Laidlaw Leadership Blog

Introduction

This summer I had the pleasure of spending 6 weeks working with Neontribe, a UK based digital agency working with charities and the third sector.

Having read about Neontribe's mission of 'tech for good' - including projects such as their work with the Alexandra Rose Charity distributing food vouchers – I knew it would be a good fit for the project.

The Martingale Foundation was selected as the project client. They are a national charity which offer postgraduate scholarships in STEM subjects to students from low socio-economic backgrounds.

The value of their work extends beyond equality; with the rise of AI and the continued digitalisation of society, people in STEM are often making decisions with far-reaching consequences. If those people are a fairly homogenous group from similar backgrounds, with similar experiences, then these decisions are not exposed to a truly representative set of perspectives. Ultimately this leads to systems which unfairly affect certain demographics of people. By broadening access to STEM at the postgraduate level, the Martingale Foundation are contributing to creating a fairer society for everyone.

I first met with the team from the Martingale Foundation in June, a few weeks before my project started. They are a very lovely group of just 7 people. Our key contact was the brilliant Mary Henes, COO, and she walked us through their workflow and the problems that they were facing.

As a small, but rapidly scaling, charity (just a few years old!), some of their processes are still evolving. With such a small team, its important to streamline these processes where possible. Ultimately this will allow them to scale their operations, increasing their impact without extra cost. That's where we came in.

Project Methodology

I was working with Senior Engineer, Charles Strange and, less directly, with Neontribe Head, Harry Harrold. I am grateful that I was given a huge amount of autonomy in this project, steering the project with the guidance of them both.

Harry is a big fan of the Design Council's Double Diamond methodology, which can be thought of in four stages: discover, define, develop, deliver. This is how we structured our project.

Discover

We spent the first two weeks understanding the Martingale team's problems, and exploring what a potential solution could involve. Time constraints were tight, and something small but useful is better than something ambitious but incomplete. In conversations with the team, we identified two areas which we thought we could help with.

1) Scheduling interviews. The Martingale team hold several in-person interview days across the country. Candidates meeting their socio-economic and academic criteria have to be interviewed by two academics, as well as by the Martingale team themselves.
   After only a few years in operation, they already interview 140+ candidates, presenting an enormous scheduling task to match candidates and academics to the appropriate slots. Additionally, the Martingale foundation pay for candidates expenses, so it is important to match candidates to their nearest interview assessment centre.
   Potential impact: save a week of person-hours each year, as well as hundreds of pounds in candidates expenses.
2) Validating household income. Scholarships are only offered to candidates meeting certain socio-economic criteria. One factor considered is household income, for which student finance assessments are used as a proxy. Candidates submit these documents for Martingale team to validate. Currently the task of checking that the values in these documents match or imply those stated by the candidates is done manually. This costs hours of time each week over the application period and beyond (lasting several months), with most applications received close to the deadline.
   Only after a candidate has been screened for socio-economic factors, can they be assessed academically.
   Potential impact: approx. a week of person-hours each year, as well as reducing bottle-necks in their work-flow.

The use of AI was suggested for both tasks. However, I believed that a traditional algorithm would be the best approach for the scheduling task.

AI carries with it a whole host of further considerations and ethical implications, which I explored during the discovery period. These include data protection & privacy, hallucination, security concerns, bias and energy consumption.

With all of this in mind, a balanced approach is required. LLMs can be incredibly useful, and in the context of a small charity, can scale impact significantly. However, they must be used carefully.

We deemed the second task to be an appropriate ethical use of AI, with systems in place to ensure that it was ultimately always a human having the final decision – reducing the risk of hallucination or bias interfering. The project is structured such that the LLM does not have access to the Martingale team's CRM software. I investigated the privacy and data protection approaches taken by various companies to select a model which would be consistent with the data protection requirements of our clients.

Define

We discussed our proposed approaches with our clients, including the **risks presented by using LLMs.** We agreed upon the following:

1) A locally-run scheduling tool using traditional algorithms. Running locally means they are not handing data over to third parties, reducing data protection concerns.
2) An LLM based automation to check candidates' student finance documents.

We also managed expectations: reliable and maintainable software is not usually produced in 6 weeks, let alone two separate projects. This was an important part of maintaining the client relationship and **demonstrating honesty and integrity**.

Develop

Having decided what we wanted to do, we needed to determine how. My focus during development was on the scheduling task, while Charles worked on the LLM automation. I experimented with a few different algorithms, settling on using Google's OR-Tools to do some constraint programming. I tested the program to ensure it behaved as expected and also worked to optimise it.

I also spent some time researching prompt engineering and constructed a prompt for Charles to include in the LLM automation task. I sought feedback from several different LLM models themselves to evaluate the prompt, building it up to consider all of the nuances of the task, and how to handle certain "edge cases"; this acts to reduce the likelihood of unexpected (and unwanted) behaviour from the LLM, by filling in any gaps in its instructions. Proper prompt engineering is a vital part of **ethical LLM use**.

The LLM outputs to a JSON file, a type of structured output. This restricts the possibility of unwanted behaviour and acts as a form of validation/input sanitisation.

Deliver

We demoed the scheduling tool and the Martingale team were thrilled – they could see how it would save them so much time. I packaged the scheduling tool up to an executable (.exe file) and shared with the client, alongside instructions for use (I intend to follow up with more complete documentation soon).

Given other constraints on Charles's time and my own inexperience with API's, the second project is yet to be completed.

Final reflections

The way the Neontribe team interacted with the Martingale clients taught me a great deal about leadership. Namely that *listening* is vital to delivering impact and really helping people, allowing you to identify their needs (even if they are unsure which needs are feasible to address). Furthermore, *asking the hard questions up front* has shown to be essential.

Finally, the experience has highlighted to me how impactful *curiosity* is in a leadership context. Being curious about everyone's point of view, being curious about the risks, being curious about potential. Curiosity is a powerful mindset that can help keep your point of view balanced, while exploring every avenue.