AGA-Goethe Fellowship Report 2019 – Neuroscience-inspired Navigation for Robots

Dominic Dall'Osto

21 December, 2020

Acknowledgements

I would like to thank the Australian German Association and Goethe Institute for all the amazing opportunities I received from the 2019 AGA-Goethe Fellowship. It was a unique and extremely valuable experience for me to meet with top researchers in the neuroscience and robotics engineering fields, and to be able to rapidly advance my German to a higher level.

Thanks must also go to the teachers, staff and volunteers at the Goethe Institutes in Munich and Berlin, both for the language lessons and cultural programs offered. My time at both institutes was incredibly rewarding from start to finish.

Thank you also to Professor Sen Cheng at the Ruhr University of Bochum, for allowing me to complete a visiting research stay as part of his group – the work being done there was inspiring, as was Sen's vision and coordination of the group as a whole.

I must also thank Professor Florian Röhrbein for introducing me to a robotics startup company in Munich. They made me feel incredibly welcome and I am proud to have contributed to their exciting project while there.

Thanks finally to all the colleagues and friends I made over the trip, without whom my time in Germany would not have been the same.

The current world situation has made me even more appreciative of my experiences thanks to the AGA-Goethe Fellowship, and how fortunate the timing for me was. I wish all the best to the AGA and Goethe Institute through these difficult times, and to the future winners of the Fellowship, who I am sure will have as valuable and rewarding an experience as I did.

Dominic Dall'Osto

Contents

About the Author

My passion is combining the fields of neuroscience and engineering to benefit both. I graduated from my undergraduate mechatronics engineering degree from Queensland University of Technology in mid-2019, only 1 day after being awarded the AGA-Goethe Fellowship which was very exciting. My connection to Germany comes from my family ties there, as two of my grandparents were born in Germany before moving to Australia, and I still have some relatives living there. I learned German for 6 years in total during school, and then went on an exchange semester during my university studies to the Technical University of Munich, a university renowned for its engineering excellence. There I was free to explore different subjects, and decided to focus on computational neuroscience – how we can apply science and engineering techniques to better study the brain, and at the same time be inspired by the brain to learn new engineering ideas.

Upon returning to Australia, I completed my bachelor's thesis – taking inspiration from the brains of rats to develop new techniques for robotic navigation. Even small, simple animals like ants and rats can navigate the world successfully – exploring, foraging for food, then returning to their nests. In comparison, the field of robotics still faces many challenges before we can safely have robots moving around our houses doing work for us. I presented a paper on my thesis at the Australiasian Conference on Robotics and Automation in 2018 which confirmed my desire to pursue further studies after my bachelor's degree.

The opportunities afforded to me by the AGA-Goethe Fellowship were very valuable in helping me shape my future career. I have recently begun a masters degree at the University of Zürich in my dream field – combining neuroscience and engineering – and I continually feel like my time in Germany prepared me well for this.

Apart from my studies, I have completed short internships at a number of different companies including CSIRO; Fraunhofer IPA in Stuttgart; the Defence, Science and Technology Group in Sydney; and Emesent. I have also worked part time during my studies at Softbank Robotics in Brisbane, investigating the use of social robots in education, commercial and Government applications. In my free time I enjoy playing volleyball, cooking, travelling, and photography.

Figure 1: Being awarded the 2019 AGA-Goethe Fellowship on what was a very exciting evening for me.

Overview of time in Germany

My time over in Germany consisted of four parts. I started in Berlin, attending two leading computational neuroscience conferences. I then completed my 8 weeks of intensive German language courses, 4 weeks in Munich followed by 4 weeks in Berlin. After these courses I passed my B2 level German examination, which was a great achievement for me. Finally, I completed a 2-week research stay at the Ruhr University of Bochum, collaborating with their researchers. Each part of my trip is detailed more extensively later in this report, but here I provide a brief summary of the highlights.

Conferences in Berlin

I flew into Berlin on 11 September, 2019 to begin my AGA-Goethe Fellowship study and research tour. Fortunately for me, two major computational neuroscience conferences had been coordinated to take place back to back at the Technical University of Berlin – the Computational Cognitive Neuroscience Conference and the Bernstein Conference. I saw countless fascinating presentations over the 8 days of the conferences, and had many discussions with researchers from around the world. I also met Professors Sen Cheng and Florian Röhrbein, whom I would spend time with later in my trip. A particular highlight was the student symposium on the final two days of the Bernstein conference, where we discussed questions not only about our research work and aspirations, but also how we could make the most of our opportunities as young researchers. This was especially valuable for me.

Goethe Institute Munich

I then moved to Munich for the first month of my intensive German courses at the Goethe Institute. I found these courses extremely valuable, especially for my ability to articulate myself well and discuss complex topics in German. In my class alone were students from 9 countries, all choosing to learn German for different reasons, which was very inspiring. Another highlight of the Goethe Institute in Munich was the guesthouse for students of the Institute. Here I lived and studied with students from every continent (except Antarctica) – an experience which contributed greatly to my cultural immersion and language learning. In addition to my language studies, I also spent two days visiting a friend at the BMW Research Centre. It was exciting to see some of the research being carried out there, and I was lucky enough to go for a ride in one of their cars as part of a data collection trip. I also reconnected with Florian Röhrbein and worked for a startup of his, using experience I had gained from CSIRO in Australia to help them setup a robotic simulation.

Goethe Institute Berlin

Next, I moved to Berlin for the second month of intensive German learning. I found living in Berlin an amazing experience. It is a city with so much history, which I could appreciate significantly more thanks to the guided city tours provided by the Institute. Coincidentally, it also happened to be the 30th anniversary of the fall of the Berlin Wall on 9 November 1989. The city therefore had an array of expositions and cultural events to celebrate. Another standout event was Berlin Science Week. A highlight was seeing cutting edge robotics research presented next to a dinosaur skeleton in Berlin's traditional

Natural History Museum. Finally, at the end of my stay in Berlin I completed the B2 level German examination – a proud achievement for me and testament to how quickly my German had improved in only 2 months of courses at the Goethe Institute.

Uni Bochum

For the last part of my trip, I spent two weeks at the Institute for Neuroinformatics at the Ruhr University in Bochum. Here I saw many familiar faces from the conferences in Berlin who were all exceptionally welcoming. The range of research work at the INI was very impressive, from brain-inspired techniques for controlling robot arms, to computational models of how our brain is able to learn. The lab's Christmas celebration was held on my final day which made for a nice conclusion to my stay in Bochum, and to my time in Germany.

Research Background

Firstly, I will provide a higher level overview of the motivation for this research – why it is important and what the potential outcomes are. I will then review some relevant related work in the field. Finally, I will discuss my research results from this work.

Motivation

The robotics industry is currently undergoing a significant transition from traditional manufacturing robots to mobile robots which are able to interact with and work alongside people in their everyday lives. Predictions are that the global market for mobile robotics will be worth nearly $US15 billion by 2023 [1], and the share of the industry in logistic centres alone will be worth $US3 billion by 2022 [2].

One of the biggest limitations for the growth of mobile robots is that the task of navigating a robot through all but the simplest of environments is still very complicated. There is no one solution, and significant amounts of computing power are required. While some success has been experienced with the continued development of autonomous vehicles [3] – they typically have access to:

1. a predefined map of the road network in which they are operating; and
2. a GPS signal informing them of their current position and direction.

For a mobile robot operating indoors no GPS signal is available. The robot must be able to determine where it is in the environment from sensor information and past knowledge. To make this more difficult, typically no detailed map exists of an indoor environment. Indoor environments used by humans are also always changing as we open doors, move chairs, and cover windows. To us, these changes seem relatively minor and don't affect our ability to recognise where in an environment we may be. But for a robot, a chair being moved to a different location would cause a difference between that robot's map and what it perceives. In the worst case, the robot may not see the chair and run straight into it. Alternatively, it may mistake the chair for a wall, then think itself to be in a completely different location. While a robot could be programmed to identify a chair visually and respond appropriately, this is just one of the almost infinite number of objects present in human environments.

