Artificial Intelligence and Augmented Intelligence for Automated Investigations for Scientific Discovery

AI3SD Interview with Dr Keith Butler
27/01/2021
Online Interview

Michelle Pauli
Michelle Pauli Ltd

04/01/2022

Humans-of-AI3SD:Interview-15

AI3SD Interview with Dr Keith Butler
Humans-of-AI3SD:Interview-15
04/01/2022
DOI: 10.5258/SOTON/AI3SD0178
Published by University of Southampton

Network: Artificial Intelligence and Augmented Intelligence for Automated Investigations for Scientific Discovery

This Network+ is EPSRC Funded under Grant No: EP/S000356/1

Principal Investigator: *Professor Jeremy Frey*
Co-Investigator: *Professor Mahesan Niranjan*
Network+ Coordinator: *Dr Samantha Kanza*

Contents

1 Interview Details

Title	AI3SD Interview with Dr Keith Butler
Interviewer	MP: Michelle Pauli - MichellePauli Ltd
Interviewee	KB: Dr Keith Butler - STFC
Interview Location	Online Interview
Dates	27/01/2021

2 Biography

Figure 1: Dr Keith Butler

Dr Keith Butler: 'Machine learning really has the potential to unlock a lot of the value that's in data that's currently not being realised'

Keith Butler is a Senior Data Scientist working on materials science research in the SciML team at Rutherford Appleton Laboratory. SciML is a team in the Scientific Computing Division working with the large STFC facilities (Diamond, ISIS Neutron and Muon Source and Central Laser Facility, for example) to use machine learning to push the boundaries of fundamental science.

In this Humans of AI3SD interview he discusses the impact of his work, the potential of self-driving labs, the importance of explainable and interpretable machine learning systems and why early career researchers should shout about what they know (and use Linux!).

3 Interview

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

KB: I'm a data scientist at Rutherford Appleton Laboratory. I work in a team called 'scientific machine learning' based in the scientific computing department. I started off as a synthetic organic chemist; my undergraduate degree was in medicinal chemistry. I quickly realised I was an awful laboratory chemist and so I had to find an alternative. I got into computational chemistry at that point and became interested in the theoretical and computing aspects of chemistry.

I did a PhD in computational chemistry, moved into material science rather than organic chemistry and subsequently did a postdoc in computational material science. I became really interested in the work that's going on at the large national facilities. Particularly, I was working with people at ISIS Neutron and Muon Source, based at Rutherford Appleton Laboratory, and then moved to take a job there. There was a scientific machine learning team being set up, not long after I arrived, and got involved at the start of that.

MP: What are you currently working on?

Most of the work I do at the lab is involved with developing and applying machine learning for understanding and interpreting the data that's collected at the lab. I work with the scientists who come to the lab to use machine learning to unlock the potential that's in the data. I also do quite a lot of research with collaborators at universities developing machine learning for designing new materials for given purposes. You might want a new photovoltaic material – can you use machine learning to help you design it?

MP: What difference is this work going to make?

At the lab, people collect so much data on these machines and it takes so long to interpret that data, that a lot of it inevitably gets wasted. These national labs are huge investments, so there's x-ray and neutron synchrotrons at Rutherford Appleton Laboratory. They do fantastic science, but there's so much extra that's probably in there that never gets looked at, just because it's too much work to wade through all the data and understand it. Machine learning really has the potential to unlock a lot of the value that's in that data that's currently just not being realised.

MP: What's holding it back at the moment?

If you see where machine learning has made a huge impact in things like facial recognition or generating texts and predictive texts, the algorithms and the computational power have certainly improved, but the huge thing has been the availability of really good datasets. The ImageNet database kickstarted the deep learning revolution, which is why some of the slightly terrifying facial recognition neural networks are so powerful now – they have these huge labelled datasets.

One of the things that really is a challenge for us is that we have huge amounts of data, but not a corresponding amount of labelled data. So I can't put out Kaptchas of x-ray diffraction patterns and say, "pick out the FCC diffraction pattern from these," because there's not enough expertise out there in the world to label those datasets quickly. So that's one of the big challenges for us in trying to bring the full potential of machine learning to the scientific data at these national facilities.

MP: What about the potential? You've spoken about self-driving labs and hypothesis generation?

Self-driving labs are a really fascinating prospect where, by giving it some kind of input at the start, the machine actually understands the signals it is getting and starts to understand the interesting areas in the signal and drive towards them. At the moment, that usually very much relies on graduate students working at three o'clock in the morning at Diamond Light Source to make those decisions. Obviously, you can't make the right decision all the time. But what if you can have machines that do that and decide the next best experiment and use rational algorithms to decide what's the next best point to measure? What's going to give me the most information? I think that's got huge potential.

