The Knowledge Engineering Review, Vol. 11:1, 1996, 57–66

Computer science research on scientific discovery

RAÚL E. VALDÉS-PÉREZ

Computer Science Department and Center for Light Microscope Imaging and Biotechnology, Carnegie Mellon University, Pittsburgh, PA 15213, USA

Abstract

This article is an essay on directions and methodology in computer-science oriented research on scientific discovery. The essay starts by reviewing briefly some of the history of computing in scientific reasoning, and some of the results and impact that have been achieved. The remainder analyses some of the goals of this field, its relations with sister fields, and the practical applications of this analysis to evaluating research quality, reviewing, and methodology. An earlier review in this journal (Kocabas 1991b) analysed scientific discovery programs in terms of their designs, achievements and shortcomings; the focus here is research directions, evaluation and methodology, all from the viewpoint of computer science.

1 History of research on scientific discovery

The early days of artificial intelligence saw various attempts to automate creative tasks of scientific and mathematical inference. Perhaps the earliest examples (on electronic computers) of symbolic mathematical or scientific inference were master's theses at MIT (JF Nolan) and at Temple (HG Kahrimanian) in 1953 on analytical differentiation in the calculus (Hearn, 1990). Soon after came the logic theorist, whose designers (A Newell and HA Simon) submitted in 1958 an improved proof discovered by the program to the *Journal of Symbolic Logic* (Crevier, 1993). At around the same time, Gelernter et al. (1963) created the Geometry Theorem Prover. Starting in the 1960s, Lederberg invented an algorithm for generating molecular structures efficiently, which led to the Stanford Dendral project whose goal was to elucidate molecular structure on the basis of mass spectrograms and other experimental evidence (Lindsay et al., 1995). These are some of the early events in the application of computers to problems of creative scientific and mathematical inference.

A milestone in this field was the work of the 1980s at Carnegie Mellon culminating in the book *Scientific Discovery* by Langley, Simon, Bradshaw and Zytkow (1987), who described a number of programs, implemented as production systems, that provided plausible accounts of a number of discoveries mostly from the history of chemistry. This seminal book also examined in detail the philosophical implications of the proposition that human discovery can be construed as heuristic search in problem spaces.

The momentum generated by this body of work resulted in a symposium at Stanford in 1989, which led to the collection *Computational Models of Scientific Discovery and Theory Formation* by Shrager and Langley (1990). The emphasis there, as in the previous book by Langley et al., was on *computational models* rather than what we might call *computational actors* in scientific discovery, i.e., programs that are intended to take part in current scientific practice. Following in this tradition were special issues on discovery in the journal *Machine Learning* by Langley and Michalski (1986) and Zytkow (1993).

2 Recent events

In March 1995 I chaired a AAAI Spring Symposium at Stanford University on *Systematic Methods of Scientific Discovery*. The other symposium organizers were L Darden, J Lederberg, HA Simon and D Sleeman. As reflected by the title, a goal of the symposium was to encourage an emphasis on *computational actors,* i.e., programs that could take part in the scientific process, independently of whether they were, strictly speaking, a model of anything in particular. In line with this goal, we invited working scientists who do not ordinarily take part in AI events, such as the chemist Hendrickson who has developed the SYNGEN program for organic synthesis in chemistry (Hendrickson and Tocsko, 1989), and the earth scientist Oliver, who has written a non-computational book on scientific discovery that AI-oriented workers would find very congenial in its attention to heuristic methods in science (Oliver, 1991).

The spring Symposiasts included a diverse mix of computer scientists, psychologists, philosophers and natural scientists united by an interest in more-systematic methods for human and computer scientific inference. A forthcoming report in *AI Magazine* has more observations on the symposium, as well as instructions for ordering the working notes. Also, prompted by the events at the symposium, a call-for-papers has been announced for a special issue on "Scientific Discovery" in the journal *Artificial Intelligence* (submission deadline of 1 November 1995).

