Catalyzing Computing Episode 16 - Interview with Melanie Mitchell Part 2

The transcript below is lightly edited for readability. Listen to “Interview with Melanie Mitchell Part 2” here.

[Intro - 00:10]

Khari: Hello, I'm your host, Khari Douglas, and welcome to Catalyzing Computing, the official podcast of the Computing Community Consortium. The Computing Community Consortium, or CCC for short, is a programmatic committee of the Computing Research Association. The mission of the CCC is to catalyze the computing research community and enable the pursuit of innovative, high-impact research.

In this episode of Catalyzing Computing, I sit down with Melanie Mitchell. Melanie recently joined the CCC Council and is a Professor of Computer Science at Portland State University, as well as an External Professor, a member of the Science Board at the Santa Fe Institute. She is the author and editor of five books and over 80 scholarly papers in the fields of artificial intelligence, cognitive science and complex systems. Her most recent book is Complexity: A Guided Tour won the 2010 Phi Beta Kappa Science Book Award.

In this episode, we discuss genetic algorithms, Melanie's work with visual recognition systems, as well as complexity science and the art of writing a book. This is part two of my interview with Melanie Mitchell. If you haven't heard part one, and would like to, go catch that and come right back.

Enjoy.

[Genetic Algorithms - 1:16]

Khari: I guess you mentioned it earlier, but can you discuss genetic algorithms? What is a genetic algorithm and how could they be useful in designing AI systems and machine learning systems?

Melanie: Genetic algorithms are inspired by Darwin evolution. Suppose that you wanted to evolve a computer program instead of writing a computer program. What you might do is start with a population of random computer programs — and you'd have to define what that means — and then you want them to perform some task. Just by chance, one of them might be a little bit better at performing the tasks than all the others. Just by chance, because they're random. Then you take that one and you let it create offspring that are variations of itself, sort of by random mutations, let's say. Then you create a new population of these random mutations of the old one and let them compete to solve the problem. And you do this over many, many generations. You can also incorporate things like crossover between individuals and your population. And over many generations, just starting from this random initial state, you could evolve programs that actually perform very well.

They don't have to be programs, they can be other kinds of designs, for instance, protein sequences or architecture designs. People have tried them on all kinds of things. But that's the idea: using this idea of variation, selection, heritability, and competition. These are all the key ideas of Darwinian evolution and applying them to computers. So this is what John Holland came up with back in the 1950s, essentially, and other people came up with other variations and it became kind of a big subfield of AI, this idea of evolutionary computing.

Nowadays, like every other part of AI, it was kind of beat out by deep learning, deep neural networks. But, it's kind of becoming rediscovered because it seems like evolutionary approaches are a very good way to learn some aspects of neural networks. So you might evolve neural networks rather than trying to train them using other techniques. And people have been doing this quite successfully, surprisingly. So, I think it's making a comeback.

Khari: Yeah, I know the CCC recently held a workshop on thermodynamic computing. Do you think this is an area that's sort of similar in terms of allowing computing systems to evolve naturally? That's not necessarily an algorithm, but...

Melanie: Right. Well, it's a kind of computing inspired by nature. You know, I can't say I understand enough about thermodynamic computing to say how it might be similar, but there is a long history in computing of biologically inspired approaches to computing and genetic algorithms or evolutionary computing is one biologically inspired approach. One very interesting aspect of it is what people call co-evolutionary computing and that's where you might have two populations, one that you might call the hosts, and the other, which you might call the parasites, and they try and compete against each other, kind of inspired by nature.

The hosts are trying to compute something, the parasites are trying to prevent them from computing it well and they each evolve with respect to the other population. There's been quite a lot of work on that, and it's very related to some of the adversarial approaches people are trying now with neural networks.

Khari: Ok, so what role does the parasite play in terms of developing the algorithm? You said it's trying to prevent the algorithm from being successful.

Melanie: Yeah, exactly. You might say the hosts are programs and the parasites are test examples that the programs are being tested on. The hosts try to give the correct answer on the test examples and the test examples themselves evolve to be increasingly harder or trickier for the programs. So the programs get better and better and the examples get harder and harder. So that's one example.

Khari: Ok, that's clever. How logistically do you get a genetic algorithm to evolve? Obviously, like in terms of breeding, you know, people understand the birds and the bees or whatever, but how does a program advance itself?

