480/490 Project Proposals for 2005

This is a list of suggested project proposals for 480 and 490 projects. Note that some of the projects are for COSC480 only or for COSC490 only while others are for either course. Please keep your choices to the correct paper. 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. However, this requires the agreement of a supervisor before you can begin.

- Select three projects that you think will interest you.
- Rank 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!)
- Also send an email with your project preferences to me (malbert@cs.otago.ac.nz).

Michael Albert malbert@cs.otago.ac.nz
400 level project coordinator

1. **Interactive Image Editing (480/490)**
   Brendan McCane (mccane@cs.otago.ac.nz)
   Interactive image editing is the process of interactively selecting objects from an image. These techniques are useful for cutting objects from one image and pasting them into another or for specifying objects in an image for later processing (e.g. selecting anatomical structures from x-rays). A common example of this is the "magic lassoo" in Photoshop and similar applications. A paper in Siggraph 2004 describes a technique called GrabCut which is based on finding minimum energy paths in a graph-based representation for object selection (similar to Dijkstra's algorithm). The task for this project is to implement the GrabCut technique and evaluate its usefulness for specifying vertebra from an x-ray of the lumbar spine.

2. **Facial Feature Extraction (480/490)**
   Brendan McCane (mccane@cs.otago.ac.nz)
   The Watching Window (WW) is an unencumbered virtual environment which displays a different picture to the user depending on her location, giving the illusion of looking through a "Window" into a three-dimensional world. To present the correct picture, the WW needs to know the location of the user's eyes. Currently, it finds the location of the user's face and guesses the eye locations. This project will involve implementing an Active Appearance Model for more accurately locating the eyes in the WW.

3. **Twists, Exponential Maps and the Watching Window (490)**
   Brendan McCane (mccane@cs.otago.ac.nz)
   Any rigid three-dimensional motion can be described by a twist - an axis about which the object rotates and a translation along the axis. Exponential maps can be used to combine multiple twists and this turns out to be a useful representation for articulated structures (such as human bodies and hands). This project involves implementing a recent technique (using twists and exponential maps) for estimating the pose of a human body in the Watching Window, focusing on the location of the hand.

4. **Signal Processing Application (490)**
Brendan McCane (mccane@cs.otago.ac.nz)
Physical measurement of some underlying property $h(t)$ of an unknown material or composite of materials often returns a signal of the form: $y(t) = \int f(t-x)h(x)\,dx + E$ where f is some fixed response function, and E an error term. Classically, this is resolved by attempting to determine the function f and then "deconvolving". However, this approach is very prone to problems in noisy data. When the underlying material has sharp boundaries it may be possible to use training data from a known dataset to find a stable response function f and then use statistical techniques to estimate h.

This approach works well enough for most work, but we do have situation where we have training sets with $h(t)$ known, and wonder whether soft computing methods, such as genetic programming or neural nets could be used to derive the deconvolution function and if so, possibly handle cases where $h(t)$ has smoother boundary transitions?

The aim of this project is to investigate such methods and compare their effectiveness to current approaches to the problem. It will be carried out in collaboration with Phil Scadden of the Institute of Geological and Nuclear Sciences.

5. **Temporal Logic (480/490)**
Willem Labuschagne (willem@cs.otago.ac.nz)
How can one formalise statements involving time? Planning involves the future and may be informed by the past. An agent committing to a contract may be undertaking to do something by a deadline or for a period of time. What sort of language should be used to represent these ideas? There are a number of decisions that lead to alternative approaches — whether to focus on time instants or time intervals, whether to include temporal parameters in predicates or to use temporal operators on sentences, and so on. The purpose of this project will be to explore the literature in order to build up a coherent picture of what the main options are and how they work. No programming is expected, and a general interest in logic is assumed. For the COSC490 option, it is expected that recent work by Colombetti on agent communication languages be included.

6. **Data Quality (480/490)**
Richard O'Keefe (ok@cs.otago.ac.nz)
It has been estimated that 20% of the fields in real business databases are wrong. The best time to fix this is at the point of entry.

