# Knowledge Engineering Methodology:

# An Annotated Bibliography of

# NEOMYCIN Research

William J. Clancey, PhD

Department of Computer Science

Knowledge Systems Laboratory

701 Welch Road, Building C

Palo Alto, CA 94304

*[This is an outline of invited lectures I presented during Second Annual Advanced Artificial Intelligence Tutorials, with extensive annotated bibliography of my publications as of 1987. This bibliography was inadvertently omitted from the published volume, Advanced Topics in Artificial Intelligence: 2nd Advanced Course, ACAI '87, Oslo, Norway, July 28 - August 7, 1987, Rolf Nossum, Editor. The citations and annotations have been updated.]*

**Abstract**

From a broad perspective, knowledge engineering is a methodology for acquiring, representing, and using qualitative models of systems. We distinguish between systems being modeled (physical, cognitive, social, etc.), modeling tasks (such as diagnosis and control), computational methods (such as heuristic classification), and implementation languages (such as rules and frames). The pragmatic value of this perspective is illustrated by uncovering knowledge representation problems in existing expert systems. New languages make explicit the dimensions of task, system model, computational method, and implementation. In state-of-the art expert systems shells, the representation of reasoning strategy is emphasized, illustrated here with examples of enhanced explanation, student modeling, and knowledge acquisition. Beyond this, we consider philosophical limitations of the representational approach and implications for future research.

## 1. Outline of the Lectures

The following is the outline for the course on knowledge engineering presented in Norway, August 1987. This outline organizes the annotated bibliography that follows.

### Lecture 1. Introduction: Describing Expert Systems

- The modeling perspective
- Examples of AI modeling tools (Clancey, 1983a, Clancey, 1987a)
- Comparison to conventional software tools (Clancey, 1985a)

### Lecture 2. Representation: Implicit Knowledge

- Alternative representations (Clancey, 1985b)
- Examples of implicit knowledge (Clancey, 1983b, Clancey, 1985c, Clancey and Buchanan, 1982)
- Abstracting for generality (Clancey, 1983c)
- Strategic reasoning in NEOMYCIN (Clancey and Letsinger, 1984, Clancey, 1984a,
- Clancey, 1984b)
- The HERACLES shell (Clancey and Bock, 1985)
- Choosing tools, problems, and representations
- Future tools

### Lectures 3 & 4. Towards a Theory of Qualitative Models

(Clancey, 1986a, Clancey, 1986b)

- General vs. situation-specific models
- Alternative views of inference
- Heuristic classification (Clancey, 1984c)

- Relating systems, models, and tasks (Clancey, 1985d)
- Comparison to decision support analysis
- Alternative methods to compute models (Clancey, 1984d)
- Classification vs. simulation models of processes
- Conceptual structure of network representations

**Lecture 5. GUIDON2: State-of-the-art Design and Research Issues**

(Clancey, 1979a, Clancey, 1979b, Clancey, 1982a, Clancey, 1982b, Clancey, 1987b)

- The case-specific model (Clancey, 1986c)
- Managing and explaining reasoning (Clancey, 1984e)
- Graphics for browsing a knowledge base (Richer and Clancey, 1985)
- Advantages of logic representation for modeling and text generation (London and Clancey, 1982, Hasling, et al., 1984, Wilkins, et al., 1987a)
- Teaching how to detect and cope with problem-solving failures (Clancey, 1984f, Clancey, 1988, Clancey, et al., 1986)
- Summary of research objectives, projects, and theories (Clancey, 1987c)

**Lectures 6 & 7. Qualitative Student Models** (Clancey, 1986b)

- Examples of general models and inference procedures (Clancey, 1987a)
- The central role of diagnosis
- Possible kinds of assessments
- Subtraction compared to medical diagnosis (Clancey, 1987d)
- Types of process models of reasoning (Clancey and Shortliffe, 1984)
- Levels of abstraction: from behavioral to functional model
- Computational representations of processes (Clancey, 1979c)
- Classification vs. simulation models of bugs
- Arguments for functional models
- Overlay vs. Bug models
- Computational methods for constructing models
- Learning, communication, other tutoring topics
- Reinterpreting "Qualitative Reasoning" in terms of knowledge bases

**Lecture 8. Cognition and Expertise: Philosophical Considerations**

- What is the knowledge acquisition bottleneck?
- How does a knowledge base relate to what experts know? (Clancey, 1983d, Clancey, 1987e). Examples from NEOMYCIN (Clancey, 1986d)

- Knowledge engineering methodology (revisited)

