400-Level Project Proposals 2019

A list of possible 400-level projects is given below. You do not have to restrict yourself to this list, and can make up your own project topic. However, to do so you will have to arrange for a supervisor. You can find more information about the research interests of the staff members on the Computer Science Department Web pages.

You should email your project selections to the project co-ordinator, 400projectadmin@cs.otago.ac.nz, by **Friday 1st March**. Before making these selections you should talk to the supervisors of the projects you are interested in. In most cases the best way is to email them to make a time to do so – supervisors' email addresses are given with each project.

In order to allocate the projects as fairly as possible we ask that you give first, second, and third choice projects, and that you choose projects from three different supervisors.

There are two different project papers - COSC480 and COSC490. The relevant Web pages (for COSC480 and COSC490) have more details, but the short version is that if you are enrolled in the Computer Science Honours degree you take COSC490, students in the first year of a Masters take either paper, and most others take COSC480. Most projects are suitable for either COSC480 or COSC490, but a few might be only suitable for COSC480, which will be indicated after the project title.

1. **Quantifying conceptual density in text**

   Anthony Robins (anthony@cs.otago.ac.nz)
   Alistair Knott

   Can we identify the density of connections between individual concepts in a body of text? In other words, can we identify discrete concepts, and quantify the relationships between them? This project will involve investigating existing natural language processing tools, and perhaps developing our own.

2. **Attractor spaces in Hopfield nets**

   Anthony Robins (anthony@cs.otago.ac.nz)

   Hopfield networks (a kind of neural network) have dynamic behaviour that can be characterised in terms of gradient descent in a multidimensional attractor space. We know a fair bit about the structure of such spaces. This project will involve implementing and exploring a Hopfield type network to further develop our understanding of its dynamic behaviour.

3. **Serial learning in neural networks**

   Anthony Robins (anthony@cs.otago.ac.nz)

   Neural networks are very powerful learning systems. But most of them work best when all training data is available simultaneously (concurrent learning). In the real world humans and other animals learn different things on different occasions (serial learning). This project will review the current state of serial learning in neural networks, and further explore my "pseudorehearsal" solution.

4. **Learning to play**

   Michael Albert (malbert@cs.otago.ac.nz)
   Lech Szymanski

   The recent successes of AlphaGo and AlphaZero have demonstrated that with sufficient training deep learning networks can achieve superhuman performance in board games even when primed with nothing other than the ability to recognise legal moves. But, do they play perfectly? In some combinatorial games, e.g., Nim, there are simple algorithms for recognising winning moves. So what happens when a network is trained to play such a game – does it learn those algorithms somehow? If trained in games of a certain size does its performance generalise to larger games? Possibly with a shortened second training period? Using and extending an open source framework for the learning paradigm that underlies AlphaZero these questions will be addressed for Nim and one or more additional combinatorial games.

5. **Micro-optimisations**

   Andrew Trotman (andrew@cs.otago.ac.nz)
   David Eyers

   Micro-optimisations are small changes to your source code that can (apparently) decrease execution time. The web is littered with advice on which micro-optimisations you should make and which you should not. For example, some say don't bother to convert `a = b / 2` into `a = b >> 1` because the compiler will do it for you – but will it? Making the right micro-optimisations can substantially decrease the execution time (and energy consumption) of your program, but making the wrong ones just makes your program harder to read. In this project we will examine a number of micro-optimisations and determine which are (and are not) worthwhile then apply the appropriate changes to our open source search engine and measure any change in execution time when running under high load and using standard benchmarks.

6. **Query performance prediction**

   Andrew Trotman (andrew@cs.otago.ac.nz)

   Several methods of predicting the quality of a search engine results list have been suggested. Some examine the query expressiveness while others make a prediction after examining the results list. In this project we will look closely at predictors that examine the results list and measure how effective they are. If they are effective then it should be possible to make small changes to the results list and measure the change. Indeed, it should be possible to determine the relevance of a single document by using a very short results list. If the relevance of each document is known then a ranking function based on the predictor would yield a very high-quality results list. We will build this ranker and measure its performance.

