AI 4 Science Discovery Network+

AI4SD Interview with Dr Grant Hill
26/04/2021
Online Interview

Michelle Pauli
Michelle Pauli Ltd

28/06/2022

Humans-of-AI4SD:Interview-20

Contents

1 Interview Details

Title	AI4SD Interview with Dr Grant Hill
Interviewer	MP: Michelle Pauli - MichellePauli Ltd
Interviewee	GH: Dr Grant Hill - University of Sheffield
Interview Location	Online Interview
Dates	26/04/2021

2 Biography

Figure 1: Dr Grant Hill

Dr Grant Hill: 'The holy grail is for the AI to come up with something a trained chemist would not have. . . '

Grant Hill was a student at the University of York, obtaining an MChem degree in 2002 and a PhD in theoretical chemistry in 2006. He spent 2005-2008 as a postdoctoral researcher in the group of Dr. Jamie Platts at Cardiff University, and 2008-2010 at Washington State University in the group of Prof. Kirk Peterson. After a temporary lectureship and short period as a Leadership Fellow at the University of Glasgow, he joined the University of Sheffield as a lecturer in 2014.

In this Humans of AI4SD interview he discusses using AI to search for a solution to a global sustainability goal, the impact of Covid on theoretical and computational chemistry and why ECRs should be looking out for collaboration projects.

3 Interview

MP: What's been your path to where you are today?

GH: I'm a senior lecturer in theoretical chemistry at the University of Sheffield. I grew up in the Lake District and I went to university at one of our local rivals, the University of York, where I got my PhD in theoretical chemistry. I did two postdocs — at Cardiff University, then Washington State University — before I came back to the UK to try to find a permanent position. I had a series of relatively short contracts at the University of Glasgow, before I got a fellowship and moved to Sheffield in 2014, where I've been ever since.

I've been interested in computers since primary school, where I was playing around with a BBC Micro Model B. Computer programming has always interested me. During my undergraduate years, I realised that I really liked research and that I wanted a career in computational or theoretical chemistry if possible. The semantics of what's computational chemistry and what's theoretical chemistry can start a bit of a war in conferences — I'll try to avoid that!

MP: What does your research involve?

GH: My research generally involves theory and method development: coming up with new tools for other researchers to use. Obviously, these are the kinds of tools I make use of as a researcher too, so I work to apply what we develop. I've always been interested in optimisation problems, and that's what led us to AI and machine learning.

I'm also interested in intermolecular interactions: very small-scale interactions between individual molecules. It's why water stays in a glass and doesn't float off into the air, or why a table is solid, or why we're alive as humans — it's all to do with these very weak interactions that hold things together. It's fascinating and difficult to treat these accurately. Most of the methods we develop are new ways of trying to describe those interactions in a fast way. In order to describe them, we need to solve the equations of quantum mechanics, because these interactions scale really badly. For example, if you double the number of molecules that you're interested in, it doesn't just double the length of time you need. It would be more like 24, 26, or 28, depending on how accurate you need to be.

Previously these long calculations were the preserve of big national supercomputers, but with AI, computers are getting faster and they're able to handle all this data. It's really opened up possibilities for us.

MP: What problem were you trying to solve in your AI4SD-funded project?

GH: We're looking to use AI to design new materials. This case involves a desalination membrane which would take the salt out of seawater and make it into clean drinking water. Obviously, that's a global sustainability goal, aiming to have safe and secure drinking water. We're trying to create something like a very small sieve, which would let water through, but not salt.

To create this, we need to train the AI to give us the best way of assembling the membrane material that will do that kind of sieving with the perfect sized pore. We do this by working to design a new molecule or material, and then testing its properties. The process is what some term "inverse design": we start with the target property and get the AI to design the

material that fits that brief. The difficult bit is ensuring somebody can actually make the AI's design. The holy grail would be for the AI to come up with something a trained chemist would not have.

MP: How does what you're developing compare with existing desalination products?

GH: Essentially, we're designing a membrane that you could place over some kind of container. Hopefully, it would be cheaper than other materials, but it's difficult to say in these early stages. But this is only one specific application — all kinds of different materials could come out of a similar type of AI pipeline.

MP: How far along are you in developing it?

GH: We had three main goals for the project. The first was to build a dataset of different types of materials, which we can use to train the AI. We've got a big tick next to that, which fulfils a large outcome of the project: looking at ways to generate datasets.

