Reflective Report

Matthew Blakeney

The six weeks of my research project were split into a two parts – an initial four-week period, then a final two weeks after a month-long break. I was coding and training neural networks – in the *Python* language, using the *PyTorch* machine learning framework – to find optimal network structures and parameters for finding symmetries of physical systems. Coming up with candidate networks to test involved reviewing literature on the physics of my research topic as well as some trial-and-error to fine tune network parameters. I was fortunate enough to gain remote access to a graphics processing unit (GPU) through Trinity's high performance computing facilities (TCHPC). A GPU is computer hardware that speeds up the machine learning process, so having access to this facility was very helpful for my research.

My project set out to discover a method of using techniques from machine learning to find symmetries in classical mechanical systems. This aim was fulfilled – I created two methods of doing so. The development of a measurement of correctness was a key step in the progression of the project and essentially defines the method of symmetry finding I created in this project. In the end, the symmetry finding algorithms developed in this project were tested on physical systems of varying levels of complexity and were successful in identifying the symmetries of each system.

I think the next step for this research is to build on my approach to finding symmetries in physical systems to eliminate the incorrect transformations that are sometimes identified alongside the correct symmetries. After this, the extension of my algorithm to field theories would be an interesting next step, since they are generally more symmetry rich.

Given that the goal of my project was to develop something without a blueprint from other research mapping out how to do so, there was a distinct possibility at the beginning of the project that I might not find a way to deliver on that goal. I found this motivating – not out of fear that I might not deliver, but out of excitement that I could discover a new way of doing something.

However, the extent to which my project was based on discovery rather than following established experimental methods was also the source of the challenges I faced. At first, my progress was much faster than I had anticipated – in week 1, I had identified my approach to the problem at hand and successfully applied it to find the simplest class of symmetries. By the end of week 3, I had figured out a way to extend my approach to find more complicated symmetries and while this may seem great, I wasn't satisfied with what I had developed. Although I had developed a process to do everything I had outlined at the outset, I felt that the way that I extended my original approach was sidestepping the purpose of the project.

I set out to find a way of applying machine learning techniques to find symmetries, and while my initial approach to finding the simplest class of symmetry did achieve its goal using machine learning techniques, my extension of that approach in weeks 2 and 3 involved some symbolic manipulations. I felt that this was somewhat like using the established, by-hand method to find the symmetries and then having machine learning techniques follow along behind the established method to the solution. Being caught up in the process, I didn't realise this flaw in my extended approach to the problem while developing and coding it. It wasn't until I was testing the approach on physical examples and watching it find the correct symmetries that I saw the futility of what I had developed. While it was tempting to consider the working algorithm a success, my conscience rebuked such a blinkered attitude. Thus, at the end of week 3 I rejected my previous two weeks of work as misguided and returned to the drawing board.

Reflecting on why I had found myself in this predicament, I became aware of the fact that I value the process by which results are obtained over the results themselves. I think was important that I gained this awareness before continuing with the project, because it effectively changed my goal. My new goal became to find an *elegant* method of finding symmetries using <u>only</u> machine learning techniques.

Based on this knowledge, I decided to pursue the development a method of finding symmetries that would emerge from the underlying 'action principle' that Lagrangian mechanics is built upon. The process of investigating such a possibility involved developing my understanding of the (largely imaginary) construction of Lagrangian mechanics from the Principle of Least Action (Hamilton, 1834) as well as my understanding of numerical integration and differentiation of multi-variable functions. I found no promising prospects of methods using these concepts to find symmetries, so after a week of experimenting with this approach – now at the end of week 4 – I dropped the line of enquiry.

The decision to stop investigating this possible approach to finding symmetries was in part due to a conversation I had with my supervisor about when to stop investing time into something that isn't working. Prof. Kirk Soodhalter, my supervisor, had suggested deciding on an amount of time that you're willing to spend working on an issue when you initially find it and approaching the problem from a different perspective at the end of this time. I had initially set myself a week to work on approaching the problem from the perspective of the underlying action principle. However, I think that what pushed me to bite the bullet and abandon this approach was not my discipline, but a newfound optimism about another approach to the problem, ushered in by my happening across a paper on machine learning symbolic integration by a group of Facebook researchers (Lample & Charton, 2019). Reflecting on the importance of the discovery of this group's work to my project, it's clear that my continuous exploration of available literature fuelled my creativity and inspired me to consider alternative approaches to a problem.

