400-Level Project Proposals 2016

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, Steven Mills steven@cs.otago.ac.nz 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. If you find multiple projects that look equally appealing, then you can give lists of projects as first, second, and third choices. You will still need to include at least three different supervisors, and at least two in your first two lists.

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 in the project description.

1. **Crossroads, Cars and People**

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

   What is the best way to run traffic lights? In George Street we have a busy area used by cars, buses, pedestrians and delivery vehicles. Changes over the years have not reduced delays or improved safety. We do not know how any particular strategy will affect the junctions until it is tried. This takes months, law changes, expensive sign replacement and publicity.

   Are there simulators that allow us to compare strategies? How to we find them, use them, and verify them? Should we write our own?

   Analyse the problem and make some useful progress.

   More information can be found at: `http://www.cs.otago.ac.nz/graphics/Geoff/Traffic.html`

2. **Having fun with programmable hardware.**

   Richard O'Keefe (ok@cs.otago.ac.nz)
   David Eyers (dme@cs.otago.ac.nz)

   A Field Programmable Gate Array is reprogrammable hardware. You can write a hardware design in a high level language (the two main ones are VHDL and Verilog) and compile it, down load it into an FPGA, and run it. VHDL and Verilog can also be compiled to silicon, to make application specific integrated circuits, so FPGAs are often used to prototype things. One can design and implement ones own CPU, but it is more useful to implement an algorithm directly without any instruction interpretation overhead.

   We have a couple of FPGA boards and some textbooks and free software.

   It would be particularly interesting to see if these things can be used in information retrieval. For example, the inventor of the widely used Porter stemming algorithm has devised a special purpose programming language for writing stemmers; is it possible to translate stemmers into VHDL so they can be compiled into hardware? Even more basically, Unicode is enormously complicated, even a simple thing like asking "are these two strings equal" is hard. Can we implement that efficiently on an FPGA?

   If you have something of your own you'd like to try, that would be welcome too.

3. **Latent Methods**

   Richard O'Keefe (ok@cs.otago.ac.nz)

   I've recently come up with an idea called "latent methods". These are methods that you define in some class, but they do not become available until a subclass implements the interface(s) they are declared to need. This project would involve

   (a) critiquing the idea and comparing it with other approaches like multiple inheritance,

   (b) implementing latent methods in some OO programming language, possibly by modifying a compiler or writing a preprocessor, and

   (c) writing test cases for your implementation that can be used as examples to show the usefulness (or otherwise!) of the idea.

4. **Semantic Relationships in Source Code**

   Richard O'Keefe (ok@cs.otago.ac.nz)

   Some programming languages, notably C#, Java, and Smalltalk, have a standard way to attach "annotations" to code. This is usually used to provide the compiler and run-time system with more information like "this reference should not be null", "this is a test case", or "this is part of an Object-Relational Mapping". The nearest we tend to get to something for the programmer's edification is "don't use this any more".

   But chunks of source code are related to each other and to external documents in ways that the compiler and run time system do not need to know about, while maintenance programmers would love to know. The idea is to take an existing body of code that you are not familiar with, read it, and note what kinds of links are either informally specified or entirely missing but important to you, so that we can build up this catalogue. You should also review what has been done in Java and C# and Smalltalk. For more work, you might construct an actual set of annotation definitions that Java (or C# or Smalltalk) programmers might use and provide some tools (perhaps by extending an existing IDE) to exploit them.

5. **Explorations in Rosetta Code**

   Richard O'Keefe (ok@cs.otago.ac.nz)

   www.rosettacode.org is a growing collection of problems (currently 780) and solutions to them in hundreds of programming languages. It has been suggested that this is a useful resource for Software Engineering analyses. Since the site was not developed for that purpose, it is not surprising that there are some questions. For example, the size metric used to compare solutions is very weak; for C it can vary by 28% depending on the options you give to a code reformatter. Is it similarly weak for other languages? It is known that some problems are interpreted more than one way, making solutions incomparable; how common an issue is this? The 8 languages compared do not have solutions to exactly the same problems; are the conclusions of previous work still sound if we add more languages or more solutions?

   This is a fairly open-ended introduction to software engineering research and code mining that should lead to a publication.