There is a Dunedin business working on this problem. They're looking at the problem of automatically fixing up street addresses as they are entered, mainly using data provided by the Post Office. They have suggested it to us as a research topic, and it's a good one.

There are many ways to approach this:

- algorithmic: adapt ternary search trees to do nearest neighbour searches by Levenstein distance instead of Hamming distance
- algorithmic: would splay ternary search trees be better than simple ternary search trees?
- algorithmic: would N-grams do a better job than ternary search trees?
- algorithmic: given a street address to latitude-and-longitude mapping, how can you use that to efficiently compute adjacency and (for example) resolve the common error of the suburb name being the (correctly) spelled name of an adjacent suburb?

- algorithmic: given a street address to latitude-and-longitude mapping which is incomplete, can you use other postal address information to compute approximate geographical location?
- architectural: how can you use context to focus spelling correction?
- application: how can you use ideas shared with address clean-up to help with names?
- application: how can you use ideas shared with address clean-up to help with free text from a particular domain (general medical practice)?
- user interface: what's a good way to build address clean-up (or name clean-up) into a user interface?
- natural language processing: can you find New Zealand postal addresses in Web pages, e-mail, or other electronic data sources?

7. **Exploring XML Query Languages (480/490)**
Richard O'Keefe (ok@cs.otago.ac.nz)
An XML document or document collection can be viewed as a (semi-structured) data base. There are several query languages for XML including XPath, XQL, XML-QL, XIRQL, and XQuery. At the top end we can regard XSLT (a rather ugly functional language for transforming XML) as a generalised query language.

   I am interested in the usability and implementability of XML Query Languages. The scope of an exploration can be chosen to suit the needs of the student.

   A Web search for "XML Query Languages" will find more than you have time to read. Two useful starting points are http://www.perfectxml.com/XPathDotNet.asp and http://www.perfectxml.com/XPathDotNet.asp.

8. **Information Retrieval on the Cell (480/490)**
Richard O'Keefe (ok@cs.otago.ac.nz)
IBM, Sony, and Toshiba have announced a new processor architecture, the Cell. For a description, see http://www.blachford.info/computer/Cells/Cell0.html. This is to be the engine for the PlayStation 3. Each Cell contains a processing unit which is basically a POWER5 (roughly, the next generation after the G5 used in current Macs) plus 8 64-bit vector processing units. 256GFLOPS peak performance at 4.6GHz.

   The purpose of this project is to investigate ways in which an Information Retrieval engine could exploit the vector processing units. This is an architecture and design topic, not a detailed implementation one.

9. **Cleaning Bathymetric Data (490)**
Richard O'Keefe (ok@cs.otago.ac.nz)
Cleaning bathymetric involves removal of bad records (consisting of latitude, longitude, depth, source) from large datasets of around 4-10 million records. Because data is pooled from many sources, the data quality is extremely variable. All methods of bathymetric surveying are prone to particular types of systematic error and affected records must be removed before modelling.

   Mesh manipulation methods show some considerable promise. Turning the Garland SCAPE algorithm for mesh-simplification on its head (see http://graphics.cs.uiuc.edu/~garland/scape.html), points p could be cleaned by successive removal of points with a high value for metric $\mathrm{Error}(p)$ calculated on triangle for which the point is a member.

Broadly the algorithm is:

repeat
Remove the point p which has highest $\mathrm{Error}(p)$ by edge collapse
Adjust triangles that used p
Recalculate Error for all points in affected triangle
until sufficient points are removed

The devil is in the detail. Because the algorithm iterates over the data set, speed is an issue. "adjust affected triangles" requires a data structure for triangles that can rapidly find triangles used by a particular point but this must be balanced against memory usage with large datasets.