However, biology presents many robust solutions to the problem of navigation. Much experimental work has proven that even a rat can perform complex navigation tasks, learning and remembering new paths through the environment much better than most advanced robotic systems. This is despite the rat having a brain with only 200 million neurons [4], 400 times less than the 86 billion in the human brain [5].

Computational models have been developed by studying how rats navigate the world, and these have been applied to robots by the RatSLAM system [6]. SLAM is Simultaneous Localisation and Mapping, when a robot can explore a new environment while keeping track of its position (localising) and creating a map of the environment (mapping) at the same time, just as humans and animals do. This approach currently performs at state of the art levels.

Like the brain, RatSLAM uses a network of a large number of cells which represent the location the robot is currently at. Each cell is connected to many other cells with a certain strength, and each connection must be "correct" for the system to function as desired. Currently, these connections are mainly tuned by hand, but this tuning is typically very specific to the operating environment of the robotic system. Changing a network calibrated for a slow-moving tour guide robot in a museum for use by a fast-moving robot carrying objects around a warehouse would require almost a full re-tuning of parameters. There is therefore significant potential for a process to automatically calibrate such a system for a particular robot. It would have the ability to significantly speed up and simplify the development of navigation solutions for mobile robots, and to improve their overall performance.

Mobile Robot Navigation

Navigation for mobile robots is the process by which they can move from their current location towards a set goal. There are two classes of approaches to this problem: *reactive navigation* and *map-based navigation*. Reactive navigation is the simpler of the two and involves a robot making decisions based only on what its sensors can detect at any particular time. This is in contrast to map-based navigation, where a robot uses a map of the environment to plan a path to its goal. Reactive navigation is only suitable for very simple mobile robotic tasks such as vacuuming a room by moving around randomly and avoiding obstacles. More complex behaviour typically requires a map-based approach [7].

Creating a Map – Simultaneous Localisation and Mapping

In most real-world applications no prior map is available and the environment is continually changing. Therefore, to navigate in such environments, a mobile robot must generate its own map of the environment, and update it as changes are detected, all while maintaining an estimate of its own position in this map. This process is called Simultaneous Localisation and Mapping (SLAM). It relies on the assumption that the robot has some sense of its self-motion and is able to recognise particular locations in the environment as it revisits them. If so, the robot's map of the world can be improved over time. Simultaneously, the robot's estimate of its own position within this map increases in confidence [8, 9]

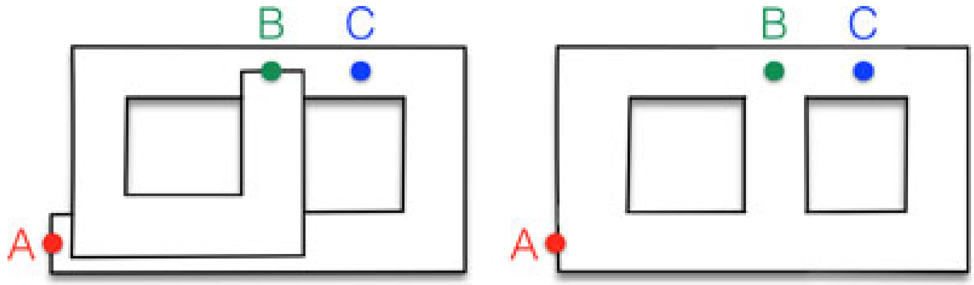

Figure 2: The SLAM problem. On the left a map has been constructed from robot odometry (self-motion), so the map is essentially a corridor from A to C to B. While B and C are close in real space, they are separated in the odometric map. On the right, SLAM has been used to recognise that B and C are close to each other, so when returning to point B the robot realises it is also next to point C. This process is called "closing the loop". Reproduced from [10].

A SLAM system involves two parts: the back end and front end. The front end analyses the sensor readings and translates them into meaningful data for the model – determining the likely location of the robot by recognising a location in an image. The back end fuses the data from multiple sensors over time to keep track of the state of the robot and the world [10]. Biologically-inspired methods for robot navigation are similar to a SLAM back end, so some traditional approaches to SLAM will first be reviewed to contextualise this research.

Traditional SLAM approaches

A traditional SLAM back end formulates the problem as a Bayesian probability distribution of the locations of the robot and landmarks in the environment. If the SLAM process is successful, these probability estimates converge to the true positions [9].

The most common approach assumes Gaussian probability models, using the Extended Kalman Filter to achieve SLAM (EKF-SLAM). It is a well-developed technique but suffers from several limitations. Landmarks detected in the environment must be correctly recognised and associated with previously detected landmarks for the solution to converge. This is especially difficult in large environments where many locations appear similar. Additionally, the computational cost of this approach increases significantly with the number of landmarks in the environment. Finally, the process assumes linear models of the robot's motion and measurements, which is never the case in practice. Linearised models for EKF-SLAM inherently causes errors in results which may or may not be acceptable [8].

A common alternative to EKF-SLAM is the Rao-Blackwellised particle filter, which uses random sampling to express more general non-Gaussian probability distributions of the robot's and landmarks' positions. While not as statistically robust as EKF-SLAM, the particle filter approach has proven effective in practice [8].

More Recent SLAM Developments

More state of the art methods exist which improve the accuracy, computational efficiency, or environmental scale over which SLAM can successfully operate. These include improved back end techniques such as factor graph optimisation, and novel front end means of feature extraction from images such as deep convolutional neural networks [10]. Additionally, some approaches produce only topological maps, instead of to-scale metric maps. This is conceptually similar to the biologically inspired SLAM techniques which will be discussed below [11].

For certain applications, current SLAM approaches achieve satisfactory performance. This typically involves robots moving at slow speeds in well defined, relatively small environments with accurate sensors. But for situations where not all 3 of these conditions are met, research work is ongoing [10].

Nature presents many solutions to the problem of navigation. Animals with relatively small brains are still able to successfully live in complex environments. Experimental work has shown that the mammalian brain does not store information in a rigid metric map like traditional SLAM approaches. This has led to the development of a number of biologically-inspired techniques for SLAM, some of which will be reviewed below [11].

Biological Navigation

The Biological Basis for the Brain's Map

Research into the neurological basis for navigation in animals can be traced back to the work of Tolman in 1948 [12]. While analysing a number of experiments on rats in mazes, he noticed that it didn't seem their behaviour could be explained by simple connections in the brain that linked stimuli to actions. Rats displayed the ability to learn from time spent in a maze without a reward, so that when a reward was introduced, they reached it more quickly and consistently than a control group. In another experiment, rats were trained to follow a certain path through a maze to a reward. When the structure of the maze was dramatically altered, they tended to follow the path leading most closely to the physical location of the reward in the initial maze. In the most amusing case, after learning a maze, some rats escaped from the starting box and ran along the top of the walls directly to the reward.

All of these observations led Tolman to propose the existence of a cognitive map – some structure in the brain capable of learning spatial information at a higher level than simple route memorisation [12]. At the time, the technology was not available to measure the activity in a rat's brain, but the idea of a cognitive map became influential in many different fields [11].

In 1958, a technique was developed to allow the recording of single neurons in unrestrained animals [13]. This led to the discovery of *place cells* in rats by O'Keefe and Dostrovsky in 1971. These were neurons only active when the rat "was situated in a particular part of the testing platform facing in a particular direction" [14]. Further research into the phenomenon found that the same place cells could be activated by a different place in a different environment [11].

