that people were better at than the user agents?

23:06 MJ: Yeah. We also had a few exploitability studies. This was more taking a fixed team of two human professional games testers and letting them play against the same agents on the same map, over and over again. When humans are in this situation, they're really... There's a lot of ingenuity here and so these players managed to come up with some strategies that some of the time worked and fooled the agents. As the humans are learning in this case and the agents are completely fixed, they can't learn from this.

23:35 SC: Ah, okay. Okay. So you got these agents, they're good at collaborating, they're good at competing, they have these strategies that seem to work against expert players but you don't actually know what they know or what they're thinking. Training is just optimizing this ability to get a reward. So you then wanted to look at what they know and what they learn. How did you do that?

23:58 MJ: Yeah. I think this is the... One of the really exciting results from this paper is that we train these things with end-to-end reinforcement learning, which means we don't prescribe anything about the way they should behave, how they should perceive the world, how they should represent the world and then the game becomes actually, trying to dig into these agents and analyze their behavior and analyze their representation of the world and what was quite surprising for us was we could actually look at the representations that the agent had learned.

24:27 MJ: So look at the activity of the neurons within the agent's neural networks and correlate those to the underlying game state of the game and we found, for example, the agents really virtually actually represented the underlying state of the game and they even had individual neurons which fired for things like I am holding the flag or my teammate is holding the flag. You found these individual neurons which fire in these different circumstances.

24:56 SC: Wow! It's just so interesting.

24:57 MJ: Exactly.

24:58 SC: I don't have a question. [laughter]

25:00 MJ: It's like we're really lucky neuro-scientists in that we can see everything about the patient's brain and we know everything about the underlying thing that they're looking at. So it's really great.

25:09 SC: Can we talk a little bit about how transferable this is? What other things we might be able to do with agents with these skills?

25:16 MJ: Do you mean in terms of transferring these agents into different games or...

25:21 SC: Yeah, sure. We can start with games but I would like to talk about the broader picture too.

25:25 MJ: I don't think we can directly transfer these trained agents into another game but I think the things that are transferable are the training methodologies that have been developed here. A lot of these things are actually very general, things like population-based training, things like learning internal rewards in these multi-time scale agents. These are components and algorithms that we can actually transfer to other games and other domains quite successfully.

25:50 SC: Max, thanks so much for talking with me.

25:52 MJ: Thank you very much.

25:53 SC: Max Jaderberg is a research scientist at DeepMind. You can find a link to his paper at sciencemag.org/podcast. Don't forget to keep listening for this month's book segment with our new reviewer, Kiki Sanford. She talks to Marcus du Sautoy about his book The Creativity Code: Art and Innovation in the Age of AI.

[music]

26:19 Dr. Kiki Sanford: Welcome to the book segment of The Science Podcast. I'm your new host Dr. Kiki Sanford and this month I had the pleasure of speaking with Dr. Marcus du Sautoy, the Simonyi Professor for the Public Understanding of Science and Professor of Mathematics at the University of Oxford about his most recent book, The Creativity Code: Art and Innovation in the Age of AI which delves into the influential role played by artificial intelligence on our understanding of creativity and as artificial intelligence exhibits more and more creativity in historically human endeavors like art, music, science and even games, Dr. Du Sautoy asks how we can use these creative tools to help us be more creative ourselves. How do you define creativity?

27:07 Marcus Du Sautoy: I think it's one of those rather mysterious way it's creativity that's quite hard to pin down and so I actually used the working definition in the book which actually comes from a woman that I talked to a little bit on this project called Margaret Boden and she suggested the idea of a creative action way one which is novel, new, well, certainly computers can make endlessly new things but it should also have an element of surprise in it and it should also have value.

27:35 MS: Now, what's interesting is that novelty is something that we can objectively identify but surprise and value are very subjective, they depend on things which can vary from one person to another so what's interesting about the AI that's emerging is that it's learning from data, it's not code that's written in a top-down manner where if someone here write the code and then it's frozen. This is code that is written in a bottom-up way where it re-parameterize itself according to its interaction with data. So this new AI can begin to learn what we find surprising, what we value and can offer it's own surprises and things which it thinks we might value and so I think that idea of novelty, surprise and value is quite a useful definition for a creative act.