[Laughter]

Melanie: Well, the first thing you have to do is figure out how to represent a program in a computer in a way that it can evolve. One way that people have done that is with tree structures. I don't know how technical I should get here, but you can represent a program as a data structure that represents a tree-like structure, then you can have programs kind of cross over with each other, where one program takes a branch from another program and adds it to itself. Then you might test these programs on some data and see how they do — that score is called their fitness. Then you select the programs that have the best fitness and let them procreate with one another, and you also have random mutations where you change the code a little bit and then you create a new generation that way and then start over. So you can go through this with many generations and see how these programs can improve.

Khari: How do you build that initial tree? Like, let's say you're trying to do image recognition.

Melanie: You might set up some primitive functions and kind of randomly put them together in a pipeline, for example, to do some kind of image processing sequences. They would do something very random because they're random programs, but just by chance, a few of them might do the task slightly better than the other ones. So people have applied this to disparate domains like creating artworks, creating music, creating designs in, say, mechanical engineering. There is a group at NASA who worked on sort of designing antennas for spacecraft, and they used genetic algorithms to design an antenna that actually was cheaper and better than anything their engineers had designed. There's all kinds of examples in the field of amazing designs that these systems have come up with.

[CCC and AAAI's AI Research Roadmap - 7:47]

Khari: Ok, so let's see. The CCC just came out, in collaboration with AAAI, this big Artificial Intelligence Roadmap. Have you read it? Have you seen it?

Melanie: Yeah, I have. I've read it, but I didn't participate in creating it.

Khari: Do you have any thoughts about things you think are pretty noteworthy that seem like they will advance the field, or things you think are missing or would not work?

Melanie: Well, it was a very long report.

[Laughter]

Melanie: It has a lot of different ideas. The thing that I'm excited about these days is having AI take more inspiration from cognitive science. You know, when I got into the field, AI and cognitive science were very close. People from AI would often go to cognitive science conferences and vice versa, and that's gotten less so since, I don't know, maybe since AI has become more commercialized. But that link where people were thinking very hard about how humans think, or how other animals think, and how we might use those ideas in AI or how we could use these ideas in AI to understand how people think, that link has been broken a little bit. In the AI Roadmap, I did see some ideas for trying to reconnect the field of AI to psychology, cognitive science, animal behavior, etc. So that's something I'm pretty excited about, getting the field back again to be more influenced by other related fields.

Khari: Okay. Yeah, that definitely seems very helpful. Anything in particular you thought seems like it's not the best idea? Not necessarily that it's bad, but maybe it would be a waste of resources.

Melanie: Well, one thing, I thought was not brought out enough in the roadmap was the way AI researchers or practitioners communicate with the public. This kind of gets back to my fallacy about how we see “AI is better than humans on reading comprehension” or those kinds of headlines. I think that part of the roadmap should be thinking a lot harder about how to communicate the abilities and limitations of the AI systems that are built and what incentives might align better with what outcomes we want. You know, right now the incentives seem to promote a lot of over-promising or over-overestimating the abilities of these systems, and then people get disappointed when they find out that the systems aren't as advanced or as reliable as the picture that was painted. So that, I think, was not really part of the roadmap and I think that kind of communication is going to be really essential as the field moves forward.

Khari: Alright. Well, hopefully with you now on the CCC council, that could be something that you can help to spread and help us to communicate.

Melanie: Yeah and I'll just mention that I just finished writing a book about AI that's going to be published in October. It’s really an account for a broad audience about the state of the field, where I see the state of the field, and I hope that helps people think more clearly about where the field is. There are several other books coming out kind of along the same lines, so I'm hoping that this sort of communication problem will start to get better.

Khari: Okay and what's the book called?

Melanie: It's called *Artificial Intelligence: A Guide for Thinking Humans*.

Khari: Good title, captivating. Is that available for preorders now?

Melanie: It's available for preorder on Amazon and any other book sellers that you want to go to.

Khari: Alright, sounds good. Well, if you're listening, you can go find that book.

[Semantic Image Retrieval - 11:57]

Another thing that I saw you had worked on, sort of recently, was this paper "Semantic Image Retrieval via Active Grounding of Visual Situations." I believe the architecture is called Situate. Is that right?

Melanie: Right, Situate.

Khari: So how does that work? Why is it novel?

