COSC480 Project Proposals for 2012

This is a list of suggested project proposals for COSC480 projects. If you have a project of your own you would like to undertake instead then you must discuss it with me or one of the staff. In the past we have often created special projects so that students can follow their interests. Additionally, many of the staff members have other interests that might be suitable for a project – you can find out about these by talking to them! In any case, suggesting a project that is not part of this list requires the agreement of a potential supervisor.

- Select (at least) three projects that you think will interest you.
- Meet with the faculty members who are offering those projects in order to discuss the possibility of taking them up (this is an **absolute** requirement – see below.)
- Rank three of them in preference order on the project request form and return it to the main office by **March 2**. Include both the project number and project title (for verification!)
- Send an email with your project preferences to me (malbert@cs.otago.ac.nz).
- You will be advised early in week 2 of your project assignment.

Michael Albert
400 level project coordinator

Meeting with possible supervisors is important. We try to give students their first choices but we have to spread the work among supervisors and match projects to your particular abilities. If you make a choice without seeing the supervisor, we will probably ignore that choice when doing the allocation.

1. **Modelling Fashion**
   Steven Mills (steven@cs.otago.ac.nz)
   Recent advances in computer vision mean that it is possible to construct models of objects from sets of images fairly reliably. One application of this is modelling objects for digital heritage. This project aims to develop tools for creating models of objects from structured sets of images. The initial application is the modelling of fashion garments for teaching and heritage applications. As well as building the models the project should consider usability and presentation issues so that non-experts can easily create and share models.

2. **Projection on to Moving Objects**
   Steven Mills (steven@cs.otago.ac.nz)
   Projectors are often used in performance to provide lighting and visual effects. If we add a camera to the mix, we can make these effects move with performers or other objects. The aim of this project is to develop a system using camera(s) and projector(s) to allow different patterns to be projected onto moving objects or people in as scene. The system should be robust to rapid motion and changing lighting conditions.

3. **Kinect Modelling**
Steven Mills (steven@cs.otago.ac.nz)
Microsoft's Kinect sensor combines a 2D colour camera with a 3D sensor based on structured light. With cheap hardware and a readily available SDK the Kinect has been applied to many areas beyond its initial gaming interface application. One that particularly interests me is in building 3D models from multiple images (each 'image' being a traditional colour image plus a depth map). The aim of this project is to design and implement a 3D scanner using the Kinect sensor as input.

4. **Transform Aware Feature Matching**
Steven Mills (steven@cs.otago.ac.nz)
Matching features (distinctive points) between images is a common task in many computer vision applications. Since the appearance of features generally changes with viewpoint and lighting conditions, methods that are invariant to these changes have become popular. In some cases, however, we have prior knowledge about how the features are transformed between images. The aim of this project is to develop a feature matching system that uses this prior knowledge to improve robustness to feature appearance changes.

5. **Extracting Structure from Point Clouds**
Steven Mills (steven@cs.otago.ac.nz)
Vision-based reconstructions from imagery typically model scenes as a large number of unconnected 3D points. However, many applications require surface models and/or more compact representations. In general estimating a surface model from point clouds is ill-posed, however it can often be simplified in specific applications. One such application is in building modelling, where flat surfaces at right angles to one another are common. The aim of this project is to investigate techniques for approximating large point cloud models by small sets of planes or other primitives.

6. **Design your own CPU (and computer)**
Andrew Trotman (andrew@cs.otago.ac.nz)
In this project we'll use an FPGA board to design a computer from scratch. We'll use the Verilog hardware description language to describe a CPU and flash it into an FPGA board (we already have the boards). We'll use the FPGA on-board RAM and the I/O to VGA monitor and USB for keyboards. We'll re-think the general purpose computer, and perhaps even build a special purpose CPU for a search engine.

