400-Level Project Proposals FYDNI 2022

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 submit your project selections via BlackBoard* by Friday 4th 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. 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.

*On BlackBoard select the COSC480_490_FYDNI_2022 paper, and then Project/Research Project/Assessment/Project Choice.

1. **Casual Collaboration**

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

   A 3D, collaborative sculpture application (Sculpture) has featured at many of the recent University Science Expos. After consulting a brief help-screen (and possibly not even doing that!) users of the provided computer terminals very rapidly begin to build (and sometimes destroy) a 3D sculpture collaboratively. The overall success of this exhibit stems in part from its ease of configuration and its usability.

   This project seeks to abstract and redesign the application platform underlying Sculpture, to help modernise the hosting of the application itself, and to support other similar collaborative systems to be built and demonstrated in future.

   Today's operating systems, web frameworks, game engines, and other forms of middleware provide tools for building cooperating software that were not available when Sculpture was first implemented. Further, high quality open source libraries are available that implement conflict-free replicated data types (CRDTs). CRDTs enable asynchronous collaborative editing such as within Google Docs. They specifically address how to avoid inconsistencies developing between distributed replicas of an artefact, as asynchronous editing occurs. The artefact in the case of Sculpture is the copy of the sculpture contained within the memory of each participating, connected computer.

   For a 490 project you would build a new application on top of the collaborative framework, or continue extension and evaluation of the framework itself.

2. **Tartini Reloaded**

   Geoff Wyvill (geoff@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 and responsively so the musician can discover technical faults such as pitch drop at the end of notes for wind players and poor bow changes for string players. 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, and continues to run on old (quarantined) computers within the Department. Geoff and David have demonstrated Tartini at the University's Science Expo, including during the event last year.

The goal of this project is to rework Tartini to run on a modern software platform. One possible target is a recent macOS version. A more ambitious target would be porting Tartini to the web. It is likely that JavaScript within a web browser now runs sufficiently quickly to handle Tartini's digital signal processing code.

More information about Tartini can be found here.

An interest in music and physics would be helpful. For a 490 research project you would also extend the analysis or simplify the user interface.

3. **Emotion Aware Dashboard**

   Veronica Liesaputra (veronica.liesaputra@otago.ac.nz)
   Claudia Ott

   Multiple studies have shown that students' emotional states influence their general well-being and their learning success. With the advances of AI and sensors, there are many techniques we can use to automatically detect human emotions through various modalities such as facial expressions, tone of voice, text, postures, etc. However, there have been limited studies done on how to effectively present the results of these analyses so that dashboards can create better learning experiences and performances for the students. In this project, you will create an emotion aware dashboard prototype and evaluate the effectiveness of your prototype at promoting changes in students' learning processes.

4. **Salient Event Detection**

   Veronica Liesaputra (veronica.liesaputra@otago.ac.nz)
   Andrew Trotman

   The ability to accurately identify salient information in a given piece of text is an important problem in natural language understanding area. Although there have many attempts at identifying salient sentences or entities in a text, most models are trained on news articles, which are short and present all important information upfront. As a result, such models are biased and do not work well for summarizing long narratives. In this project, you will be creating a classifier to recognise salient events in a long piece of text. Modelling approaches can be deep learning or classical machine learning, such as Support Vector Machine.

5. **Image analysis on food material**

   Veronica Liesaputra (veronica.liesaputra@otago.ac.nz)
   Zhiyi Huang
   Indrawati Oey

   This Honours project is in collaboration with Department of Food Science to develop methods of processing high resolution images and then tailor the method to extract useful information from the images captured from food matrices. One of the food matrix of interest is bread made from legume flour and then examining the images of crumb bread. The outcome of the project will help to draw correlation to the baking procedure raw ingredient composition on the quality of bread.

6. **Modelling with Lines and Planes**

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

   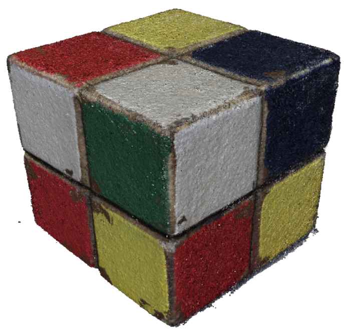

   Most techniques for modelling 3D objects from images create point clouds - thousands or millions of 3D points which lie on the surface of the object and which can then be connected to form a triangulated mesh. Many objects of interest, particularly buildings, however, are better modelled using planes or lines rather than points. These models are much simpler and more semantically meaningful. The cube in the picture, for example, is modelled with nearly 2 million 3D points, but could be well represented by a few planes - one for each face, and one for the surface below it. This project will look at methods for either converting point cloud models to simpler planar models, or for direct 3D modelling from lines or planes detected in images.