6. **Fun with Parallela**

   Richard O'Keefe (ok@cs.otago.ac.nz)

   I have a Parallela board. This is a pocket sized board with 18 cores: 2 ARM cores running Linux and 16 RISC cores, using less than 5 Watts and costing about $150. Programming this beast isn't like programming a conventional multicore CPU (although you can use OpenMP) or a GPU (although you can use OpenCL) because the 16 RISC cores have small memories with fast interconnects. The point of the project is to develop working software that does something interesting in parallel and to measure performance relative to a multicore CPU or a GPU.

   "Something interesting" might be some sort of combinatorial optimisation algorithm (other than travelling salesman), a genetic programming package, or a data mining task. If you have another idea that would be wonderful.

   If your code is clear and well described, the company that make the board might be happy to display it on their site.

7. **We can Know Everything that Happens in our House in Real Time, Remotely**

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

   Have you every worried about your house when you were away, because you were uncertain whether you have turned off the oven and have locked the door? Have you ever worried about your grandparents who live alone? Through the combination of wireless sensor networking and smartphone technologies, you can easily monitor you house remotely via your smartphone, e.g. detect fire and falls of the elderly. In last year, we have developed a prototype for such a system. In this year we are going to develop a smart client program for smartphones and a protocol for transferring photos from sensor to smartphone. This project can host two students, with each working on one function. The first one requires experience with Android programming.

8. **Tracking the Motion of the Human Body using 3D Inertial Measurement Units**

   Haibo Zhang (haibo@cs.otago.ac.nz)
   Some of: David Eyers, Zhiyi Huang, Lech Szymanski, Steven Mills

   Human motion tracking has a wide range of applications such as in sports, gaming, and medical rehabilitation. In this project we are going to develop a motion capture system using wireless sensors and 3D inertial measurement units. Please watch the following YouTube videos to get a feeling of such systems.

   - `https://www.youtube.com/watch?v=nRzS39UtTF4`
   - `https://www.youtube.com/watch?v=b9TQcv0Asys`
   - `https://www.youtube.com/watch?v=tyiuocPvUes`

   We have already designed the hardware, implemented the communication protocol, and carried out some preliminary work on animating the captured body motions using Blender. In this project we will add more functions and produce a demo for real-time body motion tracking. Background and experience with computer graphics and Python programming would be helpful to work on this project.

9. **Learning Human Contact Patterns from Bluetooth Data Traces**

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

   Smart phones are becoming popular. Bluetooth technology allows two smart phones near each other to exchange messages independently. This project aims to learn the human contact patterns by analyzing the data communication traces collected from a set of Bluetooth devices. These patterns would be very useful to understand human activities and design efficient communication protocols for smart phone networks.

10. **Convolutional Neural Network for Classification of EEG**

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

    The signals in the Electroencephalography (EEG) data have a spatial and temporal structure. In other words, you can interpret the data from the electrodes as a set of image with changes in pixel intensities over time. Given the success of convolutional neural networks in image classification, it is worth trying to apply these machine learning model to classification of EEG data in order to identify disorders such as anxiety or depression. This project will be an attempt to build a convolutional neural networks to find patterns in EEG video data.

11. **Geometric Algebra for Camera Position Estimation from Images**

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

    Given two images of the scene viewed from different angles, it is possible to compute the relative position of the camera that took the photos. This problem has been solved and it involves a bit of matrix computation based on set of corresponding features identified in the two images. In this project we want to use geometric algebra, an alternate mathematical framework for computational geometry, to solve the same problem. The objective is to investigate whether the solution, stemming from different mathematical paradigm, offers any advantages/disadvantages over the established ones.

12. **Game Engine for Object Recognition**

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

    The purpose of this project is to hook up a game engine to a machine learning system for the purpose of training it to recognise objects in images. The state of the art deep learning systems require training on millions of still images in order to generalise over the changes in object pose and lighting. Learning from a virtual scene should simplify this, and provide opportunities to manipulate objects in the scene at will. This project will involve an investigation of suitable game engines, creation of building environment in C++ to link against game engine libraries, and familiarisation with its API for the purpose of interfacing to machine learning system.