By this point I was at the end of week 4 of my research period, at which stage my supervisor and I had planned to break from the project and return to it after a month. However, the area of machine learning used in the Facebook group's research (Lample & Charton, 2019) was unfamiliar to me and I was eager to learn about it. Therefore, outside of the designated project period, I continued my work, learning about the area of natural language processing and transformer neural networks. I also built a dataset of total time and, inspired by the Facebook researchers' work (Lample & Charton, 2019), the plan at this point was to train a transformer neural network to identify the algebraic form of a total time derivative and thus classify expressions as having a given probability of being a total time derivative. This would act as a measure of how close a given transformation is to being a symmetry transformation. But alas, my initial tests were not promising and although expanding the dataset by several factors of ten may have resolved the problem, it would have been a slow process (in part because of computation time) and mightn't have worked.

However, now having a dataset of functions classed as either being or not being total time derivatives, I was inspired to find a way to correctly classify them all. I started considering whether a neural network could learn whether a function is a total time derivative or not based on its graph. I experimented with this idea, creating neural networks to learn from this information, and got a seventy percent accuracy at best. This wasn't as good as I was hoping for, but it did suggest that there was some kind of pattern there to be recognised. I investigated what this pattern could be on paper and found a process of comparing the values of a graph at specific point with each other that could (in theory) discern whether a function is a total time derivative. I coded this process and used it as a measurement of how close a transformation is to being a symmetry transformation in the machine learning process and to my relief the process I had developed worked. In this way, the final approach I landed on was discovered while investigating the feasibility of another, related approach.

At this point I reached the beginning of week 5 of my planned research period and started tying up the loose ends of my research project. Thus, I completed the process of creating an algorithm that could find the symmetries of all the systems that had inspired me to undertake this project. I thought that the challenges were over for me at this point, but as I read through examples of symmetries in other classical mechanical systems, I came across an example that I realised my algorithm would not find. I returned to my workings and realised that my process of identifying total time derivatives by comparing specific points on the function's graph needed extension to allow for finding this class of symmetry.

I was troubled by having come across this symmetry that my algorithm couldn't find because it raised the question of whether there were other symmetries that an extension of my algorithm still could not find. This question set me searching for restrictions on the algebraic form of the all-important total time derivative involved in the identification of symmetries in classical mechanical systems – a question to which I found no answers online or in literature, nor indeed did I find any record of it having been asked.

Thus, week 6 of my research project was dedicated to investigating the hypothesis that, if slightly extended, my algorithm could find any symmetry of a classical mechanical system. This investigation led me to the deepest delve into the mathematics and physics underpinning the theorem I had been working with all along. After playing with the formulae and a discussion of the topic with my advanced classical mechanics lecturer, Prof. Chaolun Wu, I believe I have a convincing reason to accept the above hypothesis. Although I was not expecting to embark on an intellectual obstacle course of this nature in the final week of my research project, I am glad that I didn't ignore the question that had occurred to me in favour of just compiling my results as soon as possible, because the understanding I gained by searching for an answer to my question is invaluable.

I think that one of my key takeaways from my research experience is that my virtues and pitfalls as a researcher both stem from the same personal attributes. I tend to become very invested in the ideas I am researching. This led me to develop an algorithm that sidestepped the point of the project without stopping to think about the validity of the approach as a whole until two weeks' worth of time had been invested into it. However, it is due to this same trait that I persevered with working on the project after several failed attempts over the course of weeks. I can also be ill-disciplined when it comes to cutting myself off from working on an idea that has been proving fruitless. Had I not happened upon a research paper that piqued my interest in a different approach, this may have caused me to continue working on developing an approach to symmetry finding from the action principle perspective. On the other hand, it is because I allowed myself additional time to investigate the possibility of training a neural network to identify total time derivatives that I came across my final approach.

I learned that when intuition guides me in another direction there is no harm in deviating from the original plan, but overall, I could do with stopping and assessing the big-picture impact of the work that I am pouring my time into. Although I've learned this lesson from working alone on a project with very little interpersonal interaction, I think that it is applicable to my leadership approach. The application of this takeaway is obvious when it comes to the project management element of leadership, but I think that it is also applicable to the personnel management element of leadership – if an initial team structure doesn't work coherently, it is a good idea to restructure with the team's long-term goal in mind.

Self-reflection was something that I singled out in my personal development plan as something to improve on. While I did constantly jot down my thoughts and progress throughout each day, it was mostly things that would be of use to refer to later in the project that I was recording. It generally wasn't until I got to a turning point with some part of the project that I would write about my big picture vision of how things will slot together, and it was at these points that I realised flaws in my work when they were present. Thus, I think I will schedule steps back to consider the big picture more frequently while recording the progress of projects in future.

References:

[1] Lample, G., & Charton, F. (2019, December). Deep Learning for Symbolic Mathematics. doi:10.48550/arXiv.1912.01412

[2] Hamilton, W. R. (1834, April). XV. On a general method in dynamics; by which the study of the motions of all free systems of attracting or repelling points is reduced to the search and differentiation of one central relation, or characteristic function. Philosophical Transactions of the Royal Society of London, 124, 247–308. doi:10.1098/rstl.1834.0017