## 2. Annotated Bibliography

Clancey, W. J. 1979. Tutoring rules for guiding a case method dialogue. *The International Journal of Man-Machine Studies*, 11, 25–49.
*[The first paper on* GUIDON, *describes the program in terms of the types of tutoring rules and the architecture of procedures and rule sets for invoking them. A simple, useful introduction, though it lacks any sense of the limitations discovered later. (Republished in Sleeman and Brown (eds), Intelligent Tutoring Systems, Academic Press, 1982)]*

Clancey, W. J. 1979. *Transfer of rule-based expertise through a tutorial dialogue.* PhD thesis, Computer Science Department, Stanford University, August.
*[The title of my dissertation is noteworthy because it reveals the emphases of the late 1970's—that expertise is something that is transferred like a commodity between people, that rules provide a complete representation language, and that a dialogue (without graphics) is an adequate means of communicating. Later work has deviated from each of these conceptions. Knowledge-Based Tutoring (see below) is an extensive revision of this dissertation in a form suitable for people interested in replicating and benefiting from* GUIDON*'s simple, paradigmatic design.]*

Clancey, W. J. 1979. Dialogue management for rule-based tutorials. In *Proceedings of the Sixth International Joint Conference on Artificial Intelligence* (pp. 155–161). Tokyo, August.
*[Presents tutoring a discourse problem and describes* GUIDON*'s knowledge in terms of a transition network of dialogue situations and rationales for transitions. My first paper with a simple theme and title, and a diagram for getting the ideas across.]*

Clancey, W. J. 1982. GUIDON. In A. Barr and E. A. Feigenbaum (Eds.), *The Handbook of Artificial Intelligence,* Applications-oriented AI research: Education (pp. 267–278). William Kaufmann, Inc., Los Altos.
*[The most concise description of the* GUIDON *program. Reprinted with revisions in Computer-Based Instruction and presented at the Society for Computer-Applications to Medical Care Conference, in Washington, DC, November, 1979. Provides examples of tutoring rules, adapting material from chapter 1 and the background material on medical computer-aided instruction in my dissertation.]*

Clancey, W. J., Bennett, J. S., and Cohen, P. R. 1982. Applications-oriented AI research: Education. In Barr and Feigenbaum (Eds.), *The Handbook of Artificial Intelligence* (pp. 223–294). William Kaufmann, Inc., Los Altos.
*[This chapter provides a concise description of early ICAI research. Every AI researcher should be familiar with programs like* SCHOLAR *and* SOPHIE, *which are described here.]*

Clancey, W. J. 1983. Communication, simulation, and intelligent agents: Implications of personal intelligent machines for medical education. In *Proceedings of AAMSI-83* (pp. 556–560).
*[A deliberately speculative paper, relates new technologies to some specific requirements for medical training. My aim was to capture the interest of people unfamiliar with early 1980's research.]*

Clancey, W. J. 1983. The epistemology of a rule-based expert system: A framework for explanation. *Artificial Intelligence* 20 (3), 215–251.

*[The definitive study of* MYCIN*'s rules, adapted from a chapter of my dissertation, with much help from John Seely Brown and Danny Bobrow and many other people over a period of three years. Describes the kinds of knowledge relating to production rules in terms of three roles: strategy, structure, and support. Strategy and structure are shown to be implicitly encoded in clause ordering, while support is not represented at all. Davis's metarules are reconsidered, and the* NEOMYCIN *architecture of tasks and metarules is introduced. The framework is applied to describe several diverse programs, including* HEARSAY *and* AM. *Reprinted in Knowledge-Based Tutoring.]*

Clancey, W. J. 1983. The advantages of abstract control knowledge in expert system design. In *Proceedings of the National Conference on Artificial Intelligence* (pp. 74–78). Washington, D.C., August.

*[Describes how control knowledge can be abstracted from domain -specific production rules. Domain knowledge is represented by explicit relations and control rules use variables rather than domain terms. The idea of structural relations introduced in the epistemology paper is thus made more precise through there implementation of* NEOMYCIN*'s metarules in predicate calculus. In my opinion, this is the best technical paper I have written to date.]*

Clancey, W. J. 1983. Review of Jonathan Miller's "State of Mind". *AI Magazine* 4 (4), 61–66.

*[Miller's brilliant interviews of a dozen brilliant minds; includes the most profound and far-reaching cognitive science ideas I have seen in a single book]*

Clancey, W. J. 1984. Methodology for Building an Intelligent Tutoring System. In Kintsch, Miller, and Polson (Eds.), *Methods and Tactics in Cognitive Science* (pp. 51–83). Hillsdale, N.J.: Lawrence Erlbaum Associates.