Later research by Ranck Jr., and by Taube *et al.* [15] found cells in the postsubiculum of the rat brain which responded based on the direction the rat was facing in a cylindrical chamber – called *head direction cells*. Analysis of this neuronal activity allowed the visualisation of a "firing curve" for each head direction cell. Peak activity would occur

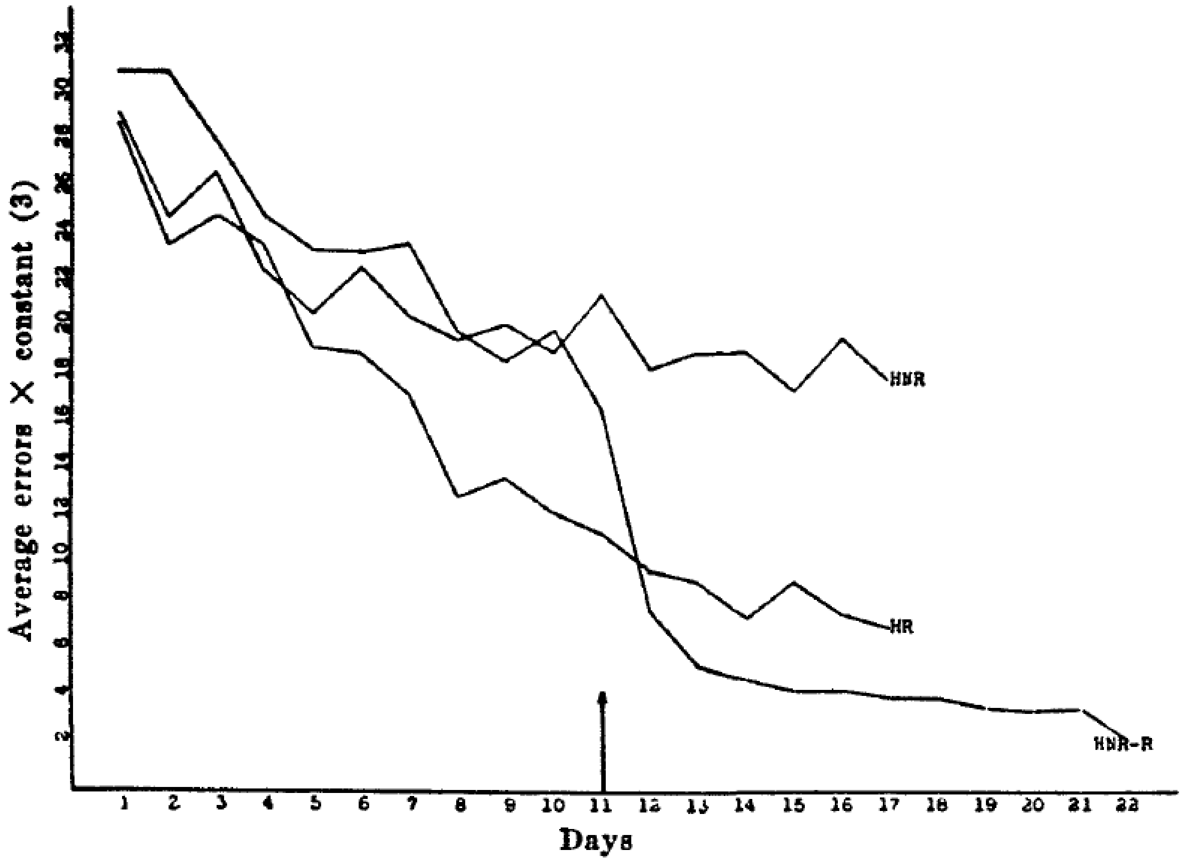

Figure 3: Rats in the HR group were rewarded food each day they navigated the maze, so learned to make fewer errors over time. Rats in the HNR group had no reward. Rats in the HNR-R group were not rewarded until day 11. Once they received a reward, their performance rapidly improved. These rats learned the layout of the maze without a reward, which allowed them to rapidly improve once the reward was introduced. Reproduced from [12].

at a set direction and fall off as the rat turned away, with an active range of about 90°. It was also shown that these cells maintained the same preferred direction within the experimental environment over periods of multiple days. This detailed quantitative analysis is the basis for many computational models of rats' navigation processes.

Finally, Hafting *et al.* in 2005 [16] discovered the existence of *grid cells* in the entorhinal cortex of rats. The firing pattern of these cells formed a tessellating grid of equilateral triangles depending on the rat's position in the testing environment. The scale of the grid varied between individual cells, and between animals, but the shape of the firing fields was consistent. The importance of visual and self-motion cues to this cell behaviour was examined. A single visual cue in the environment was rotated, causing the firing patterns of the grid cells to rotating by the same amount. A second test was performed by turning off the lights after the rats had explored the space for a short time. In the absence of visual cues, the firing patterns of the grid cells were only slightly displaced from the original pattern.

It was then proposed that the lower level information from grid cells and head direction cells was combined in the hippocampus to produce the higher level behaviour of place cells, only responding when a rat was in a certain location and facing a certain direction [16].

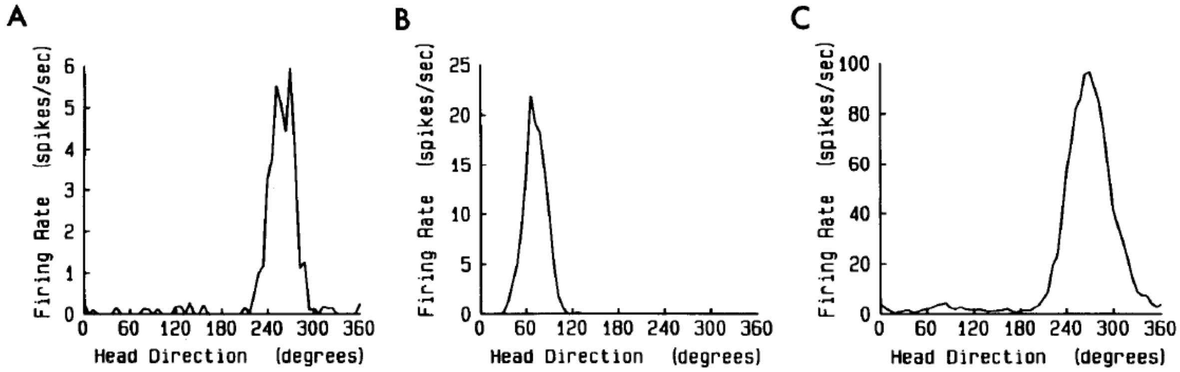

Figure 4: Firing rate curves for 3 different head direction neurons from 3 different animals. It can be seen that each cell has a distinct preferred direction at which it fires maximally, and that the peak firing rate is different for each of the cells. Reproduced from [15].

O'Keefe, Moser and Moser (under Hafting) were jointly awarded the Nobel Prize in Physiology or Medicine in 2014. Their work provided evidence for Tolman's cognitive map in the brains of rats, and later work has found evidence of similar neural structures in humans.

Computational Models

Particularly following the work of Taube *et al.* [15] to quantitatively describe the firing of head direction cells, computational models began to be developed to produce similar behaviour. Most models, such as that by Redish *et al.* [17], were based on competitive attractor networks. These networks represent a direction as a "bump" of activity, which can move in response to a rotation and to be reset to a location when seeing a visual cue.

These computational models were applied to the task of robot SLAM by Milford *et al.* in 2004 [6]. Competitive attractor networks were used to track the position (2D) and orientation (1D) of a robot like the head direction and grid cells in rats – hence the approach was called RatSLAM. Instead of operating these networks separately, a 3 dimensional grid of "pose cells" allowed multiple independent hypotheses for pose to be tracked simultaneously. Recurrent connections within the attractor network, offset to the left or right, allowed activity to be moved around the network by a velocity input coming from the robot's wheel encoder sensors. A learning process was applied based on the input from a camera so the robot knew what network activity corresponded to a certain visual input. That activity could then be injected back into the network when the same visual input was received. This provided a novel approach to the problem of SLAM compared to those discussed previously [6].