13. **Reconstruction of Stone Tools**

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

    With colleagues in the Department of Anthropology and Archaeology I am interested in building 3D models of stone tools, such as adzes. The archaeologists are interested in how these tools were made, and the shape of the indentations on the surface of the stone are an important clue as to this process. While we can make visually appealing 3D models of the tools, they lack the fine detail needed for this application. This project will investigate multi-view stereo techniques to reconstruct and analyse these surfaces.

14. **Making 3D Line Drawings from Images of Buildings**

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

    There are established techniques for creating 3D models from collections of images. These are usually based on point features in the images, which leads to a representation in terms of 3D point clouds. Scenes of and inside buildings, however, are often rich in line features. This project will look at existing techniques for using line features in 3D reconstruction, implement some of them, and work towards making 3D line sketches from images of buildings.

15. **Embodied Particle Effects**

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

    Recently I have been involved in an art installation based around the idea of being struck by lightning. The user is tracked with a Kinect sensor, and their image displayed on a screen. After a lightning strike, sparks are drawn over their image, spreading out from the point of impact.

    At the moment, these sparks are a simple 2D particle effect. The aim of this project is to make a 3D particle effect where sparks use information (from the Kinect) about the shape of the person to move around the body in a realistic fashion.

16. **Planetarium Dome Display**

    Steven Mills (steven@cs.otago.ac.nz)
    Some combination of David Eyers, Geoff Wyvill, and Oana Jones (Otago Museum)

    The Otago Museum has recently installed a planetarium, which provides a 6m wide dome display for around 50 people. As well as two 4K projectors, a Kinect v2 sensor has been installed looking at the audience. There are any number of projects that could be done with the dome, including content generation, interactive experiences using the Kinect sensor, data visualisation, and so on.

17. **Player Tracking for Team Sports**

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

    Along with Chris Button in the School of Physical Education, Sport, and Exercise Sciences, I am interested in tracking players during team sports. This is a challenging tracking task since players often appear similar, and they interact in a complex and confusing manner. This project would begin by implementing and comparing a number of standard tracking algorithms, and then look at ways to improve their performance in this application.

18. **Modelling Fashion**

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

    Many museums have collections of garments from many historical periods. Due to different fashions and styles, many of these do not fit on standard mannequins. It is possible to manufacture mannequins of a specific size, but this is expensive. The aim of this project is to model how the dimensions of a mannequin affect the fit of particular garments, and given a collection of garments, and measurements taken from them, determine the optimal set of mannequins so that the collection can be displayed.

19. **Deep Neural Networks for Fish Recognition**

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

    Recognising fish from images is useful both for the fisher-person who wants to identify what they have caught, and the environmentalist or marine scientist who wants to survey the number of different species in a body of water. This project will look at using deep neural networks for recognising fish, possibly with a view to making a mobile app.

20. **EEG, Raspberry Pi, Drone, and Machine Learning**

    Zhiyi Huang (zhuang@cs.otago.ac.nz)
    Brendan McCane (mccane@cs.otago.ac.nz)

    Electroencephalography (EEG) is a recording of electrical activity from human brain. It contains a lot of interesting and important bio-information that could be used for Brain-Computer Interface (BCI). We have implemented a simple BCI based on a single channel EEG that can control a drone through Raspberry Pi. In this project, we plan to use a 8-channel EEG device to control a drone with Raspberry Pi. This more sophisticated BCI system will involve machine learning algorithms which need to be implemented on Raspberry Pi.

21. **Accelerated Path Tracing**

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

    Path tracing is a statistical method for rendering computer graphics. Compared to ray tracing, path tracing is inherently incremental: samples are integrated into the result until a time or iteration threshold is reached. Although computationally expensive, path tracing is able to handle many complex optical effects using a simple model.

    This project will involve building a simple CPU-based path tracing algorithm, and then comparing its performance to a GPGPU or distributed implementation.

22. **Building a Scalable Distributed Entropy Infrastructure**

    David Eyers (dme@cs.otago.ac.nz)
    Student will also collaborate with Jim Cheetham and Paul Campbell as appropriate

    Software used in computer security often requires a source of randomness, e.g. for the generation of cryptographic keys. While modern CPUs include components that sample their physical environment to derive random numbers, the recent disclosures about US NSA and other similar organisations' interference have thrown doubt over the trustworthiness of these components.