3 Research results

A field with as long a history as computational scientific discovery should have produced visible results, unless its workers are on a wrong track. There have been at least two research schools: one concerned with building models of human performance in science, and another concerned with building programs that are effective players in science and which are free to use methods that are implausible in human terms, such as very large sequential searches. It is clear that the results achieved by these two research schools will differ. I will first mention briefly some substantial achievements that have affected basic theory in AI and beyond, and then some concrete achievements in domain sciences brought about by machine discovery[1].

First, the early work on logic theorist led directly to the basic theory of heuristic search in problem spaces, which emerged after analysing the behaviour of the program and other contemporary programs that demonstrated successful problem solving (Newell and Simon, 1976). Second, the Dendral project resulted in the insight that the source of much problem-solving performance was specific knowledge about a domain (Lindsay et al., 1980), contrary to the prevailing emphasis at the time on clever inference techniques which downplayed specific domain knowledge. This insight led subsequently to the expert-systems movement at Stanford and elsewhere. A third example is the wide impact that the Carnegie Mellon discovery programs, described in the book *Scientific Discovery* (Langley et al., 1987), has had in furthering a cognitive-science approach to many of the lasting problems in the philosophy of science.

I now turn to examples of specific discoveries that have been published in domain science journals, typically by those concerned with *machine* discovery rather than with modelling the human psychology of discovery. An early instance was the publication of rules for the fragmentation processes of several classes of ketoandrostanes in chemistry by the Meta-Dendral program (Buchanan et al., 1976). Some more recent examples include conjectures in graph theory contributed by the Graffiti program authored by the mathematician S. Fajtlowicz at the University of Houston, one of which was later proved and published in the *Journal of Graph Theory*. From my

[1] I will not make much effort to trace the different goals that motivated workers on scientific discovery. (A recent historical account of theorem proving (Mackenzie, 1995) does follow this approach, and its examples overlap slightly with mine.) Instead, although acknowledging the diversity of goals, the discussion of the various research lines will be centred on how they feed the concerns of automated discovery.

own work, an article (Valdés-Pérez, 1995d) in *AI Magazine* analyses three human/computer discoveries in biology, chemistry and physics; the best example therein of a machine discovery was enabled by the chemistry program MECHEM (Valdés-Pérez, 1995b, 1994). Finally, the TETRAD project (Spirtes et al., 1993), carried out by philosophers at Carnegie Mellon, has automated the inference of linear causal models from observational data; at the Spring Symposium, C Glymour described cases in which the program has improved on published causal models.

This brief overview shows that research on scientific discovery has had significant impact on basic qualitative theory in artificial intelligence and on peripheral fields such as the philosophy of science, and has led to published discoveries in domain sciences. Concerning the latter issue of published discoveries, there is still room for a program that "blows the lid off" some significant scientific problem, to borrow a phrase used by EA Feigenbaum at the Stanford Spring Symposium. Nevertheless, there are results which can justify a claim on one of the "grand challenges" proposed by Carbonell in a 1992 editorial in *Machine Learning* (Carbonell, 1992). Elsewhere I have made further observations on the future prospects of, and obstacles to, machine discovery in practice (Valdés-Pérez, 1995c).

4 The phenomena of scientific inference

Research on scientific discovery should be driven by empirical phenomena, i.e., by the inferential tasks that arise in scientific practice. Of course, these inferential tasks are not static, but evolve over time. Nevertheless, research should demonstrate a degree of actual connection to real problems of scientific inference.

In this sense, the field of scientific discovery is analogous to other empirical subfields of artificial intelligence such as robotics, vision, natural language processing, speech, and so on, which are motivated by actual problems of a robot or program sensing the world or handling human speech or written language. The first question to ask about a piece of vision research is what is learned about the problem of sensing and acting within a visual environment. Similarly, the first question to demand of scientific discovery research is what practical problem of scientific inference is being elucidated.