Melanie: So that actually is the work that I'm doing that is inspired by the old copycat project that I worked on for my dissertation. The idea is, we have computer vision systems that are really good at object recognition — they're not perfect, but they're pretty good. What we don't have as much are systems that can interpret whole visual situations. So one example that we use in the paper is a situation like walking a dog.

You have a person holding a leash, the leash is attached to the dog, they're both walking, usually outside, and that situation consists of objects that are in stereotypical relationships with one another. But you can have lots of variations of that situation, like somebody walking multiple dogs, somebody running with dogs, somebody on a bicycle holding a leash with the dog, we even see people walking cats.

[Laughter]

Melanie: So the idea is to have a program that can, using analogy, interpret an entire visual concept like that. Not just the individual objects but the whole concept. There's a lot of people working on this kind of thing because, obviously, we want to be able to have systems that make sense of entire images, not just objects, but it's a very hard problem. My work is using the assumption that we want to integrate high-level knowledge about situations with low-level perception and the way that we did in the copycat project where there's this continual dynamic interaction between the two. In that paper, we talked about the initial implementation of our ideas and compared it with some other approaches to image interpretation and showed that it actually did a better job when we tried to do image retrieval using this idea of situations. So that's ongoing work with my collaborators and my grad students.

Khari: Ok. So what level of success have you had to this point? You said it was more successful at image retrieval.

Melanie: We compared it with one other system that tried to do the same thing. How successful is it? Well, it's still in the realm of very basic and preliminary research, but I think it's a promising way to go. There's a lot of people who are working on image interpretation, either generating captions for images or trying to detect visual relationships and so on, but I don't see a lot of people trying to do this via analogy. I think that's one of the novel approaches that we're trying to do. We're trying to use prior knowledge and be able to make flexible analogies between that prior knowledge and a given situation.

[Complexity Science - 15:19]

Khari: Ok, very interesting. So I know you've also done some work with complexity science.

Melanie: Yes.

Khari: So what is complexity science? I'm assuming it's difficult, but...

[Laughter]

Melanie: Right. So complexity science is the study of emergent behavior. That is, you have a system that's composed of many relatively simple components, like the brain is composed of neurons, an ant colony is composed of individual ants or the immune system is composed of billions of cells. In all these different systems you have this phenomenon where the individual components are relatively simple compared to the system as a whole, and yet you can get this very complex, unpredictable behavior coming out of the whole system.

Complexity sciences, is really the study of how this can happen, how this kind of emergent behavior comes about. Then questions come about like, can we actually control these systems? Can we make them more effective? Can we predict what they're going to do? So it's a very broad umbrella for a lot of different fields that sort of relate to one another, ranging from cellular biology to sociology.

Complexity science is like computer science. Is it actually a science? Someone once said that any field that has science in its name isn't actually a science, with the exception of material science, which actually is a science.

[Laughter]

Melanie: That's just a joke. I don't know if that's really true, but complexity science, or people sometimes call it complex adaptive systems, is really a big umbrella for many different fields that are trying to understand what the relations are of the subsets.

Khari: Right, yeah. So is this a similar idea to chaos theory? My knowledge of chaos theory is mostly limited to reading Jurassic Park.

[Laughter]

Khari: Is it a similar sort of concept?

Melanie: Yeah, so chaos theory is one part of this bigger field called dynamical systems theory, which is one of the core frameworks people use in complex systems to understand systems. They want to understand the dynamics of these systems, what the dynamics mean, how do these systems actually change over time? And chaos is one kind of dynamic that you sometimes see. So I think of chaos theory as one particular phenomenon that we see in complex systems. That's kind of the relationship.

Khari: How is chaos different from other kinds of complex dynamical systems?

Melanie: So think about, for instance, the heart. Our heart beats regularly. All the cells kind of work together so that you get this constant heartbeat. That's a kind of dynamic, it's an oscillation, it's not chaotic. If your heart starts to behave chaotically, you're in big trouble. Those are two kinds of dynamic behaviors, as sort of regular oscillations and chaos. Then when you die, suddenly you get another kind of dynamic behavior, which is nothing. Your heart stops beating. You know, that's kind of a glib example, but it illustrates that there's these different regimes of behavior that systems can be in. Similarly, look at a dripping faucet. That's another classic example where you get a drip...drip...drip...That's an oscillation, it's periodic. But all of a sudden you start to get drip...drip.drip.drip.drip...drip...drip.drip.drip [sporadic]. It can go into chaotic behavior. So dynamical systems is the study of all those kinds of behaviors very generally across systems.