28:27 DS: You're mentioning the way that these machine learning algorithms learn through ingesting these massive amounts of data and having this bottom-up strategy, this is very similar to the way that humans learn where we read information that has been created prior to us, ingest it and then we're able to interact with that data in a way that enables us to create something new in a sense and so do you have any thoughts on this similarity in methodology between the new AI and human learning?

29:04 MS: Yeah, I think that... You've put your finger on it, exactly. That it's actually the way humans learn and develop and I think one of the challenges with AI creativity is that well, surely isn't the coder, the human who wrote the code who's the one who's truly being creative and I think in the past that certainly would be true but what's new here is that we have a very rich digital landscape for this kind of machine learning and this AI to actually learn, change, adapt and evolve inside and so I think we are seeing something which is quite similar to the learning process that a child will go through.

29:42 DS: Similar but not the same. Are there still important differences in the learning processes?

29:47 MS: If you take language, a human child does not need much exposure to data to be able to start speaking a language and so this is an indication that the human brain is still coming with quite a lot of very sophisticated pre-programming, it isn't just exposure to data but the way AI is learning language is by being exposed to data and it's picking up the structure from that data. So I think what's interesting is that tension between how much is being learned from exposure to things that are in our environment and how much is actually coming pre-programmed.

30:25 DS: What about the emotional aspect of creativity, that's still at this point is something that AI lack to a large degree. Do you think this will continue to be some element of defining the difference between creativity in humans and AI?

30:42 MS: I think unfortunately it already isn't the difference. There's a lot of evidence that AI is able to read our emotional world much better than humans can. There's some work being done out of MIT Media Lab revealing that an AI can identify a false smile in a human much better than a human can so AI is already being able to read our own emotional world but also in the book I try to explain that emotions are often a secondary response in creativity.

31:15 MS: You'll listen to a piece of music and you will have a huge emotional response to it. When you talk to the composer, they will often say, "Well, I didn't start with that emotion. I started with structure and pattern and an algorithm that I was using and something quite mathematical but out of that structure came the emotions."

31:33 DS: So the perception of creativity has more to do with the receiver than the creator?

31:39 MS: There's ambiguity in art and that ambiguity is really important in your interaction with art and some of the most successful examples of algorithmic art have been where they made the production of art into a game. One of the interesting examples is there's something called a Creative Adversarial Network where you have two algorithms somehow competing against each other. One is trying to create something new which doesn't fit into any genre that it's learned about art in the past, the other algorithm tries to say either say, "No, that's too similar to a pointillist art or cubist art" or either says, "Well that's not art at all, you've gone way too far and it's just a mess." When shown at Basel Art Fair, people weren't told that they were created by AI and they actually had quite a positive emotional response to these pieces.

32:29 DS: This sounds almost more telling about the way that we create and perceive art than the AIs themselves.

32:35 MS: The book is as much about human creativity as it is about AI creativity trying to reveal that human creativity isn't such a mysterious kind of thing as we think it is. It's perhaps got more structure and algorithm and kind of intention than we realize, actually much of human creativity has it's origins in learning from things that we've done in the past and I'm putting them together in new ways.

33:03 MS: But I think that what's interesting is where I think AI can be really helpful in being a tool to stimulate our own human creativity. I think we get very often stuck in ways of thinking and we become too much like machines and so, what is exciting is maybe the AI, although itself may not ultimately be a great artist, maybe it can just offer things to us which will help kick us out of behaving so much like machines and perhaps help us to be once again creative humans.

33:33 DS: You mentioned in a part of the book, the work that is coming from the DeepMind program, which created the AlphaGo program that beat the world's leading Go player and is now working on math! Do you think you're really threatened by these algorithms or not?

33:52 MS: I was sitting next to Demis Hassabis during a committee meeting and where we were looking at the impact that machine learning, this new AI was gonna have on society over the next 10 years and I kind of cheekily said to him, "Well, you know, do you think you could prove a mathematical theorem and get an AI elected to the Rural Society here in London, our National Science Foundation?" And he said, "Well, I'm already on the job." And so I was a bit shocked 'cause I obviously meant this as a joke.

