480/490 Project Proposals for 2010

This is a list of suggested project proposals for 480 and 490 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 5**. 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. **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.

2. **Mining problem code data**
   Anthony Robins (anthony@cs.otago.ac.nz), Nathan Rountree (nathan@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.

3. **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.

4. **Simple robot vision**
Anthony Robins (anthony@cs.otago.ac.nz)
The department has many Lego Mindstorms robots. One of the sensors that we have is a simple camera capable of returning the coordinates of blocks of specified colours in the visual field. How well does it work? In particular, can we determine location on the basis of the arrangement of colours in the environment? This project may involve developing a code library/toolbox for using the camera.

5. **Playing cleverly**
Michael Albert (malbert@cs.otago.ac.nz)
Traditional AI approaches to playing strategic games have relied on static evaluation functions. Recently, great progress has been made in Bridge and Go, using an AI based around a Monte Carlo evaluation function – which simulates random futures. The aim of this project is to describe and understand this approach, and to apply it to one or more games. Games such as Tantrix, which involve a changing board determined by the players, would seem ripe for such an approach.

6. **Puzzling logically**
Michael Albert (malbert@cs.otago.ac.nz)
Many currently popular logic puzzles (Sudoku, Kakuro, Kenken, Nurikabe) require finding the unique completion of some partial configuration. This project will concentrate on some of the following issues:

- What is known about the complexity of these types of puzzles?
- In practice, how efficiently can we solve/check specific puzzles of these types?
- How can we generate "interesting" examples of these types?
- How can we generate entirely new logic puzzles of this family?

7. **Game development - in depth**
George Sealy (george@cs.otago.ac.nz), Michael Albert (malbert@cs.otago.ac.nz)
Design and build a complete, polished game. This is a cross-discipline challenge, covering a variety of topics; game design, graphics, AI, multi-threaded programming, software architecture, iterative refinement and project management for example. This project provides an opportunity for the student to choose a particular area of interest to study in depth. It may be to create a challenging (or perhaps 'human-like') AI, to investigate advanced rendering techniques or to develop and refine a novel game play mechanic. This is offered as an individual project. However if a group of students wish to work together on a game, it might be offered as a team project.

8. **Gene networks**
Brendan McCane (mccane@cs.otago.ac.nz), Peter Dearden
Gene networks are a useful tool for modeling the interaction between genes in biological organisms. In a gene network, genes are represented as nodes and the interactions between genes (either inhibitory or excitory) are represented in the network arcs in a similar manner to neural networks. We have already developed a system for simulating gene networks. This project will involve enhancing that system to study varuous aspects of gene networks.

9. **Scratch-based interfaces**
Brendan McCane (mccane@cs.otago.ac.nz) Stethoscopes (those hearing things that doctors use) can be hooked up to microphones to provide very sensitive audio input devices. Sensitive enough that scratching a surface can easily be detected. This project is to develop the software for a dual stethoscope system to act as a mouse and gestural input device.

10. **View-oriented parallel programming (VOPP)**
Zhiyi Huang (hzy@cs.otago.ac.nz)
If you have any applications you would like to parallelize, come to discuss with me. You will enjoy both message passing and shared memory based parallel programming. You will have a chance to enjoy our novel View-Oriented Parallel Programming(VOPP).

11. **Virtual aggregated processor (VAP)**
Zhiyi Huang (hzy@cs.otago.ac.nz)
In this project we use "bare metal" cores to accelerate application performance. You will have a chance to learn x86_64 architecture and low-level OS design and implementation.

12. **Programming languages, compilers, and computer architecture**
Richard O'Keefe (ok@cs.otago.ac.nz)
Three projects in this area:

- Improving a working Smalltalk-to-C compiler, using data flow or type analysis.
- Library extensions for the compiler: graphs, sockets, internationalisation support
- Development tools, such as an optimising assembler, for an unusual system called a Transport-Triggered-Architecture.

13. **Natural language processing and information retrieval**
Richard O'Keefe (ok@cs.otago.ac.nz) What might you do with a document collection like the Wikipedia, or an even larger one? How much does spelling correction help? How could we divide a large collection into groups of related documents? Can we automatically recognise technical papers and other "stylised" document genres?

14. **Software engineering**
Richard O'Keefe (ok@cs.otago.ac.nz)
How can we read large programs? What kinds of tools can help? We have the source code of the Windows Research Kernel, 831 000 lines of C (319 000 SLOC) in 656 files. What kinds of patterns can we automatically detect? How might something like "Formal Concept Analysis" help? Could information retrieval help?