Khari: Ok, that makes sense. So you've worked on a project through the Santa Fe Institute called the Complexity Explorer. Can you talk a little bit about that?

Melanie: Right, so Complexity Explorer is an educational web platform in which Santa Fe Institute offers free online courses and tutorials for people about complex systems. This is something they do as kind of a public service, and I was one of the founders of this project because I wanted to give people who weren't at the institute a way to get the experience of being at the institute and engaging with it and learning about what the science was all about. I taught one of these courses. I made a video series of exercises and demos, and it's called Introduction to Complexity and you can get there by going to https://www.complexityexplorer.org/. And it's really accessible for all. Then there's some more advanced courses that require a little bit more of a mathematical background. But I think there's about 10 courses and maybe about 20 tutorials, which are short courses on various aspects of complex systems. It's a really great resource.

Khari: Ok. I know you guys offer complexity challenges. What are those?

Melanie: That's for people who want to do a project. The idea is the Santa Fe Institute has this network of companies who are associated with it and one of the companies presents a problem that they're facing that has to do with some aspect of complex systems. They put it out there as a challenge and this allows people to kind of learn by doing more. It's really a hands-on project course where people can work together remotely. You can meet other people who are interested in working in this field and it's a way to get involved with some complex system scientists who are interested in these topics. Rather than just watching videos, you're actually doing a project and working with other people.

Khari: Very cool. So what level do you need to be at to do one of these challenges?

Melanie: So I think to do a challenge, you probably have to know how to do some programming and some mathematics. Maybe a college level science/math background or programming or computer science.

Khari: Okay, but it's accessible to anyone with some programming and math experience? You don't need a P.h.D or anything?

Melanie: No, no. I think even undergrad, a really good undergrad student, or even possibly a high school student could be involved in some of these projects.

Khari: Okay, so definitely check those out. Do you get anything if you win the challenges as a prize or anything like that?

Melanie: You know, I don't remember. This is something that I didn't spearhead. You probably get a T-shirt and maybe you get a job offer from the company, I don't know.

[Laughter]

[Science Outreach and Writing a Book - 23:13]

Khari: Well, that's a pretty good incentive, I guess. Have you done any other projects with regards to complexity recently that you want to discuss, any interesting outcomes or anything?

Melanie: So I wrote a book called *Complexity: A Guided Tour*, which is really what I based my complexity course on and is another broad audience introduction to these ideas. That was a big project and all of these were trying to communicate some very complex ideas, if you will, to people who don't necessarily have a science or math background. As you can see, I kind of split my time between research, education, and outreach. I try to do a lot of all three where I'm really interested in outreach to the general public about scientific ideas.

Khari: So how do you manage those three different things in terms of education, research and outreach?

[Laughter]

Melanie: Not very well.

Khari: Do you have any tips for listeners that might be trying to do all three things?

Melanie: Well, for me, the most important thing in research is to have a good group of students and collaborators. That's the most important part because you can't do it all yourself and that also gives you some time to try and do some of these other kinds of projects. All of the books I've written have taken quite a long time and also involve various sabbaticals and so on. It's not easy to do all three kinds of things, plus teaching, but it's what I find the most rewarding: trying to engage with students and the public about some of these ideas.

I'm a big fan of citizen science, these initiatives that try and get the general public involved in actually doing scientific projects. That's what the complexity challenge is and so I'm really thinking hard about how we could bring citizen science more into computer science, in particular in A.. So stay tuned for some ideas on that.

Khari: Have you done any outside of the complexity explorer? Any other citizen science kind of projects?

Melanie: No, I haven't. I was just very intrigued by it.

Khari: Yeah. And how many books have you written?

Melanie: Let's see… I've written four books and I've also co-edited two collections of essays.

Khari: Okay and how long on average does it take you to write a book?

[Laughter]

Melanie: Well, let's see. The first book was a version of my Ph.D. dissertation. The second book took about a year and the third book took about 10 years. The reason for that was that I had two kids while I was writing that book. So I was balancing being a mom with a full time job and writing. Then this last book took about two years to write, but that was with a one year's sabbatical included. So it varies, but it's definitely a long haul.