7. **Hands on Software Engineering**
Andrew Trotman (andrew@cs.otago.ac.nz)
There are several critical points in the scaling up of a software development project. One of those is when the software reaches a point at which it needs automated daily builds, regression tests, portability checks, and (above all) organized scheduling and feature management. The ATIRE open source search engine written by members of the Information Retrieval lab has reached this point. In this project we'll look at best practice in software engineering and how we can apply that project. We'll turn a garage build system into a managed development project.

8. **Memory Checkers**
Andrew Trotman (andrew@cs.otago.ac.nz)
Purify is a fantastic run-time memory checker that identifies memory leaks, over and under runs, mismatched malloc and delete calls, and a plethora of other memory and run-time issues. Problematically, it does not run on 64-bit Windows. This makes writing programs in C/C++ that use dynamic memory allocation difficult. In this project we will survey alternative tools including those that use source-code annotation, object code annotation, and even static checkers. We will design a code base to test the tools, and we will test them on a real-world code base. Finally, we will make a recommendation to the Information Retrieval research group as to a set of tools to use.

9. **Scratch-based Interfaces**
Brendan McCane (mccane@cs.otago.ac.nz)
In this project you will investigate the use of scratching on a surface as a gesture device for control of a computer or similar device. To get an idea of how such input, might work, you might like to look at the seminal work of Chris Harrison. Ideally, such a device could be used to replace the remote control of your TV by allowing scratch gestures anywhere on a wall for example. In this project, you will test the limits of the idea by investigating how many different gestures can be effectively recognised.

10. **Camera-based mobile device localisation**
Brendan McCane (mccane@cs.otago.ac.nz)
In this project you will investigate the use of a mobile device camera for determining the pose and location of the device from live images in real time. Such a system could be used as the underlying technology for a markerless augmented reality system such as Sony's SmartAR[1], although augmented reality need not be part of this project. The target platform could be Android or iOS (iPod, iPhone, iPad).

11. **Spinal cord axon segmentation**
Brendan McCane (mccane@cs.otago.ac.nz)
This project is to develop a segmentation procedure for high resolution microscopy images of the spinal cord. The images consist of many cells where each cell has an inside part (intracellular space) and a thick boundary (myelin). There is also space between the cells that is less interesting but must be accounted for. This is quite a challenging project, but it will provide a good platform to learn about many image processing techniques.

12. **Augmented reality in dentistry**
Brendan McCane (mccane@cs.otago.ac.nz)
For decades, the standard method for dental imaging has been radiography. However, the introduction of 3D augmented reality has provided us with new tools that have the potential to revolutionize this fundamental area. The aim of this elective is to help design a program that is capable of simulating a virtual reconstruction of a tooth root using its crown (top half of the tooth in the mouth) as a reference structure. The software should ideally be designed for a smartphone in order to take advantage of its built-in high-resolution camera and display, as well as its portability. There are two possible projects which could be pursued: (1) development of the 3D augmented

[1] `http://www.sony.net/SonyInfo/News/Press/201105/11-058E/index.html`

reality software, and; (2) development of superimposition algorithms for predicting root morphology from the tooths crown. This project would be a great opportunity for students to be involved in a multidisciplinary biomedical team. For more information, please contact Brendan McCane (mccane@cs.otago.ac.nz) and/or Mauro Farella (mauro.farella@otago.ac.nz).

13. **Exploring orthodontic tooth movements with four dimensions**
Brendan McCane (mccane@cs.otago.ac.nz)
There is an increasing demand for orthodontic treatment (braces) amongst teenagers and adults, with teeth being straightened with the aid of CAD/CAM methods. Orthodontic tooth movements are generally described in two dimensions using a single center of rotation. However, in reality this is an over-simplification because teeth tend to move in 3D along an axis of rotation that is constantly changing over time. The finite helical axis (FHE) can be used to describe these complex movements. The aim of this project is to describe tooth movements in four dimensions (including time) by tracking the instantaneous changes of the helical axis over time. This project may have significant implications for clinical orthodontics.