10. **Remote Process Monitoring in Cluster Computers (490)**
Paul Werstein (werstein@cs.otago.ac.nz)
The department has a cluster of 32 Linux PC's. Remote access to the cluster is by way of a gateway machine. The other nodes in the cluster are considered worker nodes. The goal of this project is to be able to monitor processes running on the worker nodes from the gateway machine. In other words, the process should appear to a user as if it were actually running on the gateway machine. The project requires kernel module programming. The systems research group has some books to guide this effort.

    Last year a 480 project proposed a design. In this project, you would evaluate that design and implement it or some other design.

11. **R-tree Design for Race Cars (480/490)**
Paul Werstein(werstein@cs.otago.ac.nz)
The R-tree is a technique to index multidimensional data. In this case, the dimensions are X-Y coordinates and time. The goal of this project is to design, implement, and test an R-tree for indexing race car data. What sets this application apart from others is the repetitive motion of the cars. In other words, they appear in the same location multiple times during a race. Also there is a need to store data very efficiently due to the huge volume of data collected. This requirement dictates the design of the R-tree must be modified from its original design. However it must still look like an R-tree.

    A 480 project would implement and test one solution. A 490 project would evaluate a few potential solutions.

12. **Bidirectional Wireless Communication for Model Airplane (480/490)**
Paul Werstein(werstein@cs.otago.ac.nz)
The department has a project involving the intellegent use of a model airplane. Potential applications include search and rescue, infrastructure survey, etc. The goal of this project is to evaluate the possiblility of bidirectional wireless communication. It would involve determining frequencies that can be used and commercially available parts. A preliminary design should be developed. Some knowledge of electronics would be beneficial.

13. **Classifying Thai Burial Pots (490)**
Chris Handley (chandley@cs.otago.ac.nz)
I have a CD from Professor Charles Higham of the Department of Anthropology containing 300 photographs of burial pots that have been excavated from the site of Ban Lum Khao in Thailand. Some of these have been classified into a number of groups by a PhD student, however any manual system of classification is idiosyncratic and often breaks down as new data are found. For this project you would be expected to research different methods of classifying objects in general and to devise a scheme that could be applied to the current data set, but that should be robust to expansion of the data set. The results of this classification could be compared with the previous classification. In addition to, or as part of, this task, we would like you to investigate ways of parameterising the shapes of the pots, so that we can fill in missing or broken bits. This parameterisation could also be used to explore the 'space' of possible pot shapes.

14. **Text in Images (2 projects) (490)**
Chris Handley (chandley@cs.otago.ac.nz)
For most of us vision is the primary source of information. Furthermore, in the developed world, an enormous amount of information about the world around us and our place in it is encoded as environmental text (street signs, bus stops, opening hours for libraries, museums and shops, restrictive notices, and so on), not to mention the enormous amount of information on products, food items and in books, magazines and the like. The problem is that most of this information is inaccessible to people with impaired vision. The ultimate goal of this project (to be run in collaboration with Information Technology department at Otago PolyTechnic) is to produce a device that will "read" such environmental text to visually impaired people. Project 1 involves collecting images of the real world that may or may not contain text (I have access to a set of images to start you off) and to devise methods to select areas that are very likely to contain text. I have a set of references on methods for doing this, and you will be expected to use these as a starting point for a comprehensive literature search. Note that, your application should, ultimately, run in real time, so any proposed method should be capable of being sped up to achieve that goal. Project 2 involves taking regions of images identified in project 1 as probably containing text, (you may have to fake this to start with) and applying an optical character recognition algorithm to it to extract the actual text. Bear in mind that the text will be in a variety of fonts, in a variety of colours and against a variety of backgrounds, and will almost certainly not be "square on" to the camera. Your task will be to find an algorithm, either "off the shelf" or of your own devising that will perform well under these conditions, for instance, you may need to include a dictionary to assist with "guessing" words when parts are obscured.

15. **Title: Geocaching (480/490)**
Chris Handley (chandley@cs.otago.ac.nz) and Andrew Trotman(andrew@cs.otago.ac.nz)
When the American military descrambled the GPS satellite system on 1st May 2000, they unwittingly invented a new game, *Geocaching*. This game is played much like orienteering, except that the compass is swapped for a GPS receiver. Playing this hide-and-go-seek game, players have hidden nearly 150,000 "caches" in over 200 countries in just a few years. What's more, there's an active Geocaching community in Dunedin.

