WEEK 6

1. Four of the next five sessions are given over to guest lecturers. In the first year I taught this course (2002), I tried to teach all sessions myself and found that I did not know enough to do a good job on topics like the current state of the art in machine vision. MIT has a number of specialists in areas like machine learning, speech recognition, etc. who have experience in presenting their work to non-technical audiences (e.g. undergraduate economics majors) and for the second time I taught this course, I lined up four of these speakers. In subsequent years, my plan is to continue to line up about four outside speakers, though I can now take over some of this year's outside topics (based on the outside presentations) on and bring in speakers on new topics – i.e. what can robots now do.

2. Randy Davis on Expert Systems – These are rules based systems. Particular example was Mycin, an expert system he developed for diagnosing infections in the blood stream. Davis gave a fascinating early history of the development of expert systems. In the early years – 1950s, 1960s – artificial intelligence theorists had a hope of programming general problem solving strategies based on the observation that a human could make at least some progress on almost any problem they confronted. Over time, theorists began to realize that people made real progress only on those problems where they had substantial knowledge. This caused the search for general problem solving rules to be replaced by much narrower software that attempted to codify knowledge in particular domains into rules. Blood infections were a good candidate because the domain was circumscribed (as distinct from psychiatry which might cover all of life), the problem was important (don't want to give antibiotics for viruses, some blood infections can be fatal unless diagnosed quickly) and there were doctors who could, through extensive interviewing, articulate the knowledge in the area.

3. Tommi Jakkolla on Machine Learning – This is the fundamentally different processing alternative to rules based logic. To motivate the distinction, we can go back to the earlier paper describing how infants can classify an image as a cat rather than a dog. It would be hard, if not impossible, to write down, a priori, a set of If-Then-Do rules to do this. It has something to do with the various dimensions of the images – the specifics of the head must be particular important – but it is hard to say how. But since most of the students are familiar with regressions, it is possible to set up a probit or logit model where the variables are various dimensions - i.e. length of nose, spacing between eyes, etc. – and estimate the coefficients of a model that uses these dimensions to predict whether a

particular image is a cat or a dog. Once estimated one existing data, these coefficients can be used to classify new images. Using other examples, Jakkolla explains this general idea. He shows how learning software – software that recognizes patterns – must be trained on existing samples before it is used on new data. He also shows how the particular weights must be set with regard to the loss function – the relative costs of Type I and Type II misclassification errors.

4.   These two guest lectures were followed by a regular class.

   a. Summarize and work through the specifics of a simple neural net including hidden nodes, etc.

   b. Use *Radiology Today* paper to describe one application – identification of potentially cancerous cell clusters in mammograms.

   c. Spend some time on Upstairs/Downstairs paper to pull previously covered ideas together: two innovations – check images (rather than paper checks) and optical character recognition of images (neural net). Technology had two immediate implications for reorganized work. In the first, optical character recognition could substitute for people who had read and keyed onto the check the amount for which the check was written (these are the people who would manually type the amount of the check onto the bottom line of the check – those funny looking characters). In the second, the use of check images meant that persons who processed exceptions – e.g. stop payments or overdrafts – did not have to spend hours looking for particular paper checks. The bottom line was that this technology had different impacts in the two back offices – it led to substitution and a division of labor in deposit processing but it led to complementarity and a consolidation of tasks in exceptions processing.

Next two lectures on speech recognition (Jim Glass) and machine vision (Eric Grimson). Both lectures showed how underlying mechanics relied on pattern recognition, what current state of the art is, what kinds of problems remain. Both these and the previous lecture (4) introduce an idea that becomes very important in the second half the course – that information can only be processed in context. For example, in speech recognition, the word BILL could mean a piece of currency, a piece of legislation, the front end of a duck, a person's name, a request for payment, etc. We use contextual knowledge to infer the right meaning but this is very hard to program. One way to get around this problem is to limit the context so, for example, rather than try to write a general speech recognition program, we write a program for, say, a telephone company where we can assume that BILL means request for payment.