*[Reflects on the cognitive principles behind* NEOMYCIN's *design, considers how the theory is related to the implemented program, and details the methods for interacting with the expert medical teacher, Tim Beckett. A final section examines pitfalls in the process of constructing such a model and speculates on how the research might proceed more effectively by interdisciplinary collaboration. Includes a glossary. Reprinted in G. P. Kearsley (Ed.), Artificial Intelligence and Instruction: Application and Methods* (pp. 193–227). *Reading: Addison-Wesley, 1987.]*

Clancey W. J. 1984. Acquiring, representing, and evaluating a competence model of diagnostic strategy. In M. Chi, R. Glaser, and M. Farr (Eds.), *The Nature of Expertise* (pp. 343–418). Hillsdale, New Jersey: Lawrence Erlbaum Associates.

*[Presented at a conference at the Learning Research Development Center, University of Pittsburgh, October 1984. Describes the content of* NEOMYCIN*'s diagnostic procedure, emphasizing how it was developed by protocol analysis. Annotated teaching dialogues illustrate the use of abstract descriptions of strategy in explanations. An appendix lists the diagnostic tasks and metarules in English, with examples. Also examines the social, mathematical, cognitive, and case-population constraints that rationalize this diagnostic procedure. Written for a broad audience, this paper describes the medical psychology aspects of the research.]*

Clancey, W. J. 1984. Classification problem solving. In *Proceedings of the National Conference on AI* (pp. 49–55). Austin, TX, August.

*[Introduces a framework for describing the chains of inference that occur in heuristic reasoning. Diverse programs are described through a sequence of diagrams, and implications for language design are considered. Nominated for the Publishers' Prize.]*

Clancey, W. J. 1984. Details of the revised therapy algorithm. In B.G. Buchanan and E.H. Shortliffe (Eds.), *Rule-based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project* (pp. 133–146). Addison-Wesley.

*[After deciding to develop a computer-aided instruction program from* MYCIN, *I believed I had to re-represent the therapy program (coded in* LISP*) in order to teach it to students. I learned that linear production rules were poorly suited for representing a recursive procedure, and settled for a re-implementation (with the extensive help of Bill van Melle) in the form of a generate and test program. The program's description, which appeared on one page in the Proceedings of IJCAI-77 (my first single-author publication), is expanded upon here. The algorithm was later studied and reformulated (on paper) by derivation from constraints by Mostow and Swartout in their AAAI-86 article.]*

Clancey, W. J. 1984. Teaching classification problem solving. In *Proceedings of the Sixth Annual Conference of the Cognitive Science Society* (pp. 44–46). Boulder.

*[Interim research report, describing the teaching programs being developed on top of* NEOMYCIN. *Presents a somewhat speculative framework of what needs to be taught and novel ways of presenting it.]*

Clancey, W. J.1984. Knowledge acquisition for classification expert systems. In *Proceedings of ACM Annual Conference* (pp. 11–14). October.

*[Based on the heuristic classification model, describes a plan for interviewing experts that respects the kinds of knowledge to be organized and human preferences for explaining reasoning. Presented at the session that won the "Best Session of the Conference" award. Also in A. Gupta and B.E. Prasad (Eds.),* Expert Systems—Principles, IEEE Press, 1988.]

Clancey, W. J. 1985. Software tools for developing expert systems. In *Proceedings of the International Conference on Artificial Intelligence in Medicine* (pp. 155–178). Amsterdam.

*[Combines the introduction to Heuristic Classification with new material describing the tools developed for* EMYCIN.*]*

Clancey, W. J. 1985. Review of Sowa's *Conceptual Structures: Information Processing in Mind and Machine*. *Artificial Intelligence*, 27 (1), 113–124.

*[Sowa's book enabled me to relate knowledge representation research to the logicists and heuristic programming. Every AI researcher should read this book. The Journal recently printed a second perspective arguing that the ideas are out-of-date (S. W. Smoliar, Artificial Intelligence, 33, 2, 259–266), but Sowa's response (Artificial Intelligence, 34, 3, 388–394) defended the strength and necessity for the synthesis. He says, "Pierce was too good a logician to commit the kinds of fallacies that Brachman describes."]*

Clancey, W. J. 1985. Review of Weiss and Kulikowski's *A Practical Guide to Designing Expert Systems*. *AI Magazine* 5 (4), 84–86.