This project consists of four parts:

(a) Playing the game. To be successful at this project you'll have to understand the game. This requires finding some pre-hidden caches, and placing at least one cache yourself.

(b) A comprehensive review of the available tools for Geocaching. This involves becoming familiar with the existing web-sites, downloadable software, and reporting on what these sites and software do.

(c) Designing and building "cross platform" Geocaching software (the final application is expected to run on at least Windows and Macintosh). This software will trawl the web, find Geocaches, download the coordinates (and logs) and plot the coordinates on topographic maps of New Zealand. Of course, the ability to download coordinates to a GPS receiver using either serial or USB is essential.

(d) Port this software to run on a hand-held computer (to be chosen, probably PalmOS) so the Geocacher can carry electronic cache descriptions with them.

For further details talk to Chris or Andrew, or take a look at the International Geocaching web site, http://www.geocaching.com.

16. **Cards and cryptography (480)**
Mike Atkinson (mike@cs.otago.ac.nz)
This project is derived from the Russian cards problem: suppose three players A, B, C are given cards from a deck of 7 cards, A receiving 3 cards, B receiving 3 cards, and C the remainng card. How can A inform B what she holds without C learning a single card of her hand or of B's hand? It turns out that this problem can be solved and it has many generalisations. However (in these generalisations), *some* information is leaked. In this project we want to find methods where we can ensure that the leaked information doesn't give C a chance of guessing A's cards with better than random chance (and various similar hypotheses). Further information about similar problems can be found at http://www.cs.otago.ac.nz/staffpriv/mike/Papers/Cards/Cards4.ps.

17. **No Three in a Row (490)**
Michael Albert (malbert@cs.otago.ac.nz)
How many points can be chosen from an $n \times n$ square grid without obtaining three in a single line? This problem of Dudeney is approaching its hundredth birthday without a full solution. The upper bound of $2n$ is attained for $n \leq 52$ but there are strong reasons to believe that it will not be attained for sufficiently large n. The goal of this project is to confirm and hopefully extend the known results on this apparently simple problem. The current records on this problem are due to Achim Flammenkamp (see URL below) and use a branch and bound algorithm to exhaustively search for possible solutions. This project would aim to examine this as well as other possible searching strategies and algorithms and to compare their effectiveness. Implementation in a parallel or cluster computing environment is a possible option. For further information see http://wwwhomes.uni-bielefeld.de/achim/no3in/readme.html.

18. **Determinizing and Minimizing Automata (490)**
Michael Albert (malbert@cs.otago.ac.nz)
Non-deterministic finite automata (NFAs) arise routinely in many theoretical and practical applications. Many operations on automata require as input deterministic finite automata (DFAs). It is well known that determinizing an NFA is possible but may lead to an exponential blowup in the number of states. Following the determinization with a minimization step often leads to DFAs whose complexity is not much greater than that of the original NFA. However, the size of the automaton produced in the determinization step is often a bottleneck. This project will aim to investigate methods of circumventing this problem. One approach is to use methods that reduce the number of states in the original NFA, while another would break the determinization step down into a sequence of partial determinizations, each of which could be followed by some reduction. Many proposed solutions to this problem in the literature are described only as abstract algorithms, so another part of the project would be to implement some of these algorithms in order to compare their practical effectiveness.

19. **CGSuite (480/490)**
Michael Albert (malbert@cs.otago.ac.nz)
CGSUITE is an open source software project that provides a virtual laboratory for investigating properties of combinatorial games (roughly, abstract two player games with perfect information). The project is designed to be easily extensible and there are a number of different aspects of it which await further development. Projects in this area could range from the quite theoretical to the much more practical. A project of the theoretical type might involve improving some of the fundamental algorithms for dealing with combinatorial games, while some of the second type would include designing and implementing a useful help system, adding a "play" mode, or improving I/O in various ways. For further information see: http://cgsuite.sourceforge.net/.

