6.871 SPRING 2006 READING LIST: Installment #3

Lecture 14: *Causal Reasoning*

1. Patil, R. S., Szolovits, P., and Schwartz, W. B., Causal understanding of patient illness in medical diagnosis. *Readings in Medical Artificial Intelligence*, Clancey, W., Shortliffe, E., Addison-Wesley, 1984, pp 229–360.
   This is another approach to the diagnosis problem, which differs in substantial ways from the view supplied by Mycin's rule-based approach. Think about issues such as, What does this system know? How does it represent that knowledge? What is its approach to diagnosis?

2. Patil, R. S., Artificial Intelligence Techniques for Diagnostic Reasoning, *Exploring Artificial Intelligence*, Shrobe, H. E. (ed), Morgan Kaufmann, 1988, Chapter 9.
   This is a good overview of the diagnostic problem in all of its complexities, demonstrating that AI programs still have some distance to go in covering the problem.

3. Simmons, R, Davis R, Generate Test Debug, *International Joint Conference on Artificial Intelligence*, 1987, pp 1071–1078.
   An interesting variation on the traditional generate and test paradigm.

Lecture 16: *Model-Based Reasoning*

1. Davis, R., Diagnostic reasoning based on structure and behavior, *Artificial Intelligence Journal*, **24**: 347–410, Dec.84
   A detailed example of model-based reasoning at work, using representations of the sort described in the previous paper. It also shows how to deal with the difficult problems that arise when the structure description itself differs from the object in unknown ways.

2. Davis, R., and Hamscher, W., Model-Based Reasoning: Troubleshooting, *Exploring Artificial Intelligence*, Shrobe, H. E. (ed), Morgan Kaufmann, 1988, Chapter 8.
   A survey of the fundamental issues in model-based reasoning and suggestions about the research problems that are still open.

3. Hamscher, W., Temporally coarse representation of behavior for model-based troubleshooting of digital circuits, *International Joint Conference on Artificial Intelligence*, 1989, pp. 887–893
   A case study in how the right level of abstraction in the representation can make all the difference.

Lecture 17: *Case-Based Reasoning*

1. Kolodner, J.L. and Leake, D.B., A Tutorial Introduction to Case-Based Reasoning, in *Case-Based Reasoning: Experiences, Lessons, and Future Directions*, Leake, D.B. (ed), MIT Press, Cambridge, 1996, pp. 31–65.
   Good introduction to the topic by leaders in the field. (This chapter is adapted from Kolodner's 1993 book, *Case-Based Reasoning*.) Note the different tasks for which case-based reasoning is well suited. Skim for details; read more thoroughly if you think case-based reasoning may be appropriate for your class project.

2. Slade S., Case-Based Reasoning: A Research Paradigm, *1991 AI Magazine*, pp. 42–55.
   Good description of the general idea and history of the field. Much of the work described in this paper came out of Schank's group at Yale (e.g. work by Kolodner, Hammond) and from Kolodner's group at Georgia Tech (e.g. work by Sycara, Simpson.)

3. Hinkle D. and Toomey C, Applying case-based reasoning to manufacturing, *AI Magazine*, Spring 1995, pp. 65–73
   Good example of using case-based reasoning in the real world.

4. Koton P., A medical reasoning program that improves with experience, *1988 SCAMC Proceedings*, pp. 32–37.
   This paper is a good example of combining case-based and model-based reasoning. Why is the combination better than either alone?

5. Chi R, et. al., Generalized Case-Based Reasoning System for Portfolio Management, *1993 Expert Systems with Applications*, vol. 6, pp. 67–76.
   This paper is a good example of combining case-based and rule-based reasoning. Why did the author use both representations?

6. Porter B, Bareiss R, Holte R, Concept learning and heuristic classification in weak-theory domains, *Artificial Intelligence*, 45:229-263, 1990.
   Skim for the basic ideas, and note the differences between the this sort of reasoning and the case-based reasoning systems described in the previous three papers.

Lecture 18: *Reasoning with Constraints*

1. Paltrinieri M, On the design of constraint satisfaction problems, pp.299-311.

2. Fox, M, ISIS: A retrospective, pp. 3–28.

3. Kumar V, Algorithms for constraint-satisfaction problems: a survey, *AI Magazine*, Spring 1992, pp. 32–44.

Lecture 19: *Reasoning About Physical Systems*

1. Forbus, K. D., Qualitative Physics: Past, Present and Future. Chapter 7, *Exploring Artificial Intelligence*, Shrobe, H. E. (ed), Morgan Kaufmann, 1988.
   This is an overview of work on the area called Qualitative Physics. The earlier part of the paper is good, but there are points later on where it gets obscure. The discussion about landmark values is obscure and the section on qualitative kinematics is also hard to follow.

2. Stahovich T, Davis R, Shrobe H, Qualitative Rigid Body Mechanics, *Proc AAAI-97*, pp. 138-144.
   This is an interesting use of the notion of qualitative representations for both space and forces, producing a system with an interesting ability to simulate the behavior of devices that have been only roughly specified in a sketch.

Lecture 20: *Human intelligence project*

No readings.

Lecture 21: *Learning*

No readings.

Lecture 22: *Ontologies and Data Mining*

1. Uschold and Gruninger, Ontologies: principles, methods and applications, *The Knowledge Engineering Review*, 1996, pp. 93–136.

2. Noy and Hafner, The State of the Art in Ontology Design, *AI Magazine*, Fall 1997, pp. 53–74.
   An overview of ontologies and the state of the art in building them.

3. Hovy, E.H. and K. Knight, Motivation for Shared Ontologies: An Example from the Pangloss Collaboration. Proceedings of the Workshop on Knowledge Sharing and Information Interchange, IJCAI-93. Chambery, France.
   This paper motivates and explores an apparently inevitable problem in building large knowledge-based systems: if we are to build truly large knowledge bases, we cannot start from ground zero every time; we must be able to build on each other's work. That means we must in some fashion be able to share and re-use the knowledge bases built by others. Why is this difficult? What can we do about those difficulties?

4. Gruber and Olsen, An ontology for engineering mathematics, Stanford KSL Report 94-18, Stanford University, 1994.
   One concrete example of an ontology in (what should be) a familiar domain.