7. **Gestural Interfaces**

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

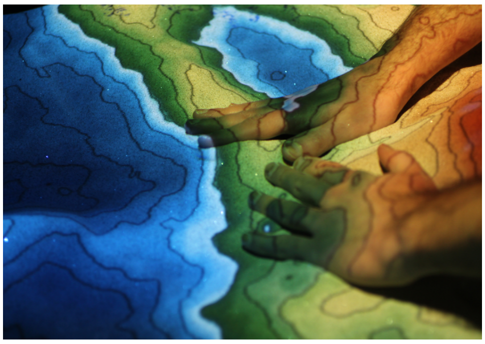

   In some applications it is not convenient for people to use hardware devices to interact with computers. Gestures are one solution to this problem - by tracking people's hands or even whole bodies we can interpret gestures to control computers. There are a couple of possible applications for this. One is our augmented reality sandbox where we would like to support pointing gestures to activate visualisations. Another would be an interactive display for the departmental foyer, where visitors could use gestures to select items of interest, swipe through information pages, etc. This project would involve a mix of gesture recognition and visualisation to create an interactive experience suitable for casual use.

8. **3D Labelling Tool**

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

   One of the big challenges in deep learning for computer vision is getting labelled training data. In the example image, 3D scans of a person have been coloured to indicate different body parts. Labelling data is a particular challenge for 3D vision, since there are no good tools for manual annotation of 3D data sets. In this project you will create a tool for this task in VR – the user would be able to move the model around, zoom in and out, and 'paint' labels with their hands. As well as supporting manual labelling, the tool could learn to predict labels as it goes, and update its predictions as you correct any mistakes.

9. **Making Masking Easier**

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

   

   We have very good software for making 3D models from images, but if we want to model a single object we need to isolate it from the background. Even with plain backgrounds, automating this is difficult, so the best models still require manual masking of images. This is tedious and can take quite a while. In this project you will investigate online learning where the system uses information from previous masks to make it easier to draw the next one. An alternative to the learning approach would be to make a capture setup that makes masking easier. One option would be to have the object sit on a flat screen, and change the display to make the masking better and/or make the 3D reconstruction pipeline more reliable.

10. **Compiled search engine indexes**

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

    A search engine usually uses an external index so that it can quickly determine which documents contain which search terms. In this project we are going to compile the index. That is, to turn the index into a program. That program will be able answer all the same queries that can be answered with an external index, but rather than process a data file, it will execute the index directly.

11. **Does stemming really work?**

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

    Search engines such as Google and Bing "stem" your search terms. That is, if you enter "running" it will also search for "run" and "ran". But does this really produce better results? There are several stemming algorithms already published, and some authors have implied that stemming can be learned using a deep learning. But when do stemmers work and why do stemmers work? In this project we will build a dictionary stemmer by scraping the Wictionary, and seeing which words are related and how. We will then use this stemmer to measure the performance of other stemmers – in particular as document collections grow in size and documents grow in lengths. Finally, we will try to learn a new stemmer using the data from the Wictionary.

12. **Real-time banjo music transcription**

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

    Some musical instruments, such as the Banjo, can either play chords or pure notes. Indeed, there is a style of Banjo (called Clawhammer) in which notes are typical and chords are rare. It is typical to play Old Time music in this style (Old Time is similar to Bluegrass). As the name suggests, these tunes are old, and they have been passed from generation to generation acoustically. In this project we will build a system that, in real time, will listen to a piece of music being played by a Banjoist and transcribe it to musical score (Tab) in real time. Ideally this will run on a mobile phone so that we can take it to skilled players and transcribe some tunes as they are being played.

13. **Messy handwriting**

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

    Tidy handwriting can make the difference between passing and failing an exam – if the examiner can't read your writing they cannot give you marks for it. Some people have very hard to read handwriting, while others is very clear. Can we use a computer to automatically identify those with messy handwriting? If we can then we can work with them to improve it (which is outside the scope of this project). In this project we will create a data set of tidy and messy writing, then we will train a computer to measure the messiness. This system might be deep-learning, or it might not. We simply do not know – that is part of the project.