The second goal was the AI design — ensuring the AI can do something with the dataset — which has a partial tick next to it! Currently, we can get it to generate new structures, which we haven't thought of, but we're running into the problem of "Could anyone ever make that?" We've got some ideas, though, as to how to improve things on that front.

The third goal is actually trying to make something that the AI has predicted, so we can test to see if it's got all the properties you would need for desalination. We're ready for a proof of concept on that, though we would have preferred to have the AI working better in terms of prediction. That being said, the AI has narrowed down a class of materials, so our aim now is to take a material from that class, try to make it, and see whether what the AI predicted could be made, and whether it works for desalination. We'd then like to have a kind of feedback loop where we can take the results of the experiment and feed them back into our dataset in a way that also improves the AI. It's what's known as reinforcement learning. We can tell the AI whether it's done a good or bad job, essentially, and use that to inform what it predicts in the future.

MP: How time-consuming is it to make that material? Do time constraints affect how you select the proof of concept?

GH: It depends on how close the chemistry of what's been predicted is to what we already know how to make. If it's relatively close, and within the skillset of our experimental collaborators, then they can probably give us the result within a week's notice. If it's something where we need to invent new chemistry in order to make it, that's a different challenge, but also a good one, because it could be a route towards exciting new chemistry. Although that's not really the goal of the project, it would be a pretty nice kind of 'sub-goal' that comes from it.

MP: What about your research has surprised you?

GH: Some of it is just how much data you can generate so quickly, and then how much time and care needs to go into looking after it all. Because a computer can generate it much quicker than you can inspect it, you have to make sure you can use what it's generating. One thing that's been clear is that a lot of the AI tools available are very specialised. You have to

figure out how to translate the existing tools for a completely different field. We have to adapt or rethink the tools, and it involves stepping back and taking a more philosophical approach. It's a new and fast-moving field and, every now and again, you need to stop and take stock.

MP: What kinds of impacts did the Covid-19 pandemic have on the project?

GH: The disruptions around Covid have thrown our timeline quite dramatically, meaning we haven't got as far as we would've liked. When we submitted the initial proposal for the work, around Christmas 2019, we wanted to begin the project with a hackfest, we imagined we'd get all of our collaborators together for a week to jumpstart the project. It was going to be everybody together in the same room, we'd buy some food and drinks, and we'd all be coding.

In March 2020, we found out that the proposal was successful, and as the pandemic struck we realised "we have no idea how to do this online." We tried to do it, but it didn't really work as well as we hoped. Online is good for getting one person's opinion, but the more collaborative and networking-related aspects of it don't work as well. So this initial boost to the project wasn't really a boost at all. You wanted to try and concentrate and then you felt compelled to keep saying something, because you had a Zoom window open, for example. Some good discussions came out of it, but that way of working was more disruptive than helpful.

MP: Do you think it might work better now, a year on?

GH: Definitely. People are easier in front of the camera now and the technologies have leapt over time. People are happier working remotely, using tools like Slack or Microsoft Teams for that collaborative aspect. With face-to-face, though, if you've got a micro-question like, "How do I do this in my code?" then you can just turn around and ask it. Working remotely, people save up those little questions and it turns into some mammoth Zoom meeting. In that sense, it's a bit of a productivity hit. Yet again, there might be some solutions to these problems. I'd like to figure out how some more tech-oriented businesses have solved them.

A lot of my experimental colleagues have said, "Oh, you've all been all right through the pandemic, you've all been able to keep working throughout it," because obviously at times they didn't have access to the laboratory. But it's had a bigger impact on theoretical and computational chemistry than people anticipated. Those small interactions that are almost social are also research, and they've taken a big hit.

MP: What advice would you have for early career researchers in your field?

GH: I think you need a slightly contradictory mix of resilience and adaptability. Resilient enough to keep going when things aren't entirely as you'd wished, but also adaptable enough so you can change when you realise things aren't working the way you'd initially thought. Finding that balance is important.

Collaboration is also going to be key as we move forward with AI or automated investigations in science. There are moves to introduce robots and continuous flow in chemistry, automating how things are produced, and that's a really exciting area for AI. But it's a different set of skills to what most chemists have at the moment, and you need to bring together different people — robotics experts, pure computational scientists, computational chemists — in order to make that happen. Early career researchers should be looking out for those collaborations pretty soon after they get into their position. How do all those pieces come together, and where do they fit into it?