14. **Indoor localisation with wireless sensor network and accelerometer**
Zhiyi Huang (hzy@cs.otago.ac.nz), Haibo Zhang (haibo@cs.otago.ac.nz)
Wireless sensor networking (WSN) is a rapidly growing technology due to its low power consumption and cheap price. It has a wide range of applications, including smart energy, remote hospital care, industrial process control, home automation, etc. In this project, we are going to apply WSN to tracking the location of an indoor moving object such as a robot. We have dozens of wireless sensors (telosb) for localisation using RSSI (received signal strength indication). We also have the acceleration signals used for measuring the acceleration and speed of the moving object. In this project, we will need to integrate all these signals together with techniques like particle filter to accurately localise the moving object carrying a wireless sensor and an accelerometer. A good math background is preferred for the project.

15. **Quantifying conceptual density in text**
Anthony Robins (anthony@cs.otago.ac.nz), Alistair Knott (alik@cs.otago.ac.nz)
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.

16. **Mining problem code data**
Anthony Robins (anthony@cs.otago.ac.nz)
We have a dataset consisting of thousands of counts of problems experienced by students in our introductory programming course. So far only a preliminary analysis has been conducted. What else lies in the data? Can we predict how well a student will do on the basis of the problems that they meet? Tools from machine learning and data mining will be relevant.

17. **Attractor spaces in Hopf eld 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.

18. **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.

19. **Games with video interaction**
Geoff Wyvill (geoff@cs.otago.ac.nz)
Since 2005, we have been demonstrating a game played in a lecture theatre. A big screen shows the audience in mirror image. Players interact with additional objects in the mirror world and the only input is from a camera on the audience. Last year, we improved the game by connecting a faster camera and fixing some minor bugs. Use the improved environment to create an game or activity that can be played in an ordinary room rather than a theatre.

20. **The dome display**
Geoff Wyvill (geoff@cs.otago.ac.nz)
We have an experimental screen in the form of a fabric dome three metres wide. So far, we have used it for two projects in virtual reality: a shopping cart simulator and a hang glider simulator. Build a new dome application. If successful, this can be demonstrated at the International Science Festival in Dunedin in July 2012.

Special note from Geoff Wyvill: I am officially retired and I do not intend to supervise as many projects as I have in the past. I do realise that a lack of graphics projects could be a disappointment to some students. So I am offering some. However, I will not take anyone who has not discussed it with me first. So don't put in a choice for one of these unless you have seen me. Email me to make an appointment. Don't leave it until Friday!

21. **A neural network model of cognitive modes and mode-switching**
Alistair Knott (alik@cs.otago.ac.nz)
There's a lot of interest in cognitive science at the moment in the idea that brain processing happens in several different 'modes', which are implemented by distinct distributed networks of brain regions. For instance, there's one brain network which is active when you are engaged in a sensory or motor task, and a quite different network which is active when you are resting quietly, and yet another network which is active when your attention is drawn to a stimulus in the world. An interesting research question concerns how the brain decides which network to engage at any given time. There must be a control mechanism which selects the network which is most appropriate, but we don't know much about what this is. In this project, you will read up a bit about large-scale brain networks, and implement a (simple) model of the mechanism which performs this selection process.

22. **Discovering large-scale brain networks using functional connectivity analysis**
Alistair Knott (alik@cs.otago.ac.nz)
This project is also on the topic of large-scale functional brain networks (see project above). One way these networks have been identified is by analysing brain activity gathered by functional magnetic resonance imaging (fMRI): they can be identified as sets of brain regions whose activity is highly correlated. (Typically the activity of the regions in one network is anti-correlated with that of regions in other networks.) In this project, you will obtain some publicly available fMRI data, and some analysis tools, and explore some hypotheses about brain networks: specifically, you will be looking for brain regions involved in deciding which network to activate. We will probably use data and tools from the functional connectomes project[2].This project would suit someone with some maths background, and an interest in correlational analyses.

23. **A model of 'properties' in the visual system**
Alistair Knott (alik@cs.otago.ac.nz)
When we look at an object, we automatically identify what category it belongs to (e.g. that it's a dog), but afterwards, we sometimes notice additional individual properties that it has (e.g. that it's brown, or big). In this project, you will implement a neural network model of how observers 'notice' or 'perceive' such individual properties. The project will involve reading up about cognitive models of object perception, and implementing a new model.