Khari: Right. Do you have any tips for both scientists or nonscientists in terms of writing? Like writing habits or...I mean, you worked for 10 years, so, sort of, persevering?

Melanie: Yeah. During those 10 years, I quit a few times. I thought, “This is never gonna happen.” But I think, you know, just persevering and trying to write everyday. I try to get even a small amount done everyday, so that I don't lose the flow of it. It doesn't always work, but you just have to kind of keep at it. I like to write in a stream of consciousness style at first just to get stuff on the page and then go back and edit it. I found, like with a lot of my students, when they're trying to write, they get blocked because they feel like they can't write anything good. I think the trick is to just write bad stuff, just to write terrible, terrible stuff and then come back and edit it. That's a lot easier than trying to write good stuff in the first place. At least for me, that's the way I usually work.

Khari: I think there's a quote, which I think is misattributed to Herman Melville, “write drunk, edit sober.” Not that you need to literally be drunk, but it might help not being so inhibited and just getting the words out.

Melanie: Yeah.

Khari: You just joined the CCC council this summer. So as a member of about a month, how did you first hear about the CCC and why did you decide to become a council member?

Melanie: I'd heard about the CCC for a long time through my professional life as a computer scientist. I'd known about the CRA and I subscribed to the CCC mailing list and received notices about what was going on. My good friend and colleague, Liz Bradley, who has been on the CCC for several years, told me that they were doing this AI Roadmap. She asked me if I would read it and give some comments, which I actually didn't do before I got on the CCC, I didn't have time. But then I guess Liz and possibly some other colleagues of mine nominated me and asked me if I wanted to be a council member. It sounded like a great way to try and engage in computer science with a broader set of people so I was excited about that.

Khari: Yeah, we're happy to have you as a member. I know it's only been one month, but do you have any specific goals or things that you hope to accomplish during your time on the council?

Melanie: I'm still kind of learning what the council does. You know, as I said, I'm really interested in communication and engagement with the public and I'm hoping that that's something that I can help focus on with the council. Today computer science has such a big role in the world, and yet so few people really understand very much about it at all and I think that's something that the CCC is really well poised to contribute to: kind of a more public understanding of what's going on in the field and also trying to help, not only the public, but computer scientists to understand what's going on in areas other than their own particular area of expertise. That's one thing that I'm interested in, but I'm still kind of getting on board.

Khari: Yeah, well, you have plenty of time so no rush, but that all sounds great. Obviously, spreading what's going on in different domains of computer science is a big part of what the CCC tries to do. So I guess related to the communications thing for your wrap up... Have you ever considered something like running for office or working on some sort of TV show or anything like that?

Melanie: No, I haven't. That's an interesting question… running for office? No. Gosh, one of the reasons I went into academia was because I like to sit in a dark room and write, so the idea of going out and giving a lot of speeches is terrifying to me. But it's an interesting idea. I haven't really been involved with policy very much and one of the things that I'm interested in with the CCC is seeing more of what that's like and trying to see if that's something that I can contribute to in some way.

Khari: Yeah. If you're interested in that, I certainly encourage... the CCC council meeting in the Fall is going to be just ahead of LISPI (Leadership in Science

Policy Institute). I'm not sure if Peter presented this during the last council meeting.

Melanie: Yeah, he mentioned it.

Khari: It's basically a day where they coach you on how to talk to Congress and congressional staffers and stuff like that, and then the congressional fly-in is sometime also in the Fall when they send you up to the Hill.

Melanie: Oh, boy. Yeah, that does sound really interesting, completely different from anything I've really done before.

Khari: Yeah, so definitely a good opportunity if that's something you or listeners are interested in participating in. So that's all the questions I had. Anything else you want to pitch or bring up?

Melanie: I think we covered a lot of stuff, so I don't have anything else.

Khari: Alright. Well, thank you for taking the time out of your day to be here.

Melanie: Well, thanks for having me on this podcast, really happy to do it.

[Outro - 33:01]

Khari: That's it for today's podcast, I hope you enjoyed it. We'll be back soon with new episodes. Until then, like, subscribe, and rate us five stars on iTunes. If you have any questions for former guests of the show, write us a letter at cccpodcastletters@cra.org and we'll try and get them to answer you. Until next time, peace.