20. **Ordering Geological Data (490)**
Michael Albert (malbert@cs.otago.ac.nz)
The construction of the geological time scale is based on data found in sections (e.g. drilling cores, cliff faces) which contain evidence of specific events. Unfortunately, these events are not always identified consistently nor do they always occur reliably in the same order between sections. The challenge is to provide the "best" synthesis of the available data. This is similar in some respects to the multiple alignment problem in computational biology and is subject to the same algorithmic difficulties.

This project will be carried out in collaboration with Phil Scadden of the Institute of Geological and Nuclear Sciences. The challenge is to develop an appropriate fitness function to measure how good a proposed synthesis of the section data is. Search techniques, both exact and heuristic (concentrating particularly on genetic algorithms) will then be used to determine a proposed event sequence. In turn, this can be used as a seed for a higher order method which seeks not only to determine the order of events, but their actual time scale.

21. **A Model of Action Recognition (490)**
Ali Knott (alik@cs.otago.ac.nz)

In a classic experiment, Gunnar Johansson (1973) attached lights to a human actor's limbs and put the actor in a dark room. If the actor stood still, an observer could not recognise what they were looking at: they just saw a collection of lights. But as soon as the actor started moving, the observer immediately saw a moving person, and recognised the action being performed. This experiment is seen as evidence that the human visual system can recognise actions just using information about the patterns of movement in the image, rather than information about shape or form.

The aim of this project is to implement a model of this mechanism. There has been lots of work in this area, so to start with a review of the literature will be needed. A model will then be chosen to implement. The goal of the project is to build a system that can recognise actions in Johansson-style point-light displays. Some useful software for modelling the early visual system has already been developed by students in the department; the project may involve extending some of this software. We also have access to the University's "Gollum-style" motion capture suite, which can create point-light displays from arbitrary actions.

For an introductory paper about action recognition, and a demo of a point-light display, see here http://www.uni-tuebingen.de/uni/knv/arl/arl-model/bm.html.

This project will involve the use of neural networks. COSC343 would be a useful background, and taking either COSC453 or COSC460 would also be useful.

22. **Environment-centred Representations (490)**
Ali Knott (alik@cs.otago.ac.nz)

Since we are mobile agents, our view of the world is always changing: when we move our heads, or move through the environment, the pattern of light which falls on our retina changes radically from moment to moment. But instead of attributing these changes to motion of objects in the environment, we are able to attribute them (where appropriate) to changes in our own position within a fixed environment. Somehow we are able to transform the retina-centred information we receive about the world into a coordinate system centred on the environment, in which we feature as one object among many others. Our own position and direction of gaze in this new coordinate system are somehow kept updated by the commands we give to our motor system to walk or to change the position of our head or eyes.

The goal of this project is to implement a model which creates and maintains an environment-centred representation of an agent's surroundings. If the agent moves through the environment, or alters its direction of attention, its retinal image will change dramatically, but the only changes to the environment-centred representation should be to the position and orientation of the agent itself. For an introduction to models of spatial cognition, see http://www.icn.ucl.ac.uk/nburgess/papers/HartleyECS.pdf.

This project will involve the use of neural networks. COSC343 would be a useful background, and taking either COSC453 or COSC460 would also be useful.

23. **Extensions to the Kaitito system (480/490)**
Ali Knott (alik@cs.otago.ac.nz)
This project involves implementing an extension to the human-machine dialogue system which we have been developing in the AI lab over the last several years. Currently, the system supports dialogue in either English or Māori; the user can make assertions or ask questions, and the system can likewise make assertions or ask questions. The system incorporates some standard natural language processing technology: grammars, lexicons, sentence parsers and generators, utterance disambiguation tools, and dialogue management routines. It has recently been extended to support multi-speaker dialogues, and tutorial dialogues to assist language learning. You can find out more about the system by looking at the project homepage, http://tutoko.otago.ac.nz:8080/teKaitito/.