*[Knowledge-Based Tutoring has a similar goal of making simple ideas available to a broad audience; however, this "practical guide" provides very little perspective on what simple methods will allow and what difficulties developers will encounter. Weiss and Kulikowski's work on* CASNET *has been tremendously influential, but the* EXPERT *shell described in their book does not incorporate their insights about casual reasoning.]*

Clancey, W. J. 1985. Heuristic Classification. *Artificial Intelligence* 27, 289–350.

*[Expansion of the original AAAI paper, including a new section on kinds of tasks (e.g., diagnosis, design), as well as a much more detailed consideration of the continuum between selecting solutions and piecing them together from primitive components. The title emphasizes the difference from classification in most earlier psychological research, as well as making clear that classification is not a kind of problem, but a method for solving any problem. Reprinted in J.S. Kowalik ( ed. ),* Knowledge-Based Problem Solving*, New Jersey: Prentice-Hall, 1985.]*

Clancey, W. J. 1989. Viewing knowledge bases as qualitative models. *IEEE/Expert* 4 (2), 9–23. KSL Report 86-27, Stanford University, 1986.

*[Argues that artificial intelligence research is distinguished from traditional programming by its methodology for modeling processes non-numerically, by relational networks. The paper is much too short for what it portends, but readers of "Heuristic Classification" and "Qualitative Student Models" will recognize the first steps in a necessary synthesis. I expect much of my future research will be built on the framework described here. A final, wide-ranging essay relates AI to General Systems Theory.]*

Clancey, W. J. 1986. Qualitative student models. In Traub, J.F. (editor), *Annual Review of Computer Science* (pp. 381–450). Annual Reviews, Inc., Palo Alto.

*[Applies the framework presented in KSL-86-27 to describe the varieties and capabilities of student modeling programs. In particular, reasoning is viewed in terms of a general model and an inference procedure for constructing a model of a system's behavior in a particular situation. Written for a general computer science audience; the AI specialist should skip to the sections on process models and inference methods.]*

Clancey, W. J. 1986. From Guidon to Neomycin and Heracles in twenty short lessons (ONR Final Report 1979–1985). *AI Magazine* 7 (3), 40–60.

*[The definitive overview describing my research from soup to nuts, also the most fun paper to write, with my favorite figures illustrating each of the twenty sections. The final section reflects on the sequence of research, characterizing it in terms of reformulation of representation languages. Reprinted in A. van Lamsweerde and P. Dufour (Eds.), Current Issues in Expert Systems, pps. 79–123, London: Academic Press, 1987.]*

Clancey, W. J. 1986. *Know-how vs. knowledge representation,* in *Proceedings of the Workshop on Knowledge Compilation* (pp. 1–2). Inn at Otter Crest.

*[Concise application of the Winograd and Flores book to* NEOMYCIN *research. Examples pointing to other papers support the claim that knowledge is not just transferred from experts to program representations, but rather for some applications it is conceptualized for the first time during the knowledge acquisition and representation process.]*

Clancey, W. J. 1987. Intelligent tutoring systems: A tutorial survey. In A. van Lamsweerde and P. Dufour (Eds.), *Current Issues in Expert Systems* (pp. 39–78). Academic Press, London, 1987.

*[Revision of a talk given as part of the IBM-Belgian Computer Science Professorship series, May 1985. The examples are mostly from a tutorial presented with Brown and Burton at AAAI-84. Descriptions of programs like* MENO, SOPHIE, *and the* MACSYMA ADVISOR *are not generally available elsewhere.]*

Clancey, W. J. *Knowledge-Based Tutoring: The Guidon Program.* Cambridge, MA: MIT Press 1987.

*[An extensive revision of my 1979 dissertation, with a preface and new final chapter bringing the work up to date. Includes new material on the limitation of knowledge-based tutoring, deriving from the Winograd and Flores review. The "epistemology" chapter is included, while its associated material on "knowledge base structures" from chapter 5 of the dissertation has been made a separate chapter. An annotated dialogue appears as an appendix and is used throughout to illustrate types of tutoring rules. The overlay student model and its limitations, the 19 question types for quizzing the student about* MYCIN*'s rules, plus experiments with* PUFF *and* SACON *have not been described elsewhere. Indexed.]*

Clancey W. J. 1987. Diagnosis, Teaching, and Learning: An Overview of GUIDON2 Research. In *Proceedings of AI-87,* Osaka, Japan.