Hypothesis generation is also a fascinating area where artificial intelligence can play a role. I personally don't think that you'll replace the scientists in the loop but I think the next generation of labs will have a lot more of this kind of stuff embedded within the laboratory.

MP: What are the pitfalls that we need to be aware of on the path to these innovations?

I would say that blind trust in these kinds of methods is not a good thing. Luckily, I find that there's a healthy amount of scepticism from people in the sciences about machine learning. Most people are interested and really keen to know about it, but they are quite sceptical that it can actually solve their problem. That's really healthy because it means that we don't fall into the trap of training an algorithm that appears to be giving the right answer and then trusting it blindly.

There are interesting examples of adversarial noise that can turn up in samples. That's where a neural network might be shown a picture of a whale and it correctly predicts that that picture is a whale. But by changing the value of just one pixel in the picture, if you choose the right pixel and switch the value around, suddenly it classifies the whale as a giraffe. That makes no sense, right? These systems are incredibly powerful but also incredibly brittle in some ways.

We really need to understand those kinds of frailties and brittles and the processes of how deep neural networks actually work. One of the great things I find from working with scientists is we've got really into trying to build explainable and interpretable machine learning systems. If you can make a prediction, can you say why you made that prediction, what was it in the data that made you say that? What kind of paths do things take through your model and what decisions did the model make along the way in order to reach that decision? Then, does that make sense physically? That's a really fascinating and rewarding area to be working in at the moment.

MP: What has surprised you in your work?

I'm going to come back again to interpretability. I found this surprising on two levels. One is an example of some work that we did recently on neutron data. This is the kind of work that really inspired us to start looking into how to interpret our models. We had a model that made the correct prediction on a neutral spectroscopy dataset. It was predicting what magnetic structure a material had, but it was making the correct prediction. What really surprised me was that it identified the precise area of the input signal and spectrum that enabled it to make the decision. I showed it to physicists who were working on that problem. They had a paper from a couple of years ago where they worked on that material and commented, "that is exactly the same area that we identified as being important."

Since then, I've tried to build interpretability into a lot of the models I have when I work with people. I'm always very pleasantly surprised by how much value people get from that. When they see that the neural network thinks the same thing as them and can explain why, then they trust it. Every time that happens, I'm always surprised by how much people really love to see that kind of thing.

MP: How has Covid-19 affected what you do?

From a day-to-day perspective, we were incredibly lucky in that with the kind of work we do it hasn't impacted us too much. There has been less data collected at the synchrotrons, but there's still plenty of it. Of course, Diamond stayed open the whole time during the pandemic because it was doing a lot of work on structural characterisation of Covid-19, and looking at the binding sites of antigens and things like that.

The way it has affected me is that I don't get to meet up with people at the lab as much as before and have those interactions with people – just the informal bumping into people in the corridor and interactions where projects get started from that kind of thing, or the chatting about something over lunch. I definitely missed the bumping into people.

It's really important to try to recreate those informal meetings and places that spark interaction between people. Because, on a day to day level, it's easy to overlook that kind of thing and not think about it, but if you overlook it and it doesn't happen for a year or two, then you'll probably notice that there's a lot less interesting stuff going on.

MP: What advice would you give to ECRs in your field?

Try to expertise yourself in something. Find something that you think is going to be important and useful, that you can be proud of working on, and find a way to expertise yourself in that. One of the important things there is choosing who you work with. Whether it's thinking about where you go to do a PhD or a postdoc or even just somebody that you collaborate with, try to identify good people to work with. I think the single thing that's made the biggest difference to me has been the people I've worked with or the people I've worked for.

If you do become expert or even quite knowledgeable in something, then let the world know about it. I have found Twitter to be very useful to tell people what you know about, highlight and say, "Oh, this is a really cool article because X, Y, Z." It doesn't have to be your own article, but could be an interesting article you've read. Write things that show people that you know what you're talking about. I've had several fruitful collaborations through Twitter. I've definitely found out about funding opportunities.

For people specifically coming into machine learning, get as good grounding as you can in statistics, because it makes a big difference to your understanding. And try to use Linux wherever possible. If you want to use the latest software in machine Learning and the latest updates and packages and things like that, people will tend to release them on Linux-type platforms first. It'll tend to be easier to set up your Linux machine to do the latest cutting-edge work, as opposed to trying to do that on Windows.