14. **Finding EEG biomarkers for depression**

    ZhiyiHuang (hzy@cs.otago.ac.nz)
    Veronica Liesaputra

    Depression is a global problem, especially under the current pandemic. Brain wave (EEG) is the most convenient way to investigate the status of the brain. It has rich information about the brain's mental status. In this project, we are going to use deep neural networks to classify the EEG data collected from depressed patients. It can help find biomarkers for depression. We have a PhD student who can assist with the machine learning algorithms. We will use public datasets in the project as a start. Part of the project involves preprocessing of EEG data with various tools available in the field.

15. **Visualising EEG data with music and nature**

    Zhiyi Huang (hzy@cs.otago.ac.nz)
    Veronica Liesaputra

    Research have shown that people's productivity is affected by their mental status. People's overall well-being can be improved if they are aware of their emotional well-being. EEG data contains a wealth of information about the brain's state of mind, such as excitement and calmness. To help people understand their mental status, in this project, we will investigate possible ways to visualise EEG data. This could be by transforming EEG information to some music and nature scenes, such as flowing water or beautiful flowers.

16. **Tracking teeth brushing using IMU and CNN**

    Haibo Zhang (haibo.zhang@otago.ac.nz)
    Zhiyi Huang
    Peter Mei (Dental School)

    Many people, especially kids, brush teeth improperly, leading to a rising number in tooth decay and periodontal disease. We have developed a prototype using a motion sensor attached on wrist to track teeth brushing. A method based on Convolutional Neural Networks (CNN) has been designed to classify brushing postures such as side-to-side and up-and-down brushes, and achieves more than 90% on classification accuracy. In this project, we aim to further increase the classification accuracy by combining CNN with Long short-term memory (LSTM). We also aim to achieve real-time brushing recognition based on the data stream generated by the motion sensor without pre-segmentation.

17. **Develop an energy model for Quadcopter**

    Haibo Zhang (haibo.zhang@otago.ac.nz)
    Zhiyi Huang

    Unmanned Aerial Vehicle (UAV) swarm systems have enabled many long-haul, delay sensitive and hazardous applications. As UAVs are powered by batteries and the airspace they work in is usually highly dynamic, planning trajectories in an energy-efficient way to avoid both UAV-to-UAV and UAV- to-obstacle collisions has become a critical and open problem. We have designed path planning methods for UAV swarms to avoid obstacles. In this project, we plan to carry out experiments using the S500 Quadcopter toolkit to build a realistic energy model to quantify the energy consumption of different collision-avoidance actions.

18. **Scheduling wireless transmissions using federated learning**

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

    A wireless channel is a shared channel. If multiple devices send data to the same receiving device over the same channel simultaneously, all transmitted data are typically corrupted and discarded due to collisions. Time Division Multiple Access (TDMA) is well-known medium access control protocol to schedule different packet transmissions in different time slots to avoid collisions. However, constructing the optimal transmission schedule to achieve high transmission reliability is a challenging problem since the channel quality sensed by each transmitting device can change over time. This project aims to investigate the feasibility of using federated learning to schedule wireless transmissions. It is expected that federated learning could enable multiple devices to collaboratively learn a shared prediction model of the wireless channel without the need of exchanging training samples among different devices.

19. **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. Such spaces have some interesting properties – and some serious problems. This project will involve implementing and exploring a Hopfield type network to further develop our understanding of its dynamic behaviour.

20. **Neural network models of memory**

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

    Artificial neural networks have been used to model many aspects of human memory. This project will involve reviewing the existing literature, and implementing, testing, and possibly extending an existing model or models. The review may focus on the representation of conceptual categories and the process of categorisation.

21. **Visualising deep networks**

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

    Deep networks are powerful – but mysterious. This project will involve reviewing current methods for visualising the behaviour of deep networks, implementing an existing method or methods, and possibly developing a new method or methods. This project will involve the use of a standard deep learning toolkit such as TensorFlow.

22. **Deep trouble**

Lech Szymanski (lech.szymanski@otago.ac.nz)
David Eyers

Transfer learning facilitates convenient reuse of machine learning models. Rather than training from scratch, in the context of reuse, partial retraining is performed over an existing source model. For example, training can be applied over a neural network with weights copied from the source network that remain fixed but for within a few final layers. Is there, from a security perspective, a potential risk of original training process embedding malicious data? Consider an unrealistically oversimplified, hypothetical scenario, in which a target neural network is trained, by transfer learning from a source neural network, to recognise faces for the purpose of controlling access to a building. Is it possible to embedded a pattern in the source capable of 'surviving" target training and later causing a strong response that may act as a "master key"?