There are several further extensions which could form the basis for a 4th year project.

- Including a model of conflict between dialogue participants, to allow the representation of inconsistencies between them, and to implement dialogue strategies for signalling and resolving inconsistencies. (with Hans van Ditmarsch)
- Extending the grammar and/or lexicon of the system, in either English or Māori or both. (N.B. the system's knowledge of Māori is very limited, so you wouldn't need to know much Māori in order to be able to do this project.)
- Making extensions to the system's interface, to add extra facilities for language-learning dialogues (e.g. requests for translations of sentences or words).
- Or any other ideas you have!

24. **Title: Model checking of knowledge (480/490)**
Hans van Ditmarsch (hans@cs.otago.ac.nz)
The goal of this project is to get some hands-on experience with the model checker MCK (see http://www.cse.unsw.edu.au/~mck/). A model checker is a program for the verification of properties of relational structures, process graphs, and such. The model checker MCK is for the verification of *knowledge* properties of agents in dynamic multi-agent systems. There are various problems and puzzles that can be implemented and verified in MCK, two of those are the "russian cards" problem and the "sum and product" problem – but there are more, and also easier start-ups. For an example, the "russian cards" problem: "From a pack of seven different cards, two players each draw three cards and the third player gets the remaining card. How can the players with three cards openly communicate each other all their cards without the third player learning from any of their cards who holds it?" This problem has been investigated in the papers cited below. We found essentially only three different solutions. These need to be implemented and verified in MCK, and apart from that there are more involved problems relating to longer protocols. This project is quite suitable for two students, who will be encouraged to collaborate but will also get individual assignments. Participation in Paper COSC462 Applied Logic is essential. The project is also related to the research project http://www.cs.otago.ac.nz/staffpriv/hans/AOARDweb.html.

H.P. van Ditmarsch. The russian cards problem. *Studia Logica*, 75:31–62, 2003.

M.H. Albert, R.E.L. Aldred, M.D. Atkinson, H.P. van Ditmarsch, and C.C. Handley. Safe communication for card players by combinatorial designs for two-step protocols. *Australasian Journal of Combinatorics*, 2005. To appear.

25. **Title: Reconstructing proofs from trees in logic programming (480/490)**
Hans van Ditmarsch (hans@cs.otago.ac.nz)

A refutation tree is a standard two-dimensional representation of resolution derivations. Given a logic program, a computation rule and a goal, different resolution derivations result from selecting different matching clauses in the program. These derivations can be visualized as separate branches in the refutation tree. This process can also be reversed: given an refutation tree where leaf nodes are labelled with either 'succeed' (for the empty clause) or 'fail', and neither other nodes nor links are labeled, one can construct a program, a computation rule and a goal that match this abstract tree. Thus we can introduce in a visual way basic logic programming concepts such as resolution strategies and finite failure. This might facilitate learning logic programming and Prolog. The goal of the project is to implement a visual tool to derive programs from refutation trees and vice versa, and to investigate some properties of programs sharing the same refutation tree-shape. Familiarity with logic programming is not essential, but recommended. This topic will be part of the second semester content of COSC462 (Applied Logic) – but has also summarily be presented in COSC343 (the Russell & Norvig textbook for COSC343 is an excellent reference). For more information, see http://www.cs.otago.ac.nz/staffpriv/hans/art2003.ps.

26. **Game theory of Pit (480/490)**
Hans van Ditmarsch (hans@cs.otago.ac.nz)

In the Pit game, between three and seven players try to corner the market on some commodity, such as barley, corn, etc. It is played with cards, and game actions consist of offers and agreements to swap cards. The first player to collect all cards of a given commodity, wins the game. This game is under continuous investigation in the Information Science department (Mariusz Nowostawski, Stephen Cranefield, Martin Purvis), partly in collaboration with me. An interface for a 12 cards version has been made by Johan Lovdahl, see http://www.ida.liu.se/~jolov/pit/. Some of its game theory can be found in the paper cited below.