Work into RatSLAM has continued, and in 2015 a version of RatSLAM using multiple pose cell networks at different scales was able to outperform two state of the art SLAM approaches [19]. But RatSLAM suffers from the same limitations as other continuous attractor networks – that they must be precisely calibrated to achieve the desired performance, often by hand. This requires assumptions and simplifications contrary to biological knowledge. For example, most networks are constructed in a completely symmetrical way, while measurements of rat neurons have shown large variability in cell activity. Additionally, attractor networks must be tuned specifically for the expected velocity inputs, and

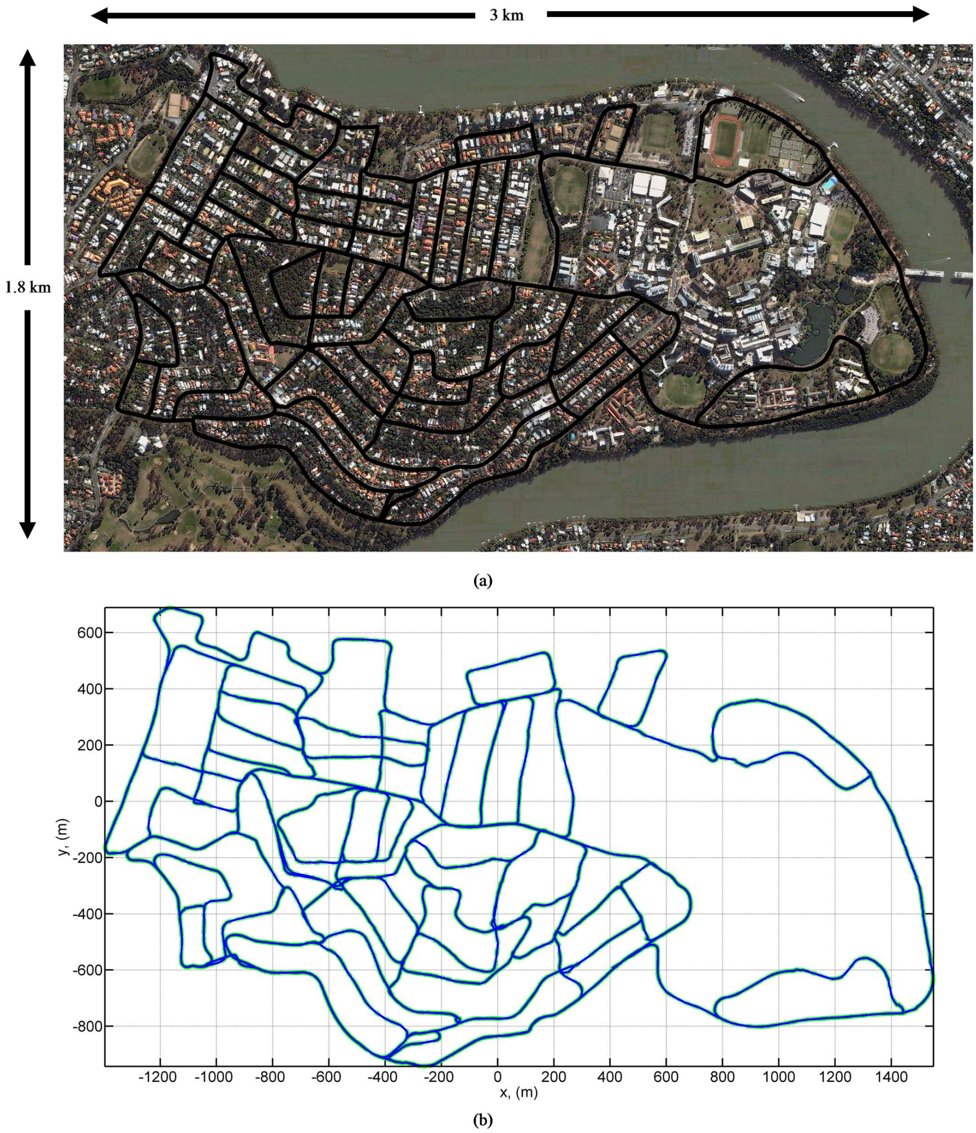

Figure 5: RatSLAM was used to map a 66 km route through the suburb of St Lucia in Brisbane. The calculated map (bottom) closely matches reality. Reproduced from [18].

only function over a fairly limited range of values [17].

This presents the opportunity to apply "artificial evolution" techniques to the task of calibrating, or even generating, a continuous attractor network. Limited work has so far been undertaken in this field, but with recent advances in available algorithms and computing power it is proposed that there is the opportunity for development.

Evolutionary Approaches – "Learning" to navigate

There are many terms for the field of applying biologically inspired or evolutionary methods of optimisation or learning. Approaches relevant to this project will be categorised as either machine learning or evolutionary algorithms. Broadly, both approaches are techniques for optimising the parameters in a model to improve that model's performance at a particular task. But there are some distinguishing differences between the two.

Machine Learning

Machine learning is a combination of statistical and computer science methods that allow a task to be described to a computer through examples, rather than by a list of rules as in a standard computer program. Several frameworks exist – including decision trees, neural networks, regression models, and support vector machines – which can classify or predict a value from input data. The outputs of these frameworks are controlled by inbuilt parameters, and statistical rules have been derived to allow these parameters to be updated to fit known example data [20].

An example machine learning problem involves categorising images of dogs and cats. A model would be chosen based on the complexity of the problem, and its parameters initialised to some default. A series of images of dogs and cats would be required, labelled correctly. Running the machine learning algorithm applicable to the model would optimise its parameters to correctly classify these images. The model would ideally then generalise to other unseen images, but this depends on the generality of the training data. In this way, a model could be produced to complete a complex task, without needing a rigid mathematical description [21].

One particular machine learning approach of note is Hebbian learning, a process for tuning the connection strengths of neurons in a network by reinforcing the connections between similarly active neurons and decreasing the connection strength between seemingly unrelated neurons [22].

Evolutionary Algorithms

In contrast to machine learning, evolutionary algorithms are typically applied in situations where mathematical or statistical optimisation is not possible, usually because we don't know exactly how a model's parameters influence its output. Instead of starting with one model and incrementally improving it, evolutionary approaches generate a "population" of possible models, evaluate the success of each, then generate a new population by altering and combining the more successful models. The solution space is therefore explored somewhat randomly, but mostly around the most successful current solutions [23].

Evolutionary algorithms can be applied to optimise almost any model with parameters that can be altered, and for which two instances of the model, the parents, can be "crossed over" to produce another individual with some of the characteristics of each parent. A means of evaluating the fitness of an individual is required, but this can be achieved by running a simulation and comparing the result to an expected or true value.

Evolutionary Algorithms for Neural Networks

While evolutionary algorithms have been used to solve many problems, including optimising the weighting parameters of neural networks, there have traditionally been difficulties using these techniques to evolve neural network structures. This is chiefly because an arbitrarily structured neural network is so general – it is not clear which neurons or connections are responsible for which aspects of the overall network's behaviour, and many different networks are possible which result in the exact same behaviour. Therefore, crossing over two neural networks is likely to give performance significantly worse than both parents. Combining networks in a meaningful way is very difficult.

NeuroEvolution of Augmenting Topologies (NEAT) is a technique proposed by Stanley and Miikkulainen [24] that addresses a number of the issues with evolving neural networks. A marking scheme for neuron mutations was proposed, tracking how the population of networks evolves over time. When crossing over two networks, their mutations can be traced back to a common point and combined effectively. A speciation system was also laid out, where innovative neural networks are only initially made to compete with those of a similar structure or "species". This allows different distinct solutions to be evolved simultaneously instead of being killed off the by the existing dominant solution. Finally, the structure is initially evolved from a network of minimum complexity, allowing optimisation and complexification to occur simultaneously.

These three developments were shown to significantly increase the performance of evolved networks compared to those generated through other then-current methods. It also showed that evolving both the weights and structure of the network led to better results than only evolving the weights. As the approach can be generalised to networks with recurrent connections, it could in theory be applied to pose cell networks. This has not yet fully been realised, but work that partially achieves this will be reviewed.