24. **Power management for a distributed data base in a box**
Richard O'Keefe (ok@cs.otago.ac.nz), David Eyers (dme@cs.otago.ac.nz)
Distributed data bases like Chord, Apache Cassandra, Tapestry, CouchDB, or Riak let you store information in multiple nodes of a network, and access it from multiple nodes, trying to stay available even when individual nodes fail. If you have a cluster computer made up of a large number of processors (say a thousand) in the same cabinet, you might want to put a node in a reduced power or very low power state, not because it has failed, but because the system isn't currently being used heavily enough to justify using all that electricity. In a full power state, a node would participate fully in data base support, accepting updates and responding to queries. In a reduced power state, a node might have enough energy to track updates, but not to answer queries. Having many copies of data is an important part of distributed data bases, so a node can stop answering queries if the information is held somewhere else. In a very lower power state, the only thing a node can respond to is a request to turn back on. Its information will then be stale, and it may be unresponsive for a while, using electricity, while it gets back up to date. What is a good policy for choosing which nodes to turn on and off? Are the distribution protocols designed to handle sporadic failures adequate to cope with frequent deliberate shut downs/restarts? What difference does having a reduced power state that can track updates but not answer queries make?

[2] http://fcon_1000.projects.nitrc.org/

25. **Patently obvious**
Richard O'Keefe (ok@cs.otago.ac.nz)
We have a number of document collections for doing information retrieval research on. I am particularly interested in a collection of 100,000 US patents. The Wikipedia collection has links between articles; and patents refer to each other. How can we visualise a collection of this many documents? How can we display them so that interesting patterns of properties or connections stand out so we notice them? This involves extracting information from XML and data visualisation.

26. **Sustainable Geocaching**
David Eyers (dme@cs.otago.ac.nz)
There are plenty of Geocaches[3] placed around Dunedin. Thanks to the Otago Regional Council, we also have access to details about where the different routes' bus stops are. Additionally, web services from LINZ now openly provide an accurate topological information about the city, which can be used to assess the relative comfort of walking routes. Such data sources as these, taken together, provide a wealth of possible constraints and preferences against which to optimise a geocacher's excursions. This project will compute and present route recommendations.

27. **Private Dropbox**
David Eyers (dme@cs.otago.ac.nz)
Products such as Dropbox provide highly useful, effective and convenient tools for synchronising data between multiple devices. However they have the downside of being run somewhere on the cloud, using server-side software that you dont have access to. At the other extreme, storing data in peer-to-peer networks avoids relying on a (logically) centralised service, however it becomes much harder to offer any quality of service guarantees. The aim of this project is to develop an efficient file synchronisation tool - at first just using a text interface, along the lines of unison or rsync - that provides an intermediate step between peer-to-peer and cloud-based services. To address data privacy concerns, it should include capabilities for explicitly controlling the subset of a distributed storage system in which particular datasets can be replicated. (There are many other possible areas of extension not discussed here.)

28. **Introspection within the Galaxy toolbox**
David Eyers (dme@cs.otago.ac.nz)
Galaxy is a web-based platform for large-scale genome analysis that is implemented in Java. In facilitates integration of a wide range of data sources with data analysis and visualisation tools. The platform aims to be extensible, allowing minimum programming to effect addition of software into the framework (e.g. locally developed tools), while still providing non techically focused researchers with a convenient and informative interface to their experiments. A current flaw in the platform is its minimal "out-of-band" monitoring. While Galaxy will inform users of the success of intermediate steps in a scientific workflow, it will not react to easily predictable system failures, such as subtasks that are heading toward use of impractical amounts of memory. This project aims to integrate a resource monitoring abstraction into the Galaxy platform, so that instrumentation can be tuned to report intermediate task progress to users, and to preemptively indicate potential oncoming software failures.