7. **Web page boilerplate removal for search**

   Andrew Trotman (andrew@cs.otago.ac.nz)

   Many web sites have boiler plating – the decoration around the page that is more a part of the site than the content (such as menus). In an eCommerce site such as eBay or TradeMe this problem is exacerbated because the sellers have their own boiler plating (such as shipping information) that is also not useful for search. It is obvious that a search engine should not index boilerplate. What is not obvious is how to remove it. In this project we will search the literature for the current state of the art in boilerplate removal, measure their effect on search performance, improve on them (especially for eCommerce search).

8. **USB keyboard and file system for xv6 (COSC480 only)**

   Zhiyi Huang (zhuang@cs.otago.ac.nz)

   Xv6 is Unix Version 6, which is used for teaching OS in many universities. I have ported it from PC to Raspberry Pi. A student last year has ported a USB driver for keyboard, which works but has many loose ends such as memory management, sound understanding of the code (documentation), etc. Since xv6 does not have a permanent file system, we would like to add a USB file system based on the existing FS of xv6. Though this project is for COSC480, the hands-on experience will enhance your working knowledge of operating systems, embedded systems, and computer architecture.

9. **Motion tracking with IMU sensors**

   Zhiyi Huang (zhuang@cs.otago.ac.nz)
   Haibo Zhang

   We have studied motion tracking with accelerometer, gyroscope, and magnetometer for several years now. We have implemented a Magdwick filter to fuse the sensors together and track, for example, arm movements with two wireless sensors. However, there are applications where magnetometer is not usable, which causes drifting issues of motion tracking. In this project we would like to investigate how to use accelerometer and gyroscope only to track motion and isplacement without losing too much precision. We will compare various filtering methods such as Kalman filter and various IMU chips to find the best approach for tracking fast movements within a short time. Some knowledge of geophysics, kinematics, math like Bayes theorem will be needed but not a prerequisite for this project.

10. **Touchscreen information panels for Orokonui Ecosanctuary**

    David Eyers (dme@cs.otago.ac.nz)
    Paul Everett (volunteer IT support and development)
    Orokonui staff (TBC)

    The Orokonui Ecosanctuary (Te Korowai o Mihiwaka) are keen to investigate deployment of self-service information displays within their main building. This project would involve building a prototype information display, along with an extensible architecture for supporting its ongoing development. The key challenges in terms of software engineering are to make the system as straightforward as possible in terms of design and reuse of existing technology. This will help ensure that the prototype is as manageable and useful as possible.

    Beyond a base platform (likely to be web-based, probably hosted on a tablet or smartphone device), there are many potential avenues to investigate in terms of location-based systems and context sensitive information display. A web-based platform that can provide context-sensitive information to visitors via their smartphones would be a stretch goal of this project's work.

11. **Interactive data visualisation using a 3D computer game engine**

    David Eyers (dme@cs.otago.ac.nz)
    Lech Szymanski

    Game engines typically aim to provide unified and optimised graphics, networking and data management functionality. Of course this feature-set is also desirable for supporting dynamic data visualisation tasks.

    This project aims to investigate the ease and efficiency of displaying and interactively exploring large-scale datasets, e.g., from research experiments, using a 3D computer game engine. (It is expected that Unity would be the game engine used, due to its popularity, practicality and portability.)

    As an example of the type of data that could be explored, we have large sets of electricity metering data. This data is currently typically displayed as a static heat-map, but the lack of interactivity precludes convenient anomaly detection and trend discovery. Some of our datasets also have geographical metadata, which might facilitate additional "map"-style interactivity to be explored, also.

12. **Distributed presentation tool**

    David Eyers (dme@cs.otago.ac.nz)

    My ideal presentation tool would: be open-source; run in any reasonably recent web browser on any typical OS without installing software; support multiple simultaneous screens with potentially different content; support vectorised graphics; have stable output (even with custom fonts, etc.); be bandwidth efficient; support animations; be responsive; allow presentation control from multiple devices; support multiple presenters' screens; and allow for embedded content such as video.

