A UROP perspective by Abigail Levison
Summer 2024 (undertaken in the Department of Physics)

Abigail was a year 3 undergraduate at Imperial College London in 2023-24: MSci Physics with a Year Abroad, Department of Physics

UROP title: Neural Wave Functions for Positronic Chemistry

At the start of the 2023-2024 academic year, I started looking for internships and research projects to explore different research topics and gain insight into potential PhD areas. After I had heard about using neural networks to represent the wave functions of molecules with more than two electrons the previous year, I reached out to Prof. Matthew Foulkes. I was eager to work in his area, which combines quantum physics and machine learning, two fields I'm particularly interested in. I was fortunate to receive the EPSRC Vacation Internship Award from Imperial, which enabled me to pursue this project. My UROP focused on using FermiNet and its new transformer variant, PsiFormer, to represent the wave function of a formaldehyde molecule (with and without a positron) and use this to calculate the positron binding energy in the molecule. In preparation, I read some papers, including one by a former PhD student in the group, whose binding energy calculation results for small molecules I would reproduce, and then add to with the formaldehyde molecule.

As a foundational exercise to understand how FermiNet and PsiFormer represent wave functions with the aim of minimising ground state energy, I implemented a Variational Monte Carlo (VMC) algorithm to optimize the parameters of an analytical wave function ansatz. I did this on Python, using JAX, a library which I had no prior experience with, on hydrogen and helium ansätze. I'm very grateful for all the support I received, especially from Andres, a PhD student in the group, who shared some JAX examples and tips about the (many!) things to watch out for. Implementing these algorithms was extremely useful, since JAX is heavily used in both FermiNet and PsiFormer. It also helped me grasp the principles of wave function optimisation in these models. I found that the ground state energy I obtained for helium was better converged than the result reported in the reference paper!

At the start of my UROP, I had some prior experience with submitting jobs on a cluster, but the one used for FermiNet simulations was very different. Andres, very helpful as always, guided me through the process of setting up Python configuration files and bash scripts for submitting jobs, and introduced me to useful bash shortcuts he invented, which inspired me to create some of my own. He also helped me debug my code and shared some very useful custom features he had developed for his own FermiNet experiments.

The formaldehyde molecule is challenging due to its small binding energy (on the order of one mH) in published calculations. Initially, my results showed that the ground state energy of formaldehyde with a positron was unexpectedly higher than that of the bare molecule, contradicting published calculations and suggesting it would be energetically favourable to release the positron to infinity. Motivated to resolve this, I worked to train the neural network further and compare results with published studies. During training with a positron, the energy values occasionally spiked. With the help of Andres and Prof. Foulkes, I investigated these spikes, and we concluded they stemmed from an electron and positron coming very close together. We considered whether this was due to the neural network architecture or initial configuration settings. With Prof. Foulkes' input, we suspected that the envelope function in the PsiFormer wave function representation might be too restrictive for positrons. Andres and I adapted the envelope function to adjust the width for positrons, with the added challenge that we couldn't add new input variables since the function is called a lot in the main FermiNet code. This was a very valuable experience in adapting code written by others.

As of one week before the end of the UROP, the total formaldehyde energy has not yet converged, and we are experimenting with adjustments to the initial configuration file to improve this. Calculating the binding energy of a positron in formaldehyde proved more challenging than anticipated, but I thoroughly enjoyed the UROP experience. It deepened my interest in quantum physics and clarified that I should pursue this field for a PhD. I also enjoyed the opportunity to network with PhD students and ask them about the PhD application process, gaining some important tips such as the importance of contacting and meeting potential supervisors beforehand. Observing my office mates puzzle over complex integrals, differential equations, and coding issues, and seeing Andres' enthusiasm and dedication to his FermiNet project, has inspired me to apply for a PhD.