Quiet Quadcopter Propellers and Financial Market Prediction using Deep Recurrent Neural Networks

Samuel Holt
Rokos-Award funded Internship
Pembroke College
University of Oxford

***Abstract*—Spent 3 months setting up experimental apparatus, improving, modelling and understanding how to reduce noise emitted from quadcopter propellers. Research culminated in identifying design selection criteria and an experimental propeller design that showed a noise reduction of 6 dB SPL lower than the best in class low noise propeller geometry. Spent an additional 3 weeks predicting financial markets using deep recurrent neural networks, specifically predicting Forex currency pairs. Experiments showed deep neural networks show potential for prediction, however are outperformed by more conventional methods such as a modified hidden Markov model with an accuracy of prediction of 74% within the next 30 minutes, to the price direction up, down or staying constant. Learning a variety of machine learning methods, through on line courses and working with fellow academics in Oxford's Machine Learning Research Group. Additionally over the summer gave me time to undertake and complete a long distance on line course in Machine Learning, by Andrew Ng at Stanford University, which I have completed with an score of 98.8%.**

I. Quiet Quadcopter Propellers

Quadcopters are experiencing significant growth in adoption and use in ever increasing applications, however the technology that they rely on has not been optimized or researched to reduce their largest inconvenience factor, that of the loud noise (100 dB sound pressure level (SPL)) that they emit, which significantly hinders their close interaction with people. Aim of the project was to use all prior research and apply this to 12" quadcopter propellers to attempt to reduce the noise as much as possible. The problem was formulated to research a propeller that can provide a nominally constant thrust to counteract the weight of an average commercial quadcopter, approximately 400 grams of thrust for a 12" propeller.

Research last examined small propeller noise directly in the 1970's, amid oil crises when propeller planes were set to overtake jet engine planes if the oil crises continued. Work by David Brown et al. (1971) formulated empirical formulae to predict the noise generated by the propeller. Applying his equations to this specific problem, it is possible to reduce the noise in theory by decreasing the tip speed velocity (by reducing RPM), increasing the blade diameter, increasing the blade solidity (blade number), and area of the blade, in effect to reduce the blade loading density for a fixed thrust.

However realistically quadcopter propellers are designed in standard sizes, often 12", and increasing the number of blades leads to significantly less aerodynamic efficient thrust during hover, verified by my experiments. Applying this understanding to well understood propeller theory, non-dimensional analysis, specifically the Thrust coefficient for a propeller, which indicates that for a given propeller to reduce the noise we want to maximise its thrust coefficient (thrust per fixed size of the propeller).

Noise reduction of aerofoils has seen a recent research interest, and has proven reliable noise reductions at large scales of fixed wing aerofoils for planes. This is achieved through the use of trailing edge serrations, which reduce the noise emitted without a too significant reduction in lift, approximately 7dB as measured by Alexandros Vathylakis et al. (2015).

Trailing edge serrations are believed to be one of the dominant noise reduction technologies found in an owl in nature, as they are believed to work by reducing vortex shedding trailing edge noise, by redistributing fluid momentum and turbulent energy towards the saw tooth tip.

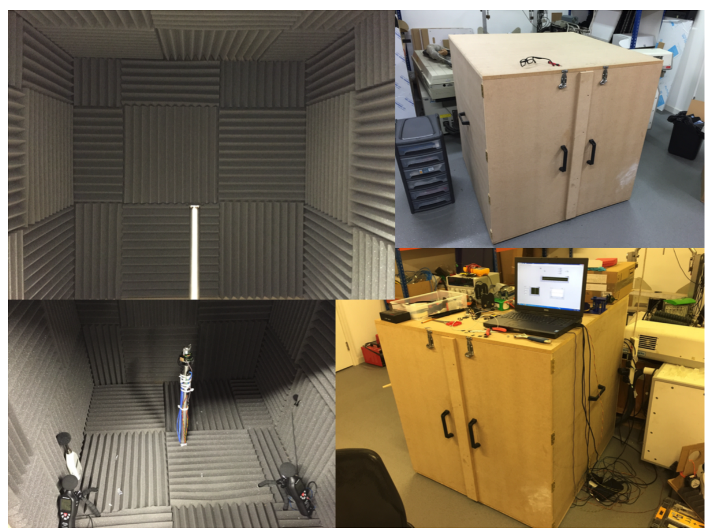

Fig. 1. Experimental Apparatus

A. Methodologies and Experiments

I designed and built a large anechoic chamber ($1.5m^3$), constructed from 8mm thick MDF sheets, lined with sound absorbing foam tiles, and designed a steel platform holder pole, which was outsourced to internal Engineering Science Metal workshop to weld together. I designed and put together a high speed (8,000 RPM) brush less quadcopter motor setup, with robust magnetic tachometer, robust load cell and static synchronized calibrated SPL Cirrus noise meter, 2 SPL level meters, and an internal synchronized microphone. I designed automated testing, and synchronization of data from multiple sensors, using various Labview Data Acquisition Devices, which fed their data into a computer for logging, control and analysis in real time. All propellers used in experiments bought propellers and the 3D printed propellers were all balanced on a mechanical propeller balancer, by sanding excess material from the underside of the aerofoil to remove excess material.