    It may not come as a surprise that I have not yet found a tool that simultaneously satisfies all of these requirements. However, there are many recent technology trends bringing all of the above closer, including HTML5's portability and reach into web browsers' capabilities, and software libraries such as pdf.js (PDF rendering in JavaScript), various web-based presentation systems (e.g., reveal.js), and open-source screen-sharing libraries.

    This project would seek to develop a workable prototype of an atypical combination of the above presentation tool features. While it is expected that a server component would be needed to support the coordination of devices within this system, this can be easily achieved using Docker containers, so server administration experience would not be required.

13. **Collaborative spreadsheet-based programming**

    David Eyers (dme@cs.otago.ac.nz)

    Spreadsheets are a highly accessible dataflow programming tool, despite most spreadsheet users not realising that they are, effectively, programming. Online collaborative tools such as Google Sheets further increase the utility of spreadsheets, by simplifying data storage and facilitating distributed interactivity.

    In this project, online, interactive spreadsheets will be embedded as a combined data and control plane within a distributed application development platform, most likely involving the newest form of cloud computing platform, namely, Function as a Service (FaaS).

    One of the open-source alternatives to Google Sheets is likely to be used in the project, to ensure that the core of the spreadsheet platform can be modified, should Google's APIs turn out to be impractically limited.

14. **Reading Historic Documents**

Steven Mills (steven@cs.otago.ac.nz)
Lech Szymanski

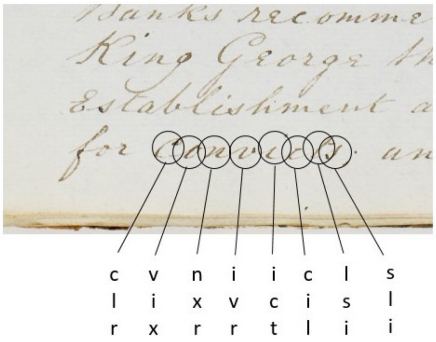

We have an ongoing project with the Hocken Collections investigating search of handwritten documents. Last year an Honours student made good progress in recognising individual letters, and this year we'd like to build on that work. There are a few directions that this could take such as looking for known patterns of letters to find keywords, or trying to identify large features such as words and lines.

15. **Gestures in the Sandbox**

Steven Mills (steven@cs.otago.ac.nz)

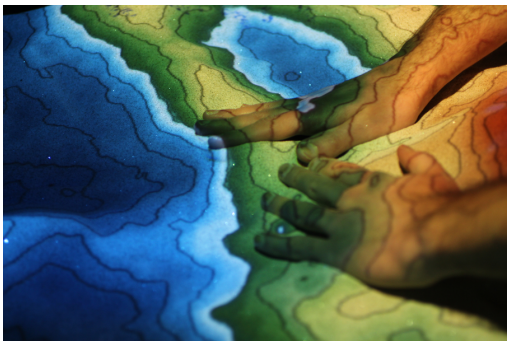

We have recently built and augmented reality sandbox. People interact by moving the sand, and there is simple recognition of hands to create a rain effect. We would, however, like to add more complex gesture recognition, to allow users to point or otherwise identify areas of interest. This could be used to add additional functionality for teaching or outreach activities. This project would look at extending an existing AR Sandbox implementation to add gesture recognition and using it to implement some demonstration applications.

16. **Robot SLAM**

Steven Mills (steven@cs.otago.ac.nz)
Brendan McCane

Simultaneous Localisation and Mapping (SLAM) is a process where a robot can build up a model of its environment (mapping) while it learns to navigate through it (localisation). There are a number of good techniques for SLAM, and the purpose of this project is to get them working to navigate a robot around the graphics lab and surrounding corridor. Ideally the robot would be able to make its way to different locations (such as "Steve's Office" or "The Tea Room") without bumping into people or things along the way.