*[Short summary of accomplishments, active research, and plans. Describes how the research might be shifted to engineering problems.]*

Clancey, W. J. 1987. Functional principles and situated problem solving. *Behavioral and Brain Sciences,* 10 (3), 479–480.

*[Commentary on John Anderson's "Methodologies for studying human knowledge," Behavioral and Brain Sciences. Argues that principles, in the form of process abstractions, are less important in mathematical domains explored by Anderson (geometry, basic lsisp programming) because the abstractions of interest (as discovered in medical and engineering applications) are about systems in the world. That is abstractions capture recurrence in processes in the world and processes for reasoning about those systems. Anderson is studying modeling tools out of context—separated from the situations in the world they are used to model.]*

Clancey, W. J. 1987. Review of Winograd and Flores's *Understanding Computers and Cognition*. *Artificial Intelligence* 31 (2), 232–250.

*[Restates and illustrates the key points of Winograd and Flores's book. The review is in the form of five positive statements, followed by five questions of serious, unresolved problems.]*

Clancey, W. J. 1988. The knowledge engineer as student: Metacognitive bases for asking good questions. In H. Mandl and A. Lesgold (Eds.), *Learning Issues in Intelligent Tutoring Systems* (pp. 80–113). New York: Springer-Verlag.

*[Develops a new model of learning relevant to the design of tutoring systems. Synthesizes diverse examples from previous papers, arguing that knowledge engineer's knowledge about types of knowledge and inference procedures serves as a framework for directing her own learning. She experimentally constructs an expert system by directing her teacher, the expert. Also describes the design of a knowledge acquisition program (*GUIDON-DEBUG*) based on debugging an improperly-formed solution.]*

Clancey, W. J. and Bock, C. 1988. Representing control knowledge as abstract tasks and metarules. In L. Bolc and M. J. Coombs (Eds.), *Expert System Applications* (pp. 1–78). Berlin: Springer-Verlag.

*[This complicated paper combines the AAAI-83 introduction with a detailed description of the* NEOMYCIN *architecture and its reimplementation in MRS. The best part of this paper is lengthy analysis of the nature of abstract control knowledge, particularly the difficulty of making its rationale explicit. We argue that the relational language of the diagnostic procedure defines a classification of domain concepts, which by definition are those useful for practical, routine problem solving. Bock's experiment in translating the task interpreter into a linear sequence of production rules demonstrates that "declarative representation" does not mean humanly readable.]*

Clancey, W. J. and Buchanan, B. G. 1982. *Exploration of problem-solving and tutoring strategies: 1979–1982.* Technical Report STAN-CS-82-910, HPP 82-8, Stanford University, 1982.

*[Summarizes the main results of the period of research immediately following the development of* GUIDON*, leading to* NEOMYCIN.*]*

Clancey, W. J. and Letsinger, R. 1981. NEOMYCIN: Reconfiguring a rule-based expert system for application to teaching. In Clancey, W. J. and Shortliffe, E. H. (Eds.), *Readings in Medical Artificial Intelligence: The First Decade* (pp. 361–381). Addison-Wesley, Reading.

*[First paper about* NEOMYCIN, *describes the rationale and architecture, with an annotated example. Appeared originally in Proceedings of IJCAI-81.]*

Clancey, W. J. and Shortliffe, E. H. 1984. *Readings in Medical Artificial Intelligence: The First Decade.* Reading: Addison-Wesley.

*[A collection of the medical AI papers that most influenced me in the development of* NEOMYCIN. *The papers by Szolovits, Feltovich, and Patil are classics that everyone interested in the modeling of problem solving should read.]*

Clancey, W. J., Richer, M., Wilkins, D., Barnhouse, S., Kapsner, C., Leserman, D., Macias, J., Merchant, A., Rodolitz, N. 1986. *Guidon-DEBUG: The student as knowledge engineer.* KSL Working Paper 86–34, Stanford University.

*[Working paper that describes a prototype programming environment that supports knowledge acquisition. A scenario illustrates how this debugging environment is used by a student, rather than a knowledge engineer.]*

Feigenbaum, E. A. and Clancey, W. J. 1981. Knowledge engineering, objectives and direction. *Mathematical Sciences* 4, 11–20.

*[Printed in Japanese. Combines a description of GUIDON research with Feigenbaum's thoughts about research in the 1980's. Primarily of historical interest. Includes a black and white photo of me playing the piano.]*

Hasling, D. W., Clancey, W. J., Rennels, G. R. 1984. Strategic explanations in consultation. *The International Journal of Man-Machine Studies* 20 (1), 3–19.