Evolutionary Algorithms applied to Head Direction Cell Network Parameters

Early work applying an evolutionary approach to tuning a head direction cell network was done by Degris *et al.* [25]. In this case, a spiking neural network model was used where neurons spike at discrete times. This is more biologically realistic than standard rate-based models, at the cost of increased computational complexity. The strength and standard deviation of Gaussian weighting curves for the head direction and odometry neurons were obtained through a genetic algorithm that optimised the network's response. This was a successful application of an evolutionary approach, but only concerned 7 discrete parameters, with many simplifying assumptions made.

Work presented by Kyriacou [26] is the most thorough in the field of using an evolutionary approach to tune an attractor network. In this case, a one-dimensional head direction network was modelled, with visual (camera), vestibular (gyroscope) and kinaesthetic (odometry) inputs. The strengths of the recurrent connections between head-direction cells were trained by a Hebbian learning process from data recorded on a real robot. The same learning process was applied for the connections from head-direction cells to vestibular and kinaesthetic cells.

A $(\mu+\lambda)$ evolutionary strategy (as described by Beyer & Schwefel in [27]) was then used to determine the relative weightings of the recurrent, visual, vestibular and kinaesthetic connections. A second series of training data was manipulated by removing combinations of various sensors for set periods of time. This represented a variety of scenarios in which the network should be able to operate. Parameters were evaluated for fitness by comparing the output of the network to truth data. When visual data was available, the network should match the robot's location; when no visual data was available, the network should integrate the vestibular and/or kinaesthetic inputs.

The work was successful in evolving the relative strengths of each data input to produce satisfactory network behaviour. But a large number of parameters were still set manually – relating to neuron time constants, learning rates, activity packet size, Δt, and others. The author states that they could be selected through evolution as well, but this was not feasible at the time within the limitations of available computational power [26].

Evolving the Structure of a Head Direction Cell Network

All the work so far presented has focuses on evolving the parameters of a network with a fixed structure. There has not been any research into evolving the complete structure of a pose cell network from scratch.

However, some research by Haferlach *et al.* [28] successfully applied an evolutionary approach to a similar field – generating homing behaviours for insects. They utilised only 3 "head directions sensors" (an abstraction of an already functioning head direction network) and evolved the connections from these to a robot's motors to allow the robot to perform path integration during motion so it could return to its original position. The genetic algorithm used was initially found not to converge, but was successful after a number of further limitations were imposed.

To increase the number of head direction cells in the network, the structure that was evolved using only 3 cells was replicated for larger numbers. This structure was then fixed and a genetic algorithm used to find the parameters for connection strengths. Overall, this work was successful in evolving the structure of a network built on top of a head direction network to achieve the desired homing behaviour, but it abstracted away the complexity of the head direction network itself. Nevertheless, it provides a useful template for applying the approach to generating pose cell networks. Additionally, many of the limitations experienced by the researchers with the evolutionary approach used are those which the NEAT approach is able to resolve or mitigate.

Recent Work into Evolving a Generic Network for Navigation

Recent work in this field has been carried out by groups from Deepmind and Columbia University. They have both trained generic recurrent neural networks for the task of path integration. The key point from both pieces of research is that neurons in the network began to exhibit firing patterns consistent with grid cells and head direction cells.

Banino *et al.* from Deepmind in [29] trained a network of 128 LSTM units to integrate velocity information: the magnitude of linear velocity, and the sine and cosine of angular velocity. The output of the LSTM layer was fed through a fully connected linear layer to produce the outputs of position and head direction. When 50% dropout was applied to the linear layer, some cells in that layer were found to exhibit responses like grid cells and head direction cells. Without dropout, no grid cell-like behaviour arose.

Cueva and Wei from Columbia University in [31] trained a network of 100 recurrently connected neurons to take the linear velocity and current direction of a simulated agent and to output its x and y position over time. Regularisation was used to penalise high weights and high cell firing rates, as a proxy for metabolic cost, and Gaussian noise was introduced into the system. It was found that training with suitable regularisation tended to evolve neurons with responses similar to those of grid cells and head direction cells in the brain.

Apart from outlining some interesting and potentially useful methodologies, the main conclusion from both of these pieces of research is that grid cells are likely an efficient means of encoding position and orientation information, if they tend to appear when

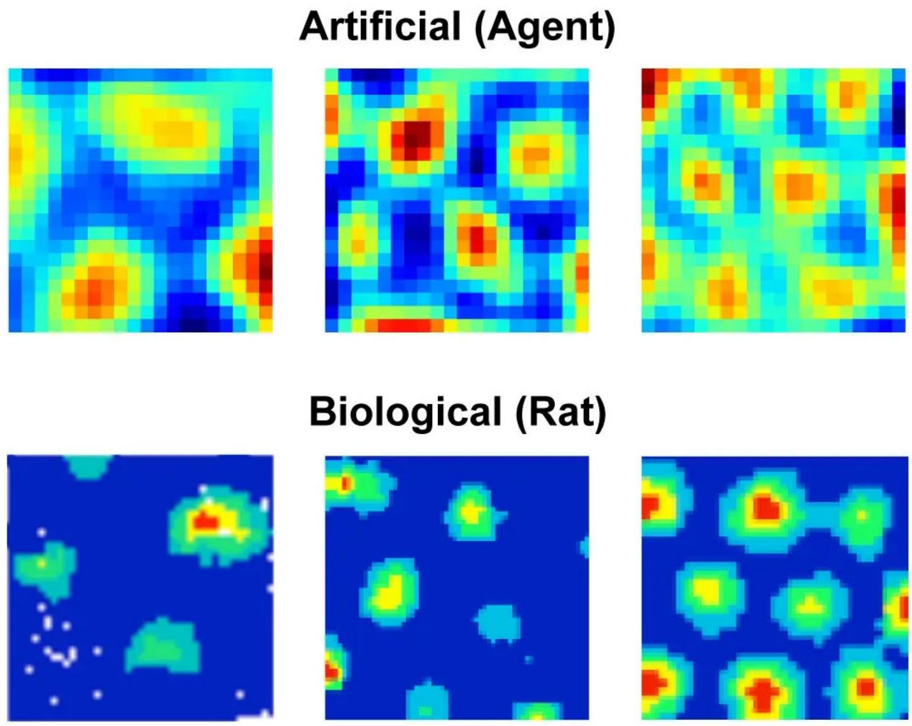

Figure 6: Results from [29], showing how their artificial evolution process gave rise to grid-like cells similar to those seen in the brains of rats. Reproduced from [30].

evolving a completely general neural network.

Research Outcomes

While still at university prior to my AGA trip, I completed research in this field and published the results at the Australasian Conference on Robotics and Automation [32] (available here). This experience of research drove me to apply for the AGA-Goethe Fellowship to further pursue international opportunities.

After returning from Germany I began work as a research assistant at QUT. In this role I continued research on biologically-inspired navigation for robots. In particular, I modelled an approach used by ants when exploring for food and applied this to robots. When returning to their nests after searching for food, ants seem to follow a fairly simple strategy – they remember a path as a series of landmarks. When repeating the path, they stop at each landmark and spin on the spot until what they see lines up with how they remember the view [33]. This allows them to robustly repeat long routes without getting lost.

I used a similar approach for robots called *teach and repeat*. I first drove the robot along a route, which it learned by capturing images and using wheel sensor information. The robot then repeated the route in a similar way to an ant, correcting for errors along the path if it noticed the scene looked different to what it remembered. More details of this work are contained in a paper I wrote, currently under review [34] (available here).