17. **Recognising 3D Shapes**

Steven Mills (steven@cs.otago.ac.nz)
Brendan McCane

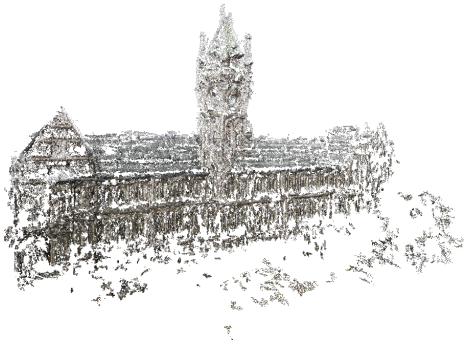

Deep Learning has been very effective in image-based recognition, but is often limited to images. Alongside these advances, it has become much easier to construct 3D models, which are typically clouds of points, such as the one in the picture. Recently a technique called PointNet has emerged, which can recognise objects on the basis of 3D point clouds. This project will investigate PointNet's capabilities, and seek to apply them to specific shape recognition tasks.

18. **Stitching 3D Models**

Steven Mills (steven@cs.otago.ac.nz)

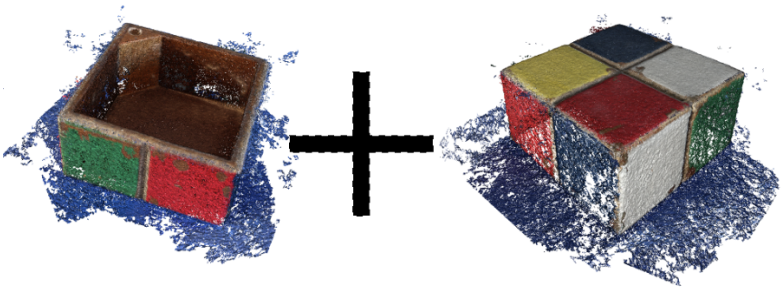

It has become quite easy to create 3D models from images, but often it is physically impractical to model a whole object at once. What we end up with are models of different parts of an object which we need to join together. There are a number of problems to be investigated here - how to line the models up, how to combine different images to texture the model, or how to refine the model once it has been stitched together.

19. **Tracking for Large-Scale AR**

Steven Mills (steven@cs.otago.ac.nz)
Stefanie Zollmann

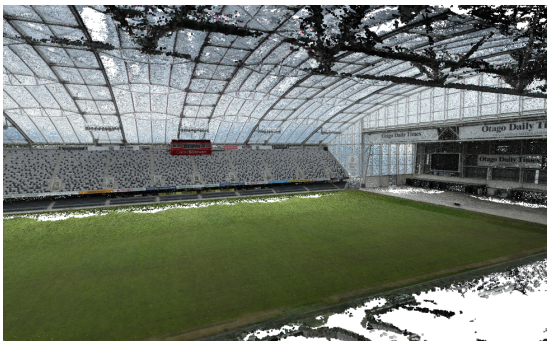

Augmented Reality (AR) often relies on camera-based tracking to align the real and virtual worlds. Most of these methods assume that the camera moves a lot relative to the scene, but we are interested in a situation where this is not the case – AR for spectators in large open environments such as at a sports game or other event in the stadium. This project will examine methods that can be used to determine where a spectator is in such a large space, and what they are looking at.

20. **Achieving centimetre-level positioning on smartphones using a combination of GNSS and MIMO**

Haibo Zhang (haibo@cs.otago.ac.nz)
Zhiyi Huang
Robert Odolinski (School of Surveying)

Many applications such as autonomous driving, UAV navigation and surveying require precise positioning at the centimetre level. This project aims to investigate the novel combination of multiple Global Navigation Satellite Systems (GNSSs) and the upcoming 5G networks, with large antenna arrays, so as to achieve such centimetre-level positioning accuracy. This GNSS+MIMO combination is essential as GNSS may become unavailable or inaccurate in urban areas. We expect to develop a prototype on smartphones to demonstrate the feasibility of this approach.