    Jim Cheetham and Paul Campbell developed the OneRNG `http://www.onerng.info` as a USB-connected source of entropy that can be verified, both in terms of hardware and software.

    This project is to design and develop a network-based, scalable random number generator. This will be useful for cases such as cloud-hosted servers, into which you cannot plug your own hardware random number generators, but to which you are able to deliver a network stream.

When higher rates of random data are required than can be safely provided, just plug in more OneRNGs, and the system should automatically reconfigure itself to use them. Note the related projects Pollen and Pollenate from Ubuntu: `http://bazaar.launchpad.net/~kirkland/pollen/trunk/view/head:/README`

23. **QR Codes for Security** David Eyers (dme@cs.otago.ac.nz)

    QR codes are a class of 2D barcodes of various sizes and error correction capabilities that have a public, extensive and extensible schema for data that can be encoded in them. Open source software is available to generate and detect QR codes.

    This project focuses on the development of software that uses QR codes for computer security functions. Forming a line-of-sight coupling between the QR code display (possibly even just on paper) and a camera has the advantage that so called "man in the middle" attacks are difficult compared to communication over a data network. This increases the convenience of operations such as delegation of security privileges.

    A previous project student successfully implemented the one-way transport of Kerberos tickets through QR codes. A promising alternative target is the OAuth2 protocol for decentralised authorisation, and/or two-way QR code interactions.

24. **Web pipes**

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

    Pipes are frequently used within shell scripting on Unix-like operating systems. They are a simple, yet highly effective inter-process communication paradigm. Shells create anonymous pipes between process invocations sequenced using the '|' character. Named pipes allow pipe stages to be connected together at a time independent from the time that they are created: the pipe endpoints sit in the filesystem. Pipes can span multiple computers, by including invocations of tools such as the Secure Shell (SSH), netcat, or wget within pipelines.

    This project integrates named pipes into web browsers. This goes beyond simply using HTTP(S) for the data transport of a pipe: wget and friends can already do this. The browser integration is to allow pipe creation, configuration, and local and remote pipe linking to be done with user involvement, and with the involvement of the browser's JavaScript engine.

25. **Secure Configuration for the Internet of Things**

    Matt McKague (mckaguem@cs.otago.ac.nz)
    Dave Eyers (dme@cs.otago.ac.nz)

    A big problem for the IoT is configuration. How can your grandma get her wifi-enabled light bulb configured? How can a contractor configure 1000 wifi light bulbs? How can you do this securely? This project aims to implement a simple and secure way to do IoT configuration using an out-of-band communication channel: a smartphone's camera. There will be lots of different components to this project, and it will have opportunities to try out cloud computing, computer vision, mobile applications, and embedded systems. There will also be an opportunity to do some research in cryptography and network security.

26. **Smarter Power Meters**

    Matt McKague (mckaguem@cs.otago.ac.nz)
    Dave Eyers (dme@cs.otago.ac.nz)

    Smart meters are not very smart: all they do is report data. They make no decisions. What if your power meter could respond to real-time electricity prices by turning on and off appliances in your house to get you the best bang for the buck? This project will look at simulating a neighbourhood equipped with such devices to answer questions like: -can changes in the price trigger undesirable swings in electricity usage? -can the system introduce oscillations in the system? If so, are there simple fixes? -what savings could there be for utilities (load balancing and increased efficiency) and consumers (lower power bill)?

27. **Playing Card Ciphers**

    Matt McKague (mckaguem@cs.otago.ac.nz)
    Dave Eyers (dme@cs.otago.ac.nz)

    Modern cryptography is mostly limited to computers because of the complexity of computation involved. Some ciphers, however, are designed to be used by humans alone (e.g. `https://www.schneier.com/cryptography/solitaire/`) but the tradeoff is often reduced security and biases. This project aims to look at some playing card ciphers and analyze their workings to identify the mechanism producing biases and hopefully modify the algorithms to reduce the bias.