The strong emphasis on generality that is often found in machine learning, or generally in AI, should be a secondary concern in the field of scientific discovery. Rather than being a meta-physical constraint on research projects, generality should emerge from the practical business of automating tasks of scientific inference. Propositions or principles of a limited, but useful, degree of generality should not be disdained because they fall short of universal scope. To do otherwise is to disregard the incremental character of science, and to hamper the field by postponing useful and illuminating partial results until the advent of a universal theory of everything, which may or may not arrive.

5 Computer science research

Here I will examine computer science (CS) research on scientific discovery, by which I mean research that computer scientists are in a special position to carry out. Of course, the intention is not to exclude others from this activity, but to describe and explicate a research direction whose values and methods fall squarely within the tradition of research in computer science.

5.1 Evaluation of research quality

How is computer science research on scientific discovery to be evaluated? A typical research contribution consists of building a program that newly automates (perhaps partially) some task of scientific reasoning, or that traverses an interesting distance in that direction. The reality and significance of the program's task are often in question, since computer science reviewers cannot be expected to know enough about the many scientific domains that may be explored by discovery research. A program's novelty is another issue, since domain scientists are themselves writing creative and complex computer programs in support of their activities, which may not be known to a computer scientist.

An obvious means to acquire positive evidence for reality, significance and novelty is, therefore, to publish the work in a journal of the relevant scientific domain. Of course, this requires the flexibility to write for a different audience, and will also imply some overlap with publication in discovery forums, but the overall benefits are clear: validation of the work, impact on the science, external visibility for discovery research, and the chance to gain new sources of collaborators, or new data to run a program on. The Stanford DENDRAL project set the standard for co-publication in computer and natural science (Lindsay et al., 1980).

By no means does such a practice imply an ethically dubious re-publication of the same material, since the distinct audiences will want to read rather different descriptions, and the link to precursor literature will also be different. Ideally, one publication is enough to inform the world of a contribution, but the pragmatics are different and will likely remain so.

The issue of *novel scope* is crucial, since it is quite possible to write useful and worthwhile programs in support of scientific discovery without demonstrating new scope for computing. For example, in fields where computers are already major players, the contribution of a new program will, by itself, add little to our understanding of what in science can be automated. There are certain areas within molecular biology and symbolic mathematics, for example, where computer programs are already dominant. Therefore, to demonstrate novel scope in molecular biology, one would need to tackle an unexplored type of inference that arises there, rather than demonstrate an incremental improvement over existing programs[2]. The property of novel scope generally correlates with surprisingness: if a mixed audience of informed computer and domain scientists is generally surprised by the fact of a discovery program and its achievements, then the program displays novel scope for computing.

5.2 How NOT to judge quality

This section will examine some fallacious judgments about discovery research which have been gleaned over the years from a variety of sources, including reviews of work by this author and his colleagues.

Frequently, the standard review forms for publication outlets in fields of AI assume that submissions are contributing a new general method, from which it follows that the submitter must give evidence why the method is an improvement over previous methods, why it is domain independent, and so on. This pigeon-holing of research contributions must not be the standard for discovery research. If it were, it would be highly ironic, since an acquaintance with the practice of science teaches that a variety of types of contribution are capable of advancing knowledge of a field, ranging from new methods and techniques, to formal and qualitative theory, informed conjecture, new positive or negative evidence about previous proposals, success stories, syntheses and reviews. At this early stage (historically speaking) of discovery research, no standard research schemata should be legislated.

If the methods that are used in a discovery program have precedents, and thus are not invented specifically for their use in the program, this does not imply that the work is not novel or significant; these are independent issues. Finding a match between scientific problems and existing methods is a valuable contribution. For example, a discovery program may represent a "clever use of heuristic search", and heuristic search is a familiar method of AI, but this fact does not detract from the contribution, which lies in identifying new scope for computing in discovery. Quite the contrary, the successful use of simple, familiar methods such as heuristic search is desirable, since it points to

[2]Of course, the incremental improvement may be relevant to a broader issue. My point assumes the absence of this separate motivation.

a desirable pattern of simplicity in the phenomena of interest: scientific inference. Conversely, novel methods that have an unknown link to actual science should seek other outlets for publication.