*[Describes the raison d'etre of* NEOMYCIN*—how explanations of diagnostic strategy are provided by the task/metarule architecture. Compares this design to* MYCIN*'s Reprinted in M.J. Coombs (Ed.),* Developments in Expert Systems *(pp. 117–133). London: Academic Press.]*

London, B. and Clancey, W. J. 1982. Plan recognition strategies in student modeling: prediction and description. *Proceedings of the 2nd American Association for Artificial Intelligence,* (pp. 335–338).

*[London describes his masters project, a program that uses* NEOMYCIN*'s tasks and metarules to parse a sequence of student requests for patient data. The program called* IMAGE, *combines top-down prediction with bottom-up descriptions of student behavior. London evaluated the program's performance in his PhD Dissertation.]*

Richer, M. H. and Clancey, W. J. 1985. GUIDON-WATCH: A graphic interface for viewing a knowledge-based system. *IEEE Computer Graphics and Applications* 5 (11), 51–64.

*[*GUIDON-WATCH *is a graphic interface for* NEOMYCIN*, consisting of several dozen windows and menus, all illustrated in this paper. This might also be the best introduction to* NEOMYCIN*'s architecture. Reprinted in Lawler, R.W. and Yazdani, M. (Eds.),* Artificial Intelligence and Education*, Volume One, pp. 373–412, New Jersey: Ablex, 1987.]*

Scott, A. C. and Clancey W. J. and Davis, R. and Shortliffe, E. H. 1977. Explanation Capabilities of Production-Based Consultation Systems. *American Journal of Computational Linguistics*, microfiche (62).

*[This early description of MYCIN's question-answer program is still useful for its taxonomy of question type and clear description of a methodology for building a classification-based QA program. Revised: Methods for Generating Explanations. (1984). In B. G. Buchanan and T. H. Shortliffe (Eds.),* Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project *(pp. 338–362). Reading, MA: Addison-Wesley.]*

Thompson, T. and Clancey W. J. 1986. A qualitative modeling shell for process diagnoses. *IEEE Software* 3 (2), 6–15.

*[Describes* CASTER, *a program for diagnosing faults in iron cast in sand molds (sandcasting). Comparing* HERACLES' *use for sandcasting and medical diagnosis, presents a generalization in terms of the modeling and diagnosing of abnormal processes. The knowledge acquisition phase of* CASTER'*s design is emphasized, particularly the value of a "structure" and "process" perspective for constructing a casual network of faults. (This paper was published with black and white figures; the original technical report version has color figures and/or improved captions.)]*

Wilkins, D. C., Buchanan, B. G., and Clancey, W. J. 1984. *Inferring an expert's reasoning by watching,* in *Proceedings of the 1984 Conference on Intelligent Systems and Machines*.

Wilkins, D. C., Clancey, W. J. and Buchanan, B. G. 1986. An overview of the Odysseus learning apprentice. In T. M. Mitchell, J. G. Carbonell, and R. S. Michalski (Eds.), *Machine Learning: A Guide to Current Research* (pp. 369–373). Academic Press, New York.

*[An excellent, accessible article describing* ODYSSEUS*.]*

Wilkins, D. C., Clancey, W. J., and Buchanan, B. G. 1987. Knowledge Base Refinement By Monitoring Abstract Control Knowledge. In B. R. Gaines and J. H. Boose (Eds.), *Knowledge Acquisition for Knowledge-Based Systems* (pp. 183–195). Orlando: Academic Press. Also: *Int J Man-Machine Studies* 27 (3), 281–293, September 1987.

*[The most mature of the papers on* ODYSSEUS, *builds on the distinctions introduced in my AAAI-83 paper to describe how the program learns new propositions within* NEOMYCIN*'s relational language by reasoning about failed metarules.]*

Wilkins, D. C., Clancey, W. J. and Buchanan, B. G. 1987. Using and Evaluating Differential Modeling in Intelligent Tutoring and Apprentice Learning Systems. In J. Psotka, L. Massey, and S. Mutter (Eds.), *Intelligent Tutoring Systems: Lessons Learned* (pp. 257–279). Lawrence Erlbaum Publishers, New York.

*[Presents a clear framework for defining what a differential model is how it can be evaluated.]*

Wilkins, D. C., Buchanan, B. G. and Clancey, W. J. 1987. *The Global Credit Assignment Problem and Apprenticeship Learning.* KSL Report 87-04, Stanford University.

*[Deals with the problem of verifying that a change to the knowledge base is plausible with respect to a general domain model, in this case using induction from a case library.]*