For more information on Pit, see http://www.hasbro.com/common/instruct/pit.pdf. The goal of this project is to compute Nash-equilibria (optimal strategies) for a 6-card version of Pit – expanding similar results by last year's Honours student Takahiro Ito, see http://www.cs.otago.ac.nz/staffpriv/hans/taka.ps, but for different game rules. In the 6-card version, commodities have different values. Independently, a project aim is to compute equilibria for a 9-card version where all cards (commodities) have the same value. The project assumes familiarization with literature in game theory (microeconomics) – this involves independent and individual study. This project can also be done by two students, who would be encouraged to collaborate and will each get an individual assignment as well. Topics from game theory will be presented in COSC451 (Artificial Intelligence), so attendance of that paper is required.

H.P. van Ditmarsch. Some game theory of Pit. In C. Zhang, H.W. Guesgen, W.K. Yeap, editors, Proceedings of PRICAI 2004 (Eighth Pacific Rim International Conference on Artificial Intelligence), LNAI 3157, pages 946-947, Springer, 2004.

27. **Music notation software (490)**
Geoff Wyvill (geoff@cs.otago.ac.nz)
There are about twenty established commercial programs that operate as 'word processors' for musicians to prepare musical scores. Last year we ran a project to make a comparison of some of these systems and, at the same time, to discover what is important to make such a system convenient for the practising musician.

My long-term aim is to create an open source music scoring system that will outperform the commercial systems. Next, we need to examine some standard formats eg: Music XML and see what is needed to support a good GUI, ease of information exchange and high quality printing. After that, we need to set some design goals and maybe build a prototype system. The purpose of this project is to take this work to the next stage. Obviously, you have to be familiar with musical notation to work on this project.

28. **Watching Window activity (490)**
Geoff Wyvill (geoff@cs.otago.ac.nz)
The Watching Window is an experimental virtual reality display that has been built by the Graphics and Vision research group. It consists of a rear projection screen and a number of TV cameras that can track the hands and eyes of the user. We have demonstration programs that provide various different experiences. We have a simple window where the user can enjoy the view. By moving the head, you see different parts of the scene as you would through a real window. We have a space invader style game where you have to dodge rocks that are aimed at your head. And we have a painting program that lets you daub the screen with virtual paint. The idea of this project is to create an interesting activity to demonstrate the watching window. It could be a continuation of a summer research project that shows a solar-system model, a simulation where the user creates virtual smoke rings or a completely new idea.

29. **Realtime Terrain Viewer (490)**
Geoff Wyvill (geoff@cs.otago.ac.nz)
The project is to research, design and develop a real time 3D terrain viewer in conjunction with a Dunedin based technology company. The desired outcome is the development of a robust, efficient, easy to use and Internet aware terrain viewer, suitable for the downloading and visualisation of small digital elevation models and orthorectified textures. The viewer will be modelled on Adobe Acrobat Reader, in that it will be a) made available at no charge, b) hosted within a browser and c) also usable as a standalone application. As an Internet aware product, it will be data location and transport agnostic, understanding that terrain and texture datasets may be stored on a local drive, a shared server, a URL or any combination thereof, and accessed via file APIs, FTP, SOAP, HTTP, HTTPS or any combination thereof.

The viewer and associated datasets will support:

- camera flyover animations
- area selection and highlighting
- area and distance measuring
- lon/lat display and local map grid conversions
- GPS display (local and remote GPS data streams)

It is anticipated that the provision and broad uptake of a free, standard terrain viewer will significantly encourage the development of both free and commercial online DEM data stores.

30. **Where does this tune come from? (490)**
Geoff Wyvill (geoff@cs.otago.ac.nz)

From the World Wide Web you can find thousands of music samples encoded as MIDI files. The MIDI format encodes the notes and timing for an ideal performance of a song or other musical work. But unlike a sound file, MIDI files are short and precise and do not carry noisy or ambiguous information. Thus we have the possibility to write a program that will find the source of a melody or fragment of a remembered tune. The idea of this project is to devise a suitable indexing system to do this job, to implement it and to collect enough songs? from the Web to create the beginnings of a musical dictionary indexed by melody.