21. **Partially Observable Markov Decision Process for scheduling wireless communication**

Haibo Zhang (haibo@cs.otago.ac.nz)

Partially Observable Markov Decision Process (POMDP) is a generalisation of Markov decision process by permitting uncertainty on Markov states. It is general enough to model a variety of real-world sequential decision processes including robot navigation and automated planning. This project aims to apply it to solve a wireless communication scheduling problem. What are expected to be done include:

- Re-implement the incremental pruning approach for POMDP based on an existing source package that is complex and unfriendly to use.
- Use the new implementation to solve a wireless scheduling problem.

22. **Simulator for Li-Fi: data communications through LED light bulbs**

    Yawen Chen (yawen@cs.otago.ac.nz)

    LED bulbs can provide more than just illumination! Recent development of "Li-Fi" is an innovative way of wireless optical communications using the signals sent through the light bulb instead of radio waves as in Wi-Fi.The lightbulb is flicked on and off very quickly, up to billions of times per second. That flicker is so fast that the human eye cannot notice it. There are many advantages and potential applications to use Li-Fi in the future. One of the problems is how to design the layout of LED bulbs according to different application scenarios and requirements (e.g.illumination level/transmission rate/energy saving). This project aims to explore the characteristics of Li-Fi and design a simulator for Li-Fi, which is useful for the usage and development of Li-Fi.

23. **Simulator for Multicasting in Optical Network on Chip (ONoC)**

    Yawen Chen (yawen@cs.otago.ac.nz)

    With the increasing number of CPU cores on a processor chip, conventional electronic interconnects for core-to-core communication are becoming a bottleneck. Recent rapid advances have made the optical-based interconnection, Optical Network-on-Chip (ONoC), an attractive solution to breakthrough electronic interconnect limitations. Multicast communication, in which packets from one source need to be sent simultaneously to multiple destinations, intensively exists in chip multiprocessors due to the need for cooperative computing and cache coherence. This project aims to simulate the behaviour of multicast traffic on Optical Network on Chip (ONoC). The simulator will be very useful for the hardware design and development.

24. **Cacophany project**

    Brendan McCane (mccane@cs.otago.ac.nz)

    The Cacophany project is an open project that is being setup to protect NZ birds by eliminating pests in a smart way. Part of the project involves identifying pests from images and thermal cameras. This project will contribute to the cacophany project in some way. The first task will be to determine what the best form of contribution will be, but something in animal recognition and machine learning is likely.

25. **EEG Recognition**

    Brendan McCane (mccane@cs.otago.ac.nz)

    EEG measures brain activity in the human cortex and can apparently be used to diagnose mental health states such as anxiety and depression. This project will involve implementing machine learning methods for estimating depression or anxiety from EEG signals.

26. **Robot Hands/Grippers**

    Brendan McCane (mccane@cs.otago.ac.nz)

    Analyse, select and build an open source robot hand/gripper with appropriate software to drive the thing. Some example repositories to source designs include: https://openbionics.org and https://www.eng.yale.edu/grablab/openhand.

27. **Teaching neural networks to write**

    Lech Szymanski (lechszym@cs.otago.ac.nz)

    Convolutional neural networks (CNNs) can outperform humans on visual tasks, such as object classification from images. However, these models often fail when presented with specially crafted adversarial examples that would not fool a human. What if, in addition to training the network to read, we were to teach it to write (by controlling a *virtual pen*). Would the CNN learn to *see* the writing more akin to the way we see it? For this project you will be developing a CNN-based agent and train it using reinforcement learning (RL) techniques to re-construct written text from images. This project will give you a background in machine learning (specifically CNNs and RL) and practical experience with Tensorflow.

28. **Unsupervised classification in convolutional neural networks using a game engine**

    Lech Szymanski (lechszym@cs.otago.ac.nz)

    "If intelligence was a cake, unsupervised learning would be the cake, supervised learning would be the icing on the cake, and reinforcement learning would be the cherry on the cake. We know how to make the icing and the cherry, but we don't know how to make the cake." – Yan LeCun. In this project, instead of teaching a CNN to recognise objects from static images, the student will attempt to train the model using simple 3D scenes with geometric shapes. Instead of the wealth of image data, the model will rely on a live scene with variable position/lighting of unlabelled objects. Instead of the labels, the training will operate on principle that the the object seen at given moment is most likely the object seen just previously. This project will give give you a background in machine learning (specifically CNNs and possibly RL) as well as practical experience with Tensorflow and working with a game engine.