28. **Bone Segmentation in MRI and CT**

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

    Magnetic Resonance Images (MRI) and Computed Tomography (CT) images are common forms of medical imaging used to help diagnose several disorders. However, discovering the pose and shape of bony structures is not well supported in MRI images. We have a set of MRI and CT images of the lumbar spine and several shape models of individual vertebra. This project will investigate and implement methods of automatically finding the pose and shape of the lumbar vertebra from MRI and/or CT images.

29. **Bird Recognition from Audio**

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

    Recognising the species of bird from images or audio is interesting and useful for birdwatchers and for conservation purposes. This project will look at automatically recognising bird species from audio signals using deep neural networks.

30. **Analysis of Beginner Programmer Errors**

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

    We have a database of lab and mastery test submissions for COMP150 from 2011-2015. This data is a potential gold mine of beginner programmer mistakes and could be used to remodel teaching in first programming courses. This project will involve curating and analysing this data, and attempt to extract out the most common problems made by beginning programmers.

31. **A Neural Network Model of Hierarchical Event Representations**

    Ali Knott (alik@cs.otago.ac.nz)

    The actions we experience in the world, and refer to in sentences, are hierarchically structured. For instance, consider the action 'Sue ate a sandwich'. Sue doesn't eat the sandwich in a single atomic action: she takes lots of separate bites, and has to chew and swallow each one. In this project, you will build a neural network model of how an agent learns to put together a sequence of low-level actions, to achieve the effect of executing a high-level action. The model will build on an existing neural network model of event representations in working memory and long-term memory. I already have an idea about how it will work, so this project is quite well defined. Its likely to lead to a publication.

32. **A Neural Network Model of Actions of Creation**

    Ali Knott (alik@cs.otago.ac.nz)

    Consider the action 'John made a cup of tea'. Its an 'action of creation', that produces something new, by assembling things together in a particular way. The object of this action is quite interesting: it only comes into existence as a consequence of the action. So what guides the action? In this project, you will build a neural network model of creation actions, that explores this question. The model will build on an existing neural network model of event representations in working memory and long-term memory. The main new idea element in the model will be a network medium that represents **goal objects**, that serve to guide creation actions. I already have an idea about how it will work, so this project is quite well defined. Its likely to lead to a publication.

33. **A neural network model of spatial navigation**

    Ali Knott (alik@cs.otago.ac.nz)

    Humans (and other animals) have some inbuilt neural machinery that allows them to find their way around the world. An important part of this machinery is a 'cognitive map', that represents the location of the agent in a coordinate system centred on its local environment. We know quite a lot about how this representation is generated: in fact last year's Nobel Prize in Physiology went to researchers studying this cognitive map. The aim of the current project is to implement a particular hypothesis about the structure of the cognitive map, using a network architecture called a recurrent Self-Organising Map (SOM). The hypothesis is novel, in that it uses exactly the same architecture as is used in a model of long-term memory for sequences of episodes. The implementation will test a simple version of the hypothesis: if it is borne out, the project is likely to lead to a publication.

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

    Yawen Chen (yawen@cs.otago.ac.nz)
    Haibo Zhang (haibo@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 lay—out 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. (http://www.lifi-centre.com/news/latest-news/)

35. **Simulator for Optical Network on Chip (ONoC): core-to-core communication by light**

    Yawen Chen (yawen@cs.otago.ac.nz)
    Haibo Zhang (haibo@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. This project aims to explore the characteristics of ONoC (number of cores, light communication path, topologies, and etc.), and design a simulator for Optical Network on Chip (ONoC). The simulator will be very useful for the hardware design and development. (http://arstechnica.com/information-technology/2015/05/ibm-demos-first-fully-integrated-monolithic-silicon-photonics-chip/).

36. **Music Projects**

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

    Computers can be used in a variety of ways to help musicians. Past music projects in this department have included:

    (a) Creation of a pitch analysis tool for musicians.
    (b) Survey of Music Notation editors.
    (c) A study of perception of vibrato.
    (d) An analysis of ragtime form.
    (e) A survey of automated composition methods with practical recommendations.

    Projects of this kind can only be handled by students who already have a background in performance and musical understanding. The ability to read music is a huge advantage here. If you have this interest and background e-mail me to arrange a meeting and we can discuss possible projects.