It is not a devastating criticism to point out that a discovery program is not completely autonomous, just as it is not a devastating criticism of the co-author of a scientific paper to point out that the co-author had co-workers. The wider the autonomy of the program, the better, but the realistic goal of achieving at least incremental progress demands that we do not set our standards too high.

Given that science is an incremental process, it is always the case that a research contribution presupposes some already-solved problems. For example, the work of most AI scientists presupposes the existence of computers, whose basic design was worked out in mid-century. Likewise, it is irrelevant to accuse discovery research of relying on previously-solved problems, e.g., of a program not deciding by itself what are the possibly relevant variables or data, since all science is similarly reliant on the work of predecessors. The key question is whether one identifies new scope for computing, or otherwise elucidates some incompletely-understood question about discovery. There are no "essences" in science that foolhardy discovery researchers pare away, there is only the issues of whether a specific problem of scientific inference is real, poorly understood, and hence interesting to automate.

From a computer-science standpoint that is concerned with automation, how scientists accomplish a given task is logically irrelevant, although possibly of great heuristic value. A valid approach to automation can and has in many cases been based on a logical analysis of the heuristic relation between data and discovery, and not on any psychological analysis.

Finally, it is too much to ask of every discovery research that it contribute a new discovery in science. There are many circumstantial obstacles to making a machine discovery, such as gaining access to promising data or even any data, and these need not be overcome as a prerequisite to publication; the justification is, again, the incremental character of scientific research in general. However, one can and should inquire about such results periodically, and I have given a brief but significant answer above. Another objection against the lofty criterion of requiring scientific discoveries is that it would rule out theoretical, conceptual and psycho-historical work, which would be most undesirable.

6 Related fields

Computer science research on scientific discovery is related to, and influenced by, various other independent, intersecting, or superordinate fields. Below I propose some ways to construe this relation.

6.1 Relation to machine learning

Research on scientific discovery is often considered a subfield of machine learning. Some of the reasons for this were discussed in editorials by Langley (1986) and by Langley and Michalski (1986) in the journal *Machine Learning.* Among their reasons were that many discovery researchers have also worked in more mainstream machine learning, and that both fields share a concern with the formulation of general law-like rules from data.

Their points are accurate, but a difficulty is that the inferential problems that arise in scientific practice include much more than the data-driven induction of rules and laws, which has been a major focus of ML research. Moreover, ML researchers have been quite concerned with generality, and less concerned with learning or discovery problems which require a substantial dose of domain-dependent elements to achieve a competent system. Scientific problems, on the other hand, abound in domain-dependent aspects, hence the tension. While generality, efficacy, elegance and so on are universal goods, the persistent reality is that science is an incremental process, and one has to choose what aspects will receive emphasis. The argument here is that the emphasis in scientific discovery research should be on efficacy and novel scope first and generality second.[3]

A further issue is that much scientific discovery work does not even fit Simon's classic definition of machine learning (Simon, 1987), according to which a program learns if it improves its performance on the basis of its own experience. A fixed program that is highly competent in some aspect of science (the Mathematica program, say) would not be a machine learning program, but may automate a very important class of tasks in science and engineering.

The conclusion to be drawn from this analysis is that scientific discovery *is not* a subfield of machine learning, much less a mere applications subfield, but is instead a separate subfield of AI which will continue to share common elements with machine learning.

6.2 Relation to computational science

A critical question is the relationship of discovery research to the computational research that scientists themselves increasingly carry out. One could make an analogy with the problem faced by academic computer architects: what research should be carried out in the face of a large, dynamic, and resource-rich computer industry? There are many scientists at work developing computer programs to support their work; what comparative advantage do computer scientists working in scientific discovery have?