29. **Fix Tartini**

    Geoff Wyvill (geoff@cs.otago.ac.nz)
    David Eyers

    Tartini is an open source application that helps practical musicians and music students. Its main function is to display the pitch and loudness of each note as a passage of music is played. Unlike an ordinary commercial tuner, it does this continuously so the musician can discover technical faults such as pitch drop at the end of notes in wind instruments and bowing errors in the strings. It records the sound so a complete passage can be analysed note by note afterwards and imperfections identified.

    Tartini was created between 2003 and 2008 as a PhD project but it has not been updated since and it no-longer runs in the latest operating systems, with parts of the code using libraries that are no longer available. The goal of this project is to get it running again properly on macOS at least, and possibly on Linux and Windows too. Note that we will be able to share access to old (quarantined) computers that are still running Tartini successfully.

    More information about Tartini can be found at http://tartini.net.

    An interest in music and physics would be helpful. As a COSC490 research project it needs to extend the analysis or investigate improvements to the user interface.

30. **Casual Capture of Stereo Panoramas using Mobile phones**

Stefanie Zollmann (stefanie.zollman@otago.ac.nz)

In this project, we will investigate the capture of stereo panoramas on mobile phones. The idea is to simplify the content creation for virtual reality viewing for normal users. While traditional 360 panorama photographs and 360 degree cameras can be viewed in virtual reality devices, they do not provide a sense of depth to the user. Traditional methods for reconstructing stereo panoramas on the other hand involve specialised hardware. In this project, we address this gap and investigate methods that allow to compute stereo panoramas with a mobile phone camera. This project will build on previous works of our group (in collaboration with Ass. Prof Jonathan Ventura from Cal Poly, San Luis) and will include mobile application development as well as computer vision, virtual reality and human-computer-interaction aspects as we will look into supporting the user during the capturing process.

31. **Computational 3D videography**

Stefanie Zollmann (stefanie.zollman@otago.ac.nz)

The amount of video footage that is captured every day is growing at a rapid rate. However, most of this video data is captured mainly in 2D and is therefore missing a sense of immersion when being replayed. In this project, we look into the options of how to compute 3D information from 2D video footage. Ultimately, these methods would allow us to experience a video captured in 2D in a VR headset. We will build on our previous work and optimise the results with regards to frame-to-frame coherence and dynamic elements in the video data. This project will include aspects from real-time computer vision and computer graphics.

32. **4D video player**

Stefanie Zollmann (stefanie.zollman@otago.ac.nz)

The replay of 4D videos involves larger amounts of data compared to traditional 2D videos. In this project, we will investigate methods for efficiently storing and replaying 4D videos that represent a 3D scene for each video frame. We will investigate if methods used for traditional video processing are suitable for this kind of data. This project will include aspects from video processing techniques, computer graphics and computer vision.

33. **Evolutionary landscapes of antibiotic resistance***

Alex Gavryushkin (alex@biods.org)

This project is aimed at developing mathematical and computational methods to analyse mutational pathways in bacteria responsible for the development of antibiotic resistance. We will specifically focus on studying the combinatorially complicated structure of interactions within and between mutational pathways/ genes. The goal of the project will be to implement a computational approach to analyse empirical fitness landscapes. No biology background is necessary.

34. **Online algorithms in computational biology***

Alex Gavryushkin (alex@biods.org)

This project will focus on certain classes of computational inference methods, within the domains of phylogenomics and fitness landscapes, with the goal to design their scalable online versions. The goal of the project will be to investigate the online version of several simple algorithms, central to computational biology. No biology background is necessary.

*This projects would be ideal for someone considering postgraduate study. Please contact Alex (alex@biods.org) if you have any questions.