The project is very open ended. One obvious extension would be to investigate similar tunes and make good guesses where a melodic fragment was entered incorrectly. A basic knowledge of music notation is needed for this project.

31. **Track, Trace and Secure - RFIDs and Art Galleries (490)**
Stewart Fleming (stf@cs.otago.ac.nz)

Ensuring the security of artworks at exhibitions and in storage is a complex and expensive activity. Insurance and requirements for art gallery operators worldwide are somewhat onerous. Inventory management of artwork not currently on display can be problematic - how do you easily identify something that must be stored securely in a controlled environment? Tracking, tracing and securing artwork (not just paintings - any form of art that could be exhibited can be considered) is critically important to the effective operation of an art gallery.

This project will investigate the risks associated with the art industry and the problems faced. The project will also investigate the opportunities offered by various technological devices (concentrating on, but not exclusive to, RFIDs) in solving these problems. This is primarily a research project, although it is possible that an able student will progress into implementation of prototype systems involving custom hardware and software.

32. **Technology Infrastructure for Mobile Tourism Operators (480/490)**
Stewart Fleming (stf@cs.otago.ac.nz)

Tourism operators in Otago face interesting problems due to location and the scope and nature of their operations. Operating in often challenging physical and communication environments, they must meet a variety of standards to enhance the tourist experience, a key part of the local economy.

Building on an existing 480 project from 2004, this project aims to specify, build systems (where appropriate) and use existing infrastructure that could support tourism operators in the Otago area. This encompasses communications design (mobile data/custom applications for Blackberry devices), database design and application design/implementation. This is a systems engineering project with possibilities to go in almost any direction - if it turns out to be a pure research project to investigate wide issues, we do that; if you want to do user requirements capture, we do that; if you want to do funky communiations apps automatically synchronized to a database, we do that.

33. **mycritic.com (480/490)**
Nathan Rountree (rountree@cs.otago.ac.nz)

I love movies. But cinema tickets are expensive these days, and don't get me started on the price of popcorn. I'd really like to be able to rely on the advice of movie critics, but most of them are idiots—except for just one or two (the ones who agree with me, of course). Rottentomatoes.com collects hundreds of critical reviews for every movie that gets released, but which critics should I listen to?

"But," I hear you say, "have you tried movielens.umn.edu, which uses collaborative filtering to match your preferences with hundreds of other users who like the same things you do?" Indeed I have, and it really irks me to have to *wait* for those hundreds of people to go see the movie I'm wondering about. I want to find a critic who agrees with me most of the time, and who will write about new movies as (or before) they come out.

Your project will be to read up on collaborative filtering, and figure out how it works. Then, you'll implement a program that builds a snapshot of my preferences, and then finds the critics at rottentomatoes.com who mostly like what I like. You get to do some data mining, some internet programming and quite a lot of user interface design. This project is suitable for both 480 or 490, depending on how far you go in terms of design and development.

34. **Data mining for High Performance Sport Databases (480/490)**
Nathan Rountree (rountree@cs.otago.ac.nz)

Data mining is a set of tools and strategies for discovering patterns that may be hidden in large amounts of data. Sometimes it is used to *describe* the data in order to gain more insight, and sometimes it is used to build *predictive models* based on the current database.

I am in contact with a company that provides database programs for high-profile international sports teams. Coaches and support staff can enter and manage data concerning their athletes, from how fast they can sprint on any given day to what tactics were employed minute-to-minute in a particular game. This company now has data on thousands of athletes. They have already implemented some data mining algorithms to allow coaches to predict the performance of their players and teams. This project involves identifying a data mining technique that might be useful for this kind of database, and implementing it as a module for use in the company's product. This is a commercially-driven project, so the requirements analysis phase will be vitally important. This project is suitable for either 480 or 490, depending on the complexity and originality of the data mining technique that we choose.