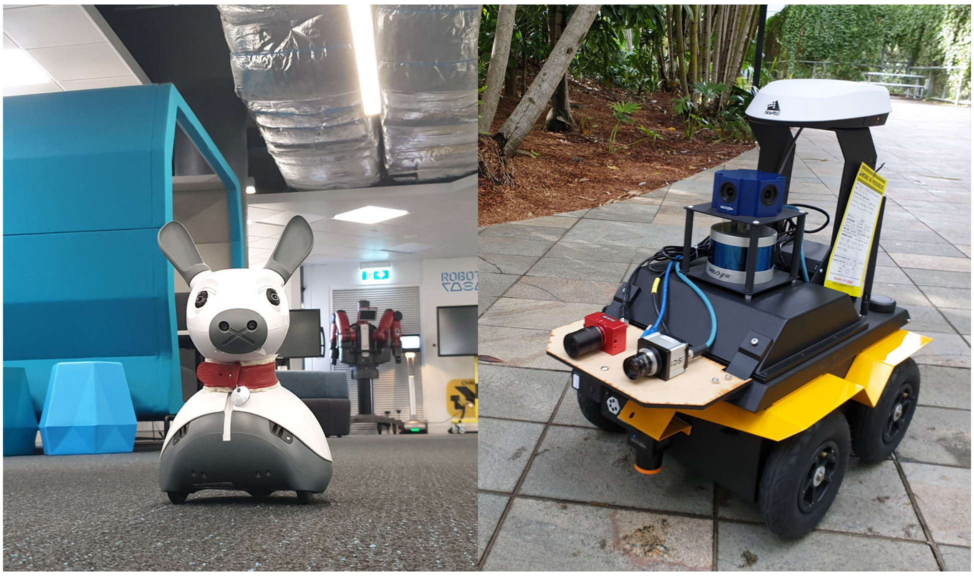

Figure 7: Robots used for my further research into bio-inspired robot navigation: Miro (left) and Jackal (right). Both used cameras to remember then repeat a route through the environment.

Attending Conferences in Berlin

I started my research and study tour in Berlin, flying in from Australia on 11 September, 2019. Two major computational neuroscience conferences were taking place at the Technical University of Berlin – the Computational Cognitive Neuroscience Conference and the Bernstein Conference.

A highlight of the conferences as a whole was how large and active the research communities in Europe were. It was amazing to see whole labs worth of PhD students all attending the conference, obviously made much easier by the small distances many had to travel compared to us in Australia. By the end of the week I had become an adopted member of the lab from Uni Tübingen – so I must give them all my thanks. We had many interesting discussions together at a variety of fantastic Berlin restaurants. Overall, I felt really welcomed by all the researchers I met over the week.

Another highlight for me, research-wise, was seeing a presentation by Chris Cueva from Columbia University. I had been following his work closely because it was strongly related to my previous research. It was really insightful to hear his presentation, and I walked away with a lot more details and ideas. In particular, he raised some salient points about how his research was still somewhat distant from the biology, and how it might be modified to be more relevant.

A huge coincidence occurred when I was waiting in the line for dinner during the first evening. I started talking to the students just behind me, only to find out they were all from the Ruhr University of Bochum, where I had already arranged to spend time as a guest researcher. They introduced me to their Professor, Sen Cheng, with whom I had spoken over email, and we ironed out the details of my visit which I will detail later in this report.

Another highlight for me was seeing the novel work from a group from University Zurich and ETH Zurich, who had developed a rat brain-inspired model for navigation for a robot. This tied in very closely with my previous research, and it was really interesting to see how they had successfully managed to run the system on a real robot.

On the final two days of the conference we had a symposium for all the computational neuroscience PhD students. It was organised by students, for students, and I found it to be a really insightful experience. We discussed the overall research landscape and techniques for how to get the most out of our time as students. A highlight of the symposium was the presentation by Ashley Juavinett, a neuroscientist who also focuses on how we can use what we know about the brain to improve teaching new students. She gave us a lot of career-related advice, backed up by statistics about student outcomes and employer feedback. She detailed the many opportunities associated with higher research studies – but also put them into perspective with alternative career paths not in academia.

(a) Flying over central Australia on the way to Germany.

(b) The Lichthof of the TU Berlin, the venue for CCN and Bernstein.

Intensive German Courses in Munich

After Berlin, I travelled by train down to Munich for the first of my two months of intensive German language courses at the Goethe Institute. Upon arriving at the institute for the first time, I was greeted and taken for an interview. I was rather amused at the end to be told I had “Wild-Deutsch”, owing to my previous time on exchange in Germany where I mostly learned the language by speaking with others, acquiring a mix of different dialects while missing many important grammatical points. So now the Goethe institute had the job of helping me cultivate my German.

I found the classes really well delivered and I learnt a lot in a short time. I was in a class with students from so many different backgrounds, all with different reasons for learning German. But we shared a similar commitment. A highlight of the class were the debates we would have about topics including travel and the environment. This really helped to build my confidence in repurposing the words I knew to speak about unfamiliar topics. I also had a lot of fun preparing a presentation for my class about different cultural celebrations in Australia.

Another great part of the experience was living in the guesthouse with 100 other students from GI Munich. Being in an environment where we were united by our motivation for learning German really helped to build my confidence in speaking and listening. Thanks must also go to the team of student interns who lived with us and also ran the cultural event program for the institute. We had events and trips almost every day, allowing us to get to know many of the sights of Munich.

(a) BMW's Research Centre in Munich. Unfortunately, pictures weren't allowed inside the building.

(b) The beautiful sunset I was greeted with on my first night at the guesthouse in Munich, which set the scene for my time there.

(c) My German class at GI Munich - B2.1 with our amazing teacher Frau Pickl.

(d) With my classmate Deniz, I was the face of GI Munich's facebook page for a day.

Apart from my time at Goethe, I also met up with a friend from my home university who was now working for BMW's research and development centre. I spent two days touring their facilities and seeing their research, and I even got to go for a drive! Most of their work was quite sensitive so unfortunately I can't explain it in detail.

I also met up again with Florian Röhrbein, the former professor from the Technical University of Munich whom I'd initially met in Berlin. He introduced me to a small robotics startup working in Munich, and I had the great opportunity to work part time with them while I stayed in Munich. I was glad to be able to contribute my knowledge of developing simulations of robots to help them in their work. Unfortunately, this work was also rather secretive, but I can share my surprise when walking into a normal suburban house for our first meeting over breakfast, only to find a large robot sitting in the living room. This definitely was a startup!

Intensive German Courses in Berlin

My stay in Berlin was with a home stay, right in the middle of Berlin. It was a unique experience to live right next door to the Hackesher Höfe, and only a few minutes' walk to the Goethe Institute. The institute itself had a different feeling to that in Munich, but was equally well organised when it came to classes. Over the month, I felt I improved immensely in being comfortable speaking German. It was a bit of a shock, but really exciting to realise on a few occasions that I'd gone a whole day speaking only German. Thanks for this must also go to the amazing cultural program offered by Goethe Berlin. Almost every day we had a tour on offer through another part of Berlin, focusing on another one of the many layers of history the city has, from its medieval roots to the present-day issues with rapid gentrification.

A nice coincidence for me was that my time in Berlin happened to coincide with the city's Science Week. I had the chance to attend a number of events as part of this. Another highlight was visiting the Futurium museum which was full of exhibits on the potential for robotics, AI and other new technologies. This made for many interesting discussions with my classmates.

Another significant event in the city while I was there was the 30th anniversary of the fall of the Berlin Wall on 9 November 1989. Historical installations were set up at significant sites throughout the city. The week culminated in a huge celebration at the Brandenburg Gate. Many of the tours offered as part of Goethe's cultural program were by guides who had lived on either side of the wall during that time. It was fascinating to hear their first-hand accounts of this pivotal moment in human history.

At the end of my month in Berlin, and my two months of intensive German courses, I successfully completed the Goethe B2 level German exam, which was a great achievement for me. But more satisfying was the fact that I had lived successfully in Berlin predominantly speaking German.

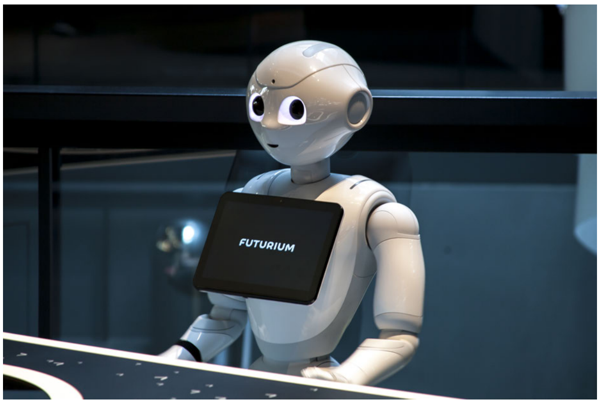

(a) A Pepper robot (which I work with in Brisbane) at the Futurium museum.