15. **Search engine thesaurus**
Andrew Trotman (andrew@cs.otago.ac.nz)
Search engines often fail because of the "out-of-vocabulary" problem. The user types "NZ cow industry" into the search engine but the web page contains "Commercial bovine farming in New Zealand" and so is not found. We can address the problem by expanding the user's query with a thesaurus. In this project we will extend the ANT search engine with thesaurus lookup using WordNet and measure the performance change to determine its effectiveness.

16. **Memory Errors**
Andrew Trotman (andrew@cs.otago.ac.nz), Richard O'Keefe (ok@cs.otago.ac.nz) Memory errors in C and C++ are some of the hardest errors to find. Luckily we can use tools such as Purify and Valgrind to help us find such errors. Unluckily we are in the process of migrating to 64-bit operating systems and many of these tools just don't work 64-bit. In this project we'll take a look at the current state of the art in memory checkers. We'll test each tool by running it over a collection of research software from within and outside the department. Which errors are found, and which tools should be adopted?

17. **Digital Music**
Andrew Trotman (andrew@cs.otago.ac.nz), Richard O'Keefe (ok@cs.otago.ac.nz)
Wikipedia claims that digital music dates to the 1960s. It in fact dates to earlier than the 1860s when such devices as the barrel organ and pianola were invented. But we can't play old digital music, whether from the 1960s or the 1860s, any more because we don't have the necessary digital devices. In this project we will build an electronic pianola and use it to convert some old digital music into MIDI files so that we can play the music on our modern devices.

18. **High performance sport**
Andrew Trotman (andrew@cs.otago.ac.nz), external
The hp sport website (`www.highperformancesport.co.nz`), sponsored by SPARC is intended for the exchange of ideas and information among the high performance sport community. They have proposed several possible projects including information discovery and visualisation. In order to propose undertaking such a project you must act *quickly* to discuss the possible areas and scope with both Andrew, and the external client.

19. **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 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.

20. **A model of quantification in semantic memory**
Alistair Knott (alik@cs.otago.ac.nz)
When we learn a new fact about an object, we store it in a form of long-term memory called *semantic memory*. Semantic memory can be queried in different ways, but one way is to activate all objects of a certain type—for instance, all dogs—so that their properties can be inspected collectively. For instance, we might find that all the dogs we retrieve have a particular property. In this project, you will implement a model of semantic memory which supports the retrieval of groups of objects and their properties.

21. **Audience games**
Geoff Wyvill (geoff@cs.otago.ac.nz)
Since 2005, we have been developing a simple concept of a computer providing a game for a group of people in a lecture theatre. The software works well using just a single camera and projector. But the range of games is limited. They are all based on hitting virtual balls around. Develop a new game for this system.

22. **Living room games**
Geoff Wyvill (geoff@cs.otago.ac.nz)
The illusion of 3D interaction in a lecture theatre using a single camera has been used successfully to play games with an audience. But the method depends on a trick of calibration, working on the assumption that the members of the audience are sitting in tiered seats. Find a similar way to be able to play games between people in an ordinary room using only a single, fixed camera. Devise and implement a suitable game to demonstrate the system.

23. **The new Watching Window**
Geoff Wyvill (geoff@cs.otago.ac.nz)
Our Watching Window display uses cameras to track head and eyes and displays an image from the user's point of view. It includes a 42 inch lenticular stereo display. This is a screen that can present a stereoscopic 3D image to the user without the use of special glasses. The software needs an overhaul to make it more reliable and to make better use of the special screen. Learn all about the Watching Window and get it working properly.

24. **Perception of vibrato**
Geoff Wyvill (geoff@cs.otago.ac.nz)
Vibrato is a variation of pitch used by musicians to decorate notes. The pitch of a note with vibrato is perceived as being somwhere between the highest and lowest pitches in the note, but where exactly? Devise a set of experiments to settle this question by getting musicians and others, to estimate the pitch of a note with vibrato. This project was attempted in 2007 but the results were inconclusive. Build on the experience of the first attempt.

25. **Make a movie**
Geoff Wyvill (geoff@cs.otago.ac.nz)
Join a team of computer animators to create a short animation suitable for showing at one of the art and animation festivals. Over the years, we have created six such movies as computer science projects. Two have made it to the SIGGRAPH Animation Festival and one also appeared in other festivals in Japan and Korea. Don't choose this project without seeing me first to discuss it.

26. **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.

27. **Mining public data sources**

Nathan Rountree (nathan@cs.otago.ac.nz)

Social networking sites, aggregator sites, and sites where users generate or rate content contain a lot of data on buying habits, artistic taste, and geographical or demographic trends. This project will involve getting some of that data (sometimes through the site's own API) and implementing a recently-proposed model from the data mining literature.