This project will involve devising means of embedding source networks with "malicious" patterns. The aim is to determine the extent to which it is possible to have these patterns survive the transfer learning and whether they can affect a target network in any meaningful way.

23. **Deep recollections**

Lech Szymanski (lech.szymanski@otago.ac.nz)

Artificial neural networks and deep learning works, surprising well sometimes. The key aspects of this success are: availability of massive amounts of training data, unreasonably large and very deep network models (convolutional, if we are dealing with images), and optimisation methods that force the network to overtrain without sacrificing its generalisation capabilities. This is surprising because the latter of these two aspects are precisely what machine learning theory says not to do if you want a model to extract *useful* patterns from data. So, does a neural network learn *useful* patterns, or does it just memorise its input?

The aim of this project is to investigate the degree of memorisation in neural networks. For example, if we repurpose a network pre-trained for classification to do auto-encoding, how much, and what type of information on the input gets passed through the network? Will this be a viable measure of memorisation? If so, can we use it to track memorisation over different layers of the network? Can get a glimpse of the patterns the network pays attention to?

24. **Learning to play Bridge with deep reinforcement learning**

    Lech Szymanski (lech.szymanski@otago.ac.nz)
    Michael Albert

    An agent trained to play against itself using deep reinforcement learning techniques is capable of mastering games like Chess and Go without any prior knowledge about the rules of these games. Let's try to train an agent to teach itself to play the card game of Bridge using the AlphaZero framework (if possible). Bridge is different (from Chess or Go) in that it is played in pairs (for a total of four players), with agent communication during the *contract bidding* phase of the game and the history of what happened up to the current point being of vital importance. The challenge of this project will be in devising neural network models and amending the AlphaZero framework to allow it to play and learn how to play Bridge.

25. **Systolic Array Simulator**

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

    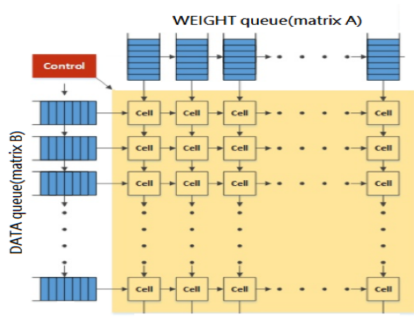

    https://github.com/leo47007/TPU-Tensor-Processing-Unit

    Systolic Arrays are one of the most popular techniques for Deep Learning accelerators today, as they provide extremely high efficiency for running dense matrix multiplications. Systolic design methodology maps an N-dimensional dependence graph to a lower-dimensional systolic architecture using a transformation. However, the research community lacks tools to provide insights on both the design trade-offs and efficient mapping strategies for systolic-array based accelerators. This project is to simulate the mapping strategies for systolic array, which can include some features for future comprehensive design space exploration.

26. **Simulation of Photonic tensor cores for machine learning**

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

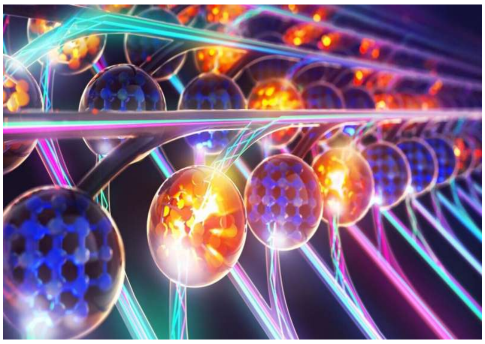

https://scitechdaily.com/ai-boosted-by-parallel-convolutional-light-based-processors/

In a study published in Nature (J. Feldmann, N. Youngblood, et al., *Parallel convolution processing using an integrated photonic tensor core*, Nature, 07 January 2021), scientists developed a hardware accelerator, photonic tensor core, which are used for machine-learning algorithms. Since different light wavelengths (colors) don't interfere with each other, the researchers could use multiple wavelengths of light for parallel calculations. To do this, they used another innovative technology, developed at EPFL, a chip-based "frequency comb", as a light source. This project aims to investigate and simulate some features for this new photonic tensor core.