[3] `http://en.wikipedia.org/wiki/Geocaching`

29. **Management of cloud computing nodes**
David Eyers (dme@cs.otago.ac.nz)
For many years, clusters of mass-produced PCs have been employed as an alternative to more specialised supercomputers. In order to allow a collection of PCs to appear more like a single, virtual supercomputer, it has been common to install an operating system tuned for this task, such as a cluster-oriented Linux distribution. Recently, the shape of some computing clusters has changed significantly. On one hand, is dedicated "high throughput" computing - e.g. map-reduce systems. On the other, are cloud computing environments. Open source software has started to catch up of these newer cluster forms, and it is possible to set up a "private cloud" or a Hadoop cluster (for example). Do these newer cluster initialisation and control systems subsume the jobs of previous cluster control software? This project investigates the comparative capabilities of these different cluster platforms.

30. **XML parsing within programming languages: safety, eff ciency, and convenience**
Richard O'Keefe (ok@cs.otago.ac.nz), David Eyers (dme@cs.otago.ac.nz)
Parsing and generating XML is a common task within today's software. Libraries for parsing and manipulating XML are available for most common programming languages. However, some languages, such as Scala, have chosen to handle certain XML operations right down at the level of their language compiler. The co-supervisors of this project have both been involved in efforts to retrofit Scala-like XML support into other programming languages, such as Python. Is it worth tightly integrating XML into a language? Is enough expressiveness provided in the low-level language features to capture the majority of use cases? What are the degrees of difference in speed and memory usage, and security against injection, from handling XML in different ways within a language? This project aims to study these questions, and potentially implement improved language support for XML.

31. **Graph layout on a sphere**
Shawn Martin (smartin@cs.otago.ac.nz)
A graph consists of nodes connected by edges. Graph layout algorithms are designed to draw graphs subject to various aesthetic criteria. Examples of such criteria include minimal number of edge crossings, low density of nodes, symmetry, etc. Most graph layout algorithms draw graphs in two dimensions. There are some layout algorithms that work in three dimensions, but the resulting drawing is typically very hard to visualise. What about drawing a graph on a sphere? The layout would be three-dimensional, but would be easier to visualise. In addition it would have some interesting geometrical advantages over two-dimensional layouts. In this project you would learn about graph drawing and implement a spherical layout algorithm. An interactive visualisation component would also be considered.

32. **Molecular Humpty-Dumpty**
Shawn Martin (smartin@cs.otago.ac.nz)
Like Humpty-Dumpty, molecules can be broken into little pieces. Unlike Humpty-Dumpty, they can be put together again. The interesting thing is, however, that you might not have the same molecule you had when you started. In this project you will implement code to read Simplified Molecular Line Entry Specification (SMILES) data describing a molecule, produce the corresponding molecular graph (atoms are nodes, bonds are edges), break the molecule into pieces, and output new SMILES for each piece. How the pieces can be put together will then be examined.

33. **Cycloalkane generator**
Shawn Martin (smartin@cs.otago.ac.nz)
Cycloalkanes are carbon rings. They can have different numbers of carbon atoms and can assume interesting shapes, known as conformations. Cyclohexane, for example, can assume a chair and a boat shape. Cyclooctane can assume chair, boat, and crown shapes, as well as numerous other possibilities. We have code for computing the conformations of cycloalkanes, but it is unwieldy and difficult to use. Your job will be to overhaul the code so that it is easier to use and can sample cycloalkane conformations intelligently.

34. **Cyclooctane visual explorer**
Shawn Martin (smartin@cs.otago.ac.nz)
Cyclooctane is a molecule with eight carbons in a ring. We have enumerated all possible shapes, known as conformations, that this molecule can assume. We have discovered that the conformations occupy an interesting and unusual space known as a Klein bottle. It is difficult to understand how the conformations vary over the Klein bottle because the Klein bottle lives in four dimensions. In this project, you will develop an interactive visualization that can be used to explore how cyclooctane confromations vary over the Klein bottle.