To design custom propellers, I created code that took propeller measurement data (beta angles along the chord of the propeller, and chord to length ratios, pictures of the top down outline and side pictures) from a the large online Illinois model propeller database, to construct 3D CAD models of selected propellers. Once I have the 3D model in point cloud format, I created various algorithms that can adapt that 3D propeller to add trailing edge serrations, warp the geometry and change the dimensions and or the aerofoil NACA shape. Once the desired propeller was created, I can then use the labs high resolution (Formlabs 2) resin 3D printer to print the propeller. Once the 3D models had cured, I was able to test them in the designed and built experimental apparatus, collect data, and analyse any noise reductions for the same amount of lift. This process was repeated iteratively testing various designs, theories and ideas, to investigate what would reduce propeller noise.

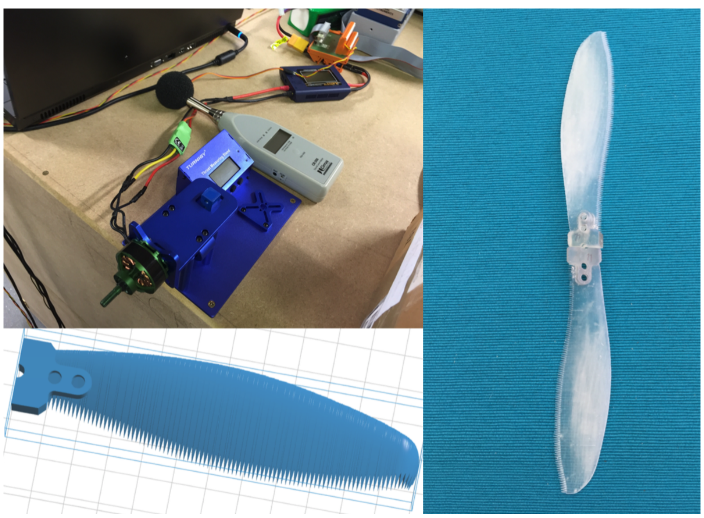

Fig. 2. Thrust stand, and a 3D printed propeller

Evident from theory that the lowest noise propellers were those with the largest thrust coefficient, this was experimentally verified by testing 11 different off the shelf propellers. The propeller geometry which showed the lowest noise for the same amount of lift was the ”Slow Fly” geometry. Therefore I searched and used the propeller geometry with the highest thrust coefficient from the online propeller database (270 propellers), which had a higher predicted thrust coefficient 1.4 compared to the slow fly model 1.1. However when I 3D printed this model, and tested it, it underperformed against the off the shelf ”Slow Fly” geometry, however this could be due to the 3D printed material is different from common propeller injection moulded plastic, and additionally the surface roughness of the 3D printed propeller would have been different compared to the injection moulded plastic.

The technical problem of how to further reduce the noise for a fixed amount of thrust, incorporating trailing edge serrations on the propellers was then worked on. Initially experimenting with trailing edge serrations along the entire edge of the propeller, showed reduction in noise of 6dB for the same geometry, however a significant loss in thrust, due to decreased blade lift area. Also experimented with q-tip, warping the propeller around on itself, however the results showed this reduced lift more than noise reductions.

Through many iterations (approx 40) a optimized propeller was empirically researched and designed that shows lowest noise for its geometry, and trailing edge adaptions, whilst maintaining the same thrust of 400 grams. This lead to a reduction comparative noise of 6 dB SPL. This design was verified in various materials, and appears on small tests to be more efficient than conventional comparative blade geometries. The design will not be published in this paper, as I am still under discussion about the possibility securing Intellectual Property on the design, with a patent from the University.