(b) My class at GI Berlin – B2.3 with our teacher Herr Wilkening.

(c) Alexanderplatz, celebrating the 30th anniversary of the fall of the Berlin Wall.

(d) 30th Anniversary celebration concert at the Brandenburg Gate.

Research Stay at Ruhr University Bochum

After leaving Berlin I had a relatively short train journey to Bochum, in the west of Germany. Here I would spend two weeks at the Ruhr University thanks to Professor Sen Cheng, the head of the computational neuroscience research group at Bochum whom I'd met earlier during my visit to Berlin. I arrived just in time for the weekly lab meeting which served to introduce me to the group. Again, thanks to the conferences in Berlin, I saw some familiar faces. Over my short visit I would get to know their individual areas of research, which I found very interesting. But the main impression on me was how all these research areas fit together into a bigger picture for the group as a whole.

I can't list all the researchers with whom I had the chance to speak. But some of the highlights were: Dr Mathis Richter, who works with brain-inspired models of computation and applies them to controlling robotic arms in a robust way; Professor Tobias Glasmachers, who specialises in numerical optimisation techniques and who was able to provide very helpful advice on my use of these techniques; Dr Sebastian Houben, who conducts research with drones and in particular took part in a competition very related to a previous internship I had done at CSIRO back in Brisbane; and PhD student Sandhiya Vijayabaskaran and postdoctoral researcher Dr Tobias Walther who were working in a similar area to me, investigating how the brain can perform spatial navigation tasks and how it represents information while doing this. I also greatly enjoyed discussing with the large group of PhD students in the lab to learn about their research and any advice they had for choosing a direction after completing an undergraduate degree.

A cultural highlight of my time was visiting the musical Starlight Express. There is a custom-built theatre in Bochum that only stages this production, which is about trains played by actors wearing roller skates and performing stunts. It was a very funny and quirky experience, In fact, it holds the record for being the most visited musical running in a single theatre, with some 15 million visitors.

On my last day in Bochum, the Institute of Neuroinformatics held their Christmas party to end the year. It was a fantastic opportunity to thank all those who had made my stay such a great experience. This also marked the end of my time in Germany and the beginning of my trip back to Australia.

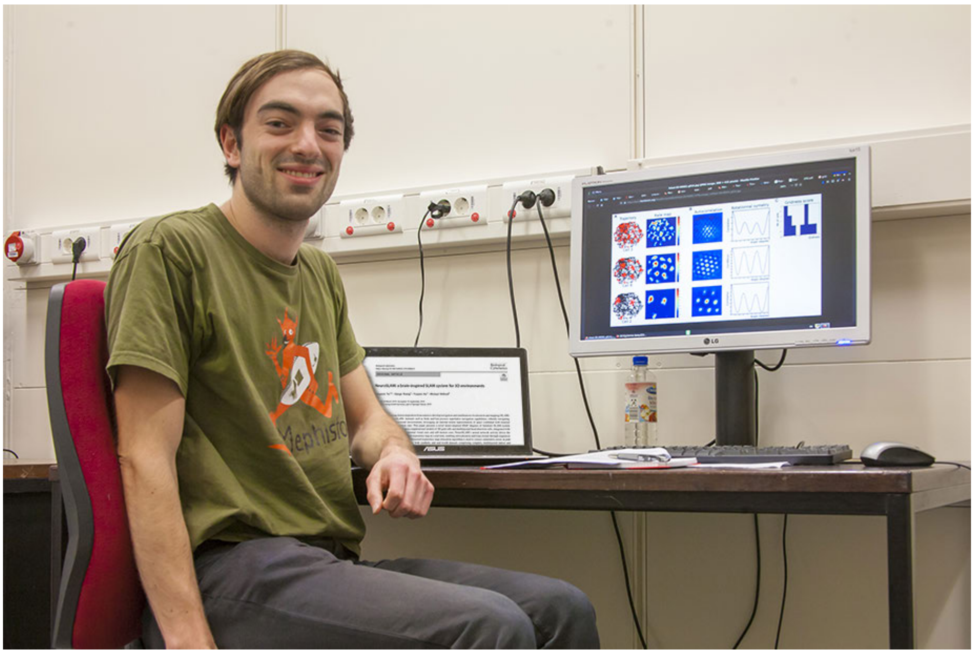

Figure 11: My desk at the Ruhr University of Bochum.

Professional Benefits and Outcomes

The experience of the AGA-Goethe Fellowship has provided me many benefits since returning from the trip. In particular, I continued research in biologically-inspired robot navigation after returning home, and I could leverage the knowledge I gained in Germany to further this research. The results of this work are under review for publication.

The AGA trip also affirmed my decision to pursue further studies in the area of neuroscience and robotics. In fact, I have recently started a masters degree at the University of Zurich, studying a combination of neuroscience, engineering and informatics. The experiences during my trip have continued to serve me well as I am already familiar with many research topics in the area, particularly after having seen them presented at the conference. Additionally, my German knowledge has made my life here significantly easier, but I am still acclimatising to the Swiss German dialect.

Finally, I remain in contact with those researchers I met on the trip. In particular with Florian and the team in Munich and with Professor Sen Cheng and some of his students at the Institute for Neuroinformatics in Bochum. Unfortunately, now is not the best time for me to travel and meet with them again, but when the situation allows I will jump at the chance.

Figure 12: Hiking in the alps in Switzerland – my new home.

Conclusion

Overall I can only repeat what a great experience the AGA-Goethe Fellowship was for me. I have benefited immeasurably from the trip, both personally and professionally. I have gained a wide network of friends and collaborators and dramatically improved in my German. I continue to reflect on my experiences and how they have helped prepare me for my current masters studies in Zurich.

The current world situation and difficulty of travelling has made me appreciate even more how lucky I was to have had this opportunity when I did. Unfortunately I have not yet been able to return to Melbourne to talk about my experiences and express my thanks to AGA and Goethe Institute in person. But I wish all the best to both organisations and look forward to the time when the AGA-Goethe Fellowship can continue to provide young people such as myself with amazing opportunities for international collaboration.

Bibliography

[1] "Global Mobile Robotics Market Research Report- Forecast 2023 — MRFR." [Online]. Available: https://www.marketresearchfuture.com/reports/mobile-robotics-market-1946

[2] A. Sharma, "Logistics Mobile Robots Forecast to Surpass $3bn in 2022." [Online]. Available: https://www.interactanalysis.com/logistics-robots-forecast-growth/

[3] L. Adouane, *Autonomous Vehicle Navigation: From Behavioral to Hybrid Multi-Controller Architectures.* Boca Raton: CRC Press, 2016. ISBN 978-1-4987-1558-4

[4] S. Herculano-Houzel, "Isotropic Fractionator: A Simple, Rapid Method for the Quantification of Total Cell and Neuron Numbers in the Brain," *Journal of Neuroscience*, vol. 25, no. 10, pp. 2518–2521, Mar. 2005. https://www.doi.org/10.1523/JNEUROSCI.4526-04.2005

[5] F. A. C. Azevedo, L. R. B. Carvalho, L. T. Grinberg, J. M. Farfel, R. E. L. Ferretti, R. E. P. Leite, W. J. Filho, R. Lent, and S. Herculano-Houzel, "Equal numbers of neuronal and nonneuronal cells make the human brain an isometrically scaled-up primate brain," *Journal of Comparative Neurology*, vol. 513, no. 5, pp. 532–541, 2009. https://www.doi.org/10.1002/cne.21974