34:21 MS: I think there's a very exciting brand emerging, a lot of the AI of course is aimed at business applications, health, politics but certainly DeepMind and here in London and Demis in particular is also interested in how this can be a tool for fundamental research. To spoil the ending of the book, what I discover is that the sort of mathematics that is coming out of this new work, I don't think is a threat. I don't think it is what I regard as mathematics and I try to make this point, that mathematics isn't just about proving all the true theorems about numbers and geometry, it's a lot about making choices though telling the stories about numbers and geometry that surprise us, that we get excited by.

35:04 MS: I think we are story tellers and our characters are numbers and geometry and ultimately it looks like one of the hardest things for AI to do, although it can make quite convincing music, it can paint quite well. What it's very bad at is storytelling. So I think once I see AI being able to write a prize-winning novel, that's the moment I think I'd really worry about the fact that mathematics might be possible to also be done by a computer.

35:30 DS: I'm fascinated by the idea of humans and the language ability but that being also related to metaphor and imagery and that we can imagine things because of words and that it's our brain being able to turn something into a visual or auditory experience that makes it possible, do you travel? Do you travel in your brain when you're imagining different maths descriptions?

35:56 MS: I think it's a very important metaphor for me, actually, that one of exploring a landscape. I'll often feel like a conjecture is a far distant peak and I've got to find some way, a little bit like Frodo in Lord of the Rings and it feels like an epic journey that you're making to prove a theorem. And I think that's when I'm doing my mathematics, I sort of try to get into a meditative state where I am sort of almost exploring in my mind that landscape, 'cause I have the freedom to make things which have no physical reality, I can make worlds that are different to ones that describe our physical universe.

36:31 DS: What about having a brain connected to a body that allows sensory feedback with and from the outside world, the concept of embodiment, is it important?

36:43 MS: A lot of the research in AI at the moment, is trying to explore the importance of having a body to understand things and interact with our environment around us, that even the mathematics that we create, although it seems un-embodied, very abstract, still very often and has it's origins in measurements or counting in things which are embodied. Could AI produce a different way of seeing things because it isn't embodied very often and I think that's what's exciting to, actually not replicate our own intelligence, we know how to replicate our own intelligence, we have children and students but perhaps to create a new intelligence that might be able to take us in new directions.

37:22 DS: You bring that up in the book a little bit, especially with regards to the speed of computer processing and the fact that computers can be much faster than us. And so will they potentially create music at some point that we can't really hear or will they create language that we need to interpret? Do you think there's going to be AI culture?

37:43 MS: There's already kind of an example of that emerging. There's a story in the book about robots creating their own language amongst themselves, this interesting idea that there can be convergence. They all make up their own languages to start with but by sharing their languages by the end of the week, they've actually converged on a common language. We're not at that level yet because I think the AI is still connected to us as humans. It may be doing things that we don't really understand but the intention is still coming from us as humans.

38:10 MS: We should see this as collaboration but that collaboration needs us to be able to understand what we're doing. The book is trying to give people the tools to be able to understand the algorithms, what machine learning is. I really try and give people access to looking underneath the bonnet and not just taking on trust what these things are. Going forwards we're gonna need to know how our algorithms are pushing and pulling us around if we want to have a positive collaboration with the emerging AI.

38:43 DS: Thank you for joining me for this interview with Dr. Marcus Du Sautoy about his book, The Creativity Code: Art and Innovation in the Age of AI. I'm Dr. Kiki Sanford and I hope that you'll join us again next month for a peek between the pages of another science book.

[music]

39:00 SC: And that concludes this edition of The Science Podcast. If you have any comments or suggestions for the show, write to us at sciencepodcast@aaas.org. You can subscribe to the show on iTunes, Stitcher, Spotify, many other places or you can listen on the Science website at sciencemag.org/podcasts. There you'll find links to the research and news discussed in the episode.

39:23 SC: To place an ad on the podcast, contact midroll.com. This show was produced by Sarah Crespi and edited by Podigy. Jeffrey Cook composed the music. On behalf of Science Magazine and its publisher, AAAS, thanks for joining us.