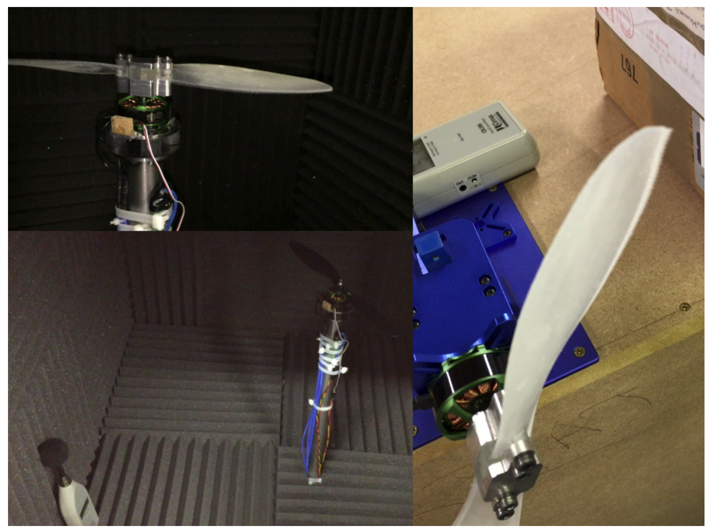

Fig. 3. Propellers being tested

Working on this project, with supervision from Post Doc James Turner, gave me great insight into all phases of an experimental research project. I learnt many lessons, among one of which is best to always plan out exactly what you plan to achieve, and there are often many ways to design experimental apparatus, as I discovered when testing propellers originally I designed an optical tachometer, to measure the propeller speed, however was larger than expected and collided with a propeller, upon the propeller undergoing larger deflections than anticipated when operating. This was replaced with a much smaller, more accurate magnetic tachometer. Likewise the same was for the 3D propeller manipulation, I coded all my algorithms in Matlab, however I needed to blend two CAD components a blade holder and the custom blade itself, originally I wrote code to write out Solidworks macro commands to create the blade, however soon ran into Solidworks macro size constraints, with high resolution point cloud models, (2.5 million points), therefore I switched to using CAD package called MeshLabs to blend the two CAD models together, which I could then create the final 3D model, which was sent to the 3D printer for printing. Also in some blade geometries, I drafted sharp angles, which lead to sever stress concentrations, which would often break the propeller under load at those areas. This project helped me gain real insight following a project from conception, design and build, testing, experiments to write up. I think these skills will be invaluable in my future research when carrying out my dPhil (PhD).

II. Financial Market Prediction using Deep Recurrent Neural Networks

Forecasting financial markets is of great interest to hedge funds, and financial organisations, for opportunities and risk mitigation. Adapting the work on time series forecasting by a dPhil candidate, and working with Stephen Roberts, to predict the price trend of financial prices, specifically tested on easily obtainable high resolution Forex currency pair GBP/USD minute data. Previous work forecasting stationary time series signals showed that deep recurrent neural network (RNN), Long Short Term Memory (LSTM) architectures were always outperformed by conventional proven forecasting methods, such as Attractor based Gaussian Processes.

However LSTM architectures showed best in class for predicting repetitive patterns, such as seasonal variations, over conventional methods. My small internship in the Machine Learning Research Group in the University of Oxford, was to apply this work to predict non-stationary financial data, to improve the results of the deep RNN LSTM, it was suggested that I transform the continuous minute data into symbols, where each symbol corresponds to a price movement direction up, down or flat. Coding up a symbol extractor that sorted all price points into hourly blocks, and used 50% of the data points randomly permuted to perform linear regression, and then used the remaining 50% data points to act as a cross validation cost, using the mean squared error cost criterion. This was done to attempt to avoid over fitting. Created a Markov like model adaption to act as my benchmark, to compare against the RNN LSTM.

Initially I tried a symbol character set of 5, which were up, down, flat, quadratic up and quadratic down. I extracted symbols for 4 weeks of data (This was as much as I reasonably could extract using my limited computer processing power, within the given time frame), 95% of this data was training data, and the last 5% was testing data. Running the RNN LSTM output a best accuracy of 27.27%, however this was outperformed by the benchmark of an accuracy of 30.57%. Interestingly both of these methods outperform random, which has a default accuracy of approximately 20% with 5 characters to choose from.

Analysing the data, it was clear I could increase the accuracy by reducing the symbol character set. I reduced the symbol character set to 3 (Up, Down, Flat), and overlapped the symbols by 50% (i.e. the symbols are one hour wide, however increment 30 minutes in time). I ran the same RNN LSTM which output a best accuracy of 45.45%. However this underperformed compared to the benchmark model with an accuracy of 74.14%. Both of these methods outperform random, which has a default accuracy of approximately 33% with 3 characters to choose from.

Working in Stephen Roberts lab gave me great insight to how world class machine learning and quantitative financial research is performed. This has improved my machine learning skills, research skills, report writing, team work and data science skills. I am currently in talks with an application for a dPhil within the machine learning research labs here in the University of Oxford.