[6] M. Milford, G. Wyeth, and D. Prasser, "RatSLAM: A hippocampal model for simultaneous localization and mapping," in *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA '04. 2004.* New Orleans, LA, USA: IEEE, 2004, pp. 403–408 Vol.1. https://www.doi.org/10.1109/ROBOT.2004.1307183. ISBN 978-0-7803-8232-9

[7] P. Corke, *Robotics, Vision and Control*, ser. Springer Tracts in Advanced Robotics, B. Siciliano and O. Khatib, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, vol. 73. ISBN 978-3-642-20143-1 978-3-642-20144-8

[8] H. Durrant-Whyte and T. Bailey, "Simultaneous localization and mapping: Part I," *IEEE Robotics & Automation Magazine*, vol. 13, no. 2, pp. 99–110, Jun. 2006. https://www.doi.org/10.1109/MRA.2006.1638022

[9] T. Bailey and H. Durrant-Whyte, "Simultaneous localization and mapping (SLAM): Part II," *IEEE Robotics & Automation Magazine*, vol. 13, no. 3, pp. 108–117, Sep. 2006. https://www.doi.org/10.1109/MRA.2006.1678144

[10] C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. J. Leonard, "Past, Present, and Future of Simultaneous Localization and Mapping: Toward the Robust-Perception Age," *IEEE Transactions on Robotics*, vol. 32, no. 6, pp. 1309–1332, Dec. 2016. https://www.doi.org/10.1109/TRO.2016.2624754

[11] S. Lowry, N. Sunderhauf, P. Newman, J. J. Leonard, D. Cox, P. Corke, and M. J. Milford, "Visual Place Recognition: A Survey," *IEEE Transactions on Robotics*, vol. 32, no. 1, pp. 1–19, Feb. 2016. https://www.doi.org/10.1109/TRO.2015.2496823

[12] E. C. Tolman, "Cognitive maps in rats and men." *Psychological Review*, vol. 55, no. 4, pp. 189–208, 1948. https://www.doi.org/10.1037/h0061626

[13] F. Strumwasser, “Long-Term Recording from Single Neurons in Brain of Unrestrained Mammals,” *Science*, vol. 127, no. 3296, pp. 469–470, 1958.

[14] J. O’Keefe and J. Dostrovsky, “The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat,” *Brain Research*, vol. 34, no. 1, pp. 171–175, Nov. 1971. https://www.doi.org/10.1016/0006-8993(71)90358-1

[15] J. Taube, R. Muller, and J. Ranck, “Head-direction cells recorded from the post-subiculum in freely moving rats. I. Description and quantitative analysis,” *The Journal of Neuroscience*, vol. 10, no. 2, pp. 420–435, Feb. 1990. https://www.doi.org/10.1523/JNEUROSCI.10-02-00420.1990

[16] T. Hafting, M. Fyhn, S. Molden, M.-B. Moser, and E. I. Moser, “Microstructure of a spatial map in the entorhinal cortex,” *Nature*, vol. 436, no. 7052, pp. 801–806, Aug. 2005. https://www.doi.org/10.1038/nature03721

[17] A. D. Redish, A. N. Elga, and D. S. Touretzky, “A coupled attractor model of the rodent head direction system,” *Network: Computation in Neural Systems*, vol. 7, no. 4, pp. 671–685, Jan. 1996. https://www.doi.org/10.1088/0954-898X_7_4_004

[18] M. Milford and G. Wyeth, “Mapping a Suburb With a Single Camera Using a Biologically Inspired SLAM System,” *IEEE Transactions on Robotics*, vol. 24, no. 5, pp. 1038–1053, Oct. 2008. https://www.doi.org/10.1109/TRO.2008.2004520

[19] Z. Chen, S. Lowry, A. Jacobson, M. E. Hasselmo, and M. Milford, “Bio-inspired homogeneous multi-scale place recognition,” *Neural Networks*, vol. 72, pp. 48–61, Dec. 2015. https://www.doi.org/10.1016/j.neunet.2015.10.002

[20] M. Mohammed, M. B. Khan, and E. B. M. Bashier, *Machine Learning: Algorithms and Applications.* Boca Raton: CRC Press, Taylor & Francis Group, 2017. ISBN 978-1-4987-0538-7

[21] J. Watt, R. Borhani, and A. Katsaggelos, *Machine Learning Refined: Foundations, Algorithms, and Applications.* Cambridge: Cambridge University Press, 2016. ISBN 978-1-316-40227-6

[22] D. O. Hebb, *The Organization of Behavior : A Neuropsychological Theory.* Psychology Press, Apr. 2005. ISBN 978-1-4106-1240-3

[23] A. Petrowski, *Evolutionary Algorithms.* Hoboken, NJ: ISTE Ltd/John Wiley and Sons Inc, 2017. ISBN 978-1-84821-804-8

[24] K. O. Stanley and R. Miikkulainen, “Evolving Neural Networks through Augmenting Topologies,” *Evolutionary Computation*, vol. 10, no. 2, pp. 99–127, Jun. 2002. https://www.doi.org/10.1162/106365602320169811

[25] T. Degris, L. Lachèze, C. Boucheny, and A. Arleo, “A Spiking Neuron Model of Head-Direction Cells for Robot Orientation,” in *From Animals to Animats 8: Proceedings of the Eighth International Conference on the Simulation of Adaptive Behavior.* The MIT Press, 2004. ISBN 978-0-262-29144-6

[26] T. Kyriacou, “Using an evolutionary algorithm to determine the parameters of a biologically inspired model of head direction cells,” *Journal of Computational Neuroscience*, vol. 32, no. 2, pp. 281–295, Apr. 2012. https://www.doi.org/10.1007/s10827-011-0352-x

[27] H.-G. Beyer and H.-P. Schwefel, “Evolution strategies – A comprehensive introduction,” *Natural Computing*, vol. 1, no. 1, pp. 3–52, Mar. 2002. https://www.doi.org/10.1023/A:1015059928466

[28] T. Haferlach, J. Wessnitzer, M. Mangan, and B. Webb, “Evolving a Neural Model of Insect Path Integration,” *Adaptive Behavior*, vol. 15, no. 3, pp. 273–287, Sep. 2007. https://www.doi.org/10.1177/1059712307082080

[29] A. Banino, C. Barry, B. Uria, C. Blundell, T. Lillicrap, P. Mirowski, A. Pritzel, M. J. Chadwick, T. Degris, J. Modayil, G. Wayne, H. Soyer, F. Viola, B. Zhang, R. Goroshin, N. Rabinowitz, R. Pascanu, C. Beattie, S. Petersen, A. Sadik, S. Gaffney, H. King, K. Kavukcuoglu, D. Hassabis, R. Hadsell, and D. Kumaran, “Vector-based navigation using grid-like representations in artificial agents,” *Nature*, vol. 557, no. 7705, pp. 429–433, May 2018. https://www.doi.org/10.1038/s41586-018-0102-6

[30] A. Banino, D. Kumaran, B. Caswell, and U. Benigno, “Navigating with grid-like representations in artificial agents,” May 2018. [Online]. Available: https://deepmind.com/blog/article/grid-cells

[31] C. J. Cueva and X.-X. Wei, “Emergence of grid-like representations by training recurrent neural networks to perform spatial localization,” *arXiv:1803.07770 [cs, q-bio, stat]*, Mar. 2018.

[32] D. Dall’Osto, S. Hausler, A. Jacobson, and M. Milford, “Automatic Calibration of a Biologically Inspired Neural Network for Robot SLAM,” in *Australasian Conference on Robotics and Automation 2018*, Christchurch, 2018. ISBN 978-1-5108-7958-4

[33] B. Webb, “The internal maps of insects,” *The Journal of Experimental Biology*, vol. 222, no. Suppl 1, p. jeb188094, Feb. 2019. https://www.doi.org/10.1242/jeb.188094

[34] D. Dall’Osto, T. Fischer, and M. Milford, “Fast and Robust Bio-inspired Teach and Repeat Navigation,” *arXiv:2010.11326 [cs]*, Oct. 2020.