Part of the answer is this: of all the disciplines, computer scientists are, or should be, the most qualified to perceive new scope for computing in any sphere, since they, generally speaking, are armed (in their heads) with the most computational tools and are best capable of developing the complex algorithms and programs that many tasks require. There are, of course, notable exceptions, such as Lederberg's innovative development of the first algorithm to generate molecular structures, which led to the DENDRAL project.

To carry out discovery research, computer scientists must be aware of scientific practice enough to identify novel scope for computing, and they should also be armed with sufficient applied mathematics and statistics, which clearly also find wide scope within science. Given the current state of affairs, computer scientists are likely to find the most novel scope within *inference,* since still today the common perception is that the role of computers within science is largely confined to *modelling,* not inference. This perception prevails despite the numerous successful counter-examples such as the inferential tasks automated by Mathematica or by the many programs of molecular biology, e.g., those that infer similarities between DNA strings or that infer phylogenetic trees. This perception is even fostered by some computer scientists: a recent article in *Communications of the ACM* (Stevenson, 1994) actually defines computational science in terms of modelling, thus completely ignoring inference, which is the larger of the two spheres. Most of what goes on in science is reasoning of some sort or another, not modelling nature.

An example may clarify the intended distinction between inference and modelling: my MECHEM program in chemistry (cited earlier) *reasons* in the course of finding a *model* of a reaction mechanism. The reasoning that leads to a model does not itself model any natural process. If a second program simulates the behaviour of a model found by MECHEM using the laws of chemical kinetics, then this second program is modelling nature. Similarly, any other program that searches a space of models is carrying out inference. The broader scope belongs to the class of programs that reason, not to the class of programs that model nature.

Besides an emphasis on novel scope, another realm of comparative advantage can be a concern with generality *across scientific domains,* since a computer scientist is not necessarily tied to a specific discipline for reasons of training and colleagues. There are strong generic elements in science, i.e., an inferential problem in one science very often has strong analogues in another science, so that there is a potential for the sharing of design knowledge and of formulating concepts

[3]Zytkow (1993) has argued that *autonomy* characterizes the difference between machine learning and discovery: discovery programs should be more autonomous than learning programs. The emphasis here on tasks of scientific inference is rather orthogonal to the issue of autonomy.

and propositions that pertain to classes of scientific tasks. Elsewhere I have argued this point in more detail by articulating the concept of a *generic scientific task* (Valdés-Pérez, 1995a).

In brief, the dimensions of likely comparative advantage for computer science research on scientific discovery lie in perceiving novel scope for computing and in formulating general knowledge about computing in science. The prerequisites for fulfilling this potential are a close attention to scientific practice, a willingness to enter deeply into scientific problems as the opportunity warrants, and a broad acquaintance with the intellectual tools of algorithms, artificial intelligence, applied mathematics, and statistics.

6.3 Relation to computer science

The relation of research on scientific discovery to basic computer science is an important issue. CS-oriented workers should not be viewed as doing something other than computer science, but instead expanding what doing computer science means. The subject matter of a field changes over time; one can browse the *Communications of the ACM* of the 1960s and contrast its contents with present-day CS research to get an idea, even without seeking analogies to other fields with longer and more evolved histories.

An intellectual, CS-oriented justification of discovery research is as follows. A very interesting question about computers is: What is their scope, what are their limits? This question can be studied theoretically and abstractly or practically.

To study it practically is to show that the scope of computers is wider than is ordinarily perceived, even by experts. This demonstration proceeds by building systems to automate practical affairs. The quality of the (basic) research correlates with the surprisingness of the achievements that are enabled by the program.

Practical affairs include science, art, programming, chess, and so on, i.e., affairs that are pursued with some seriousness by people (tic-tac-toe and noughts-and-crosses are not practical affairs). To obtain solid achievements will often require going deeply into other domains, which then accounts for the interdisciplinary character of such work. But, this character does not detract from the real, disciplinary CS content.

6.4 Relation to psycho-historical approaches

Much of the work on scientific discovery has followed a psychological or cognitive-science approach, by which I mean it has been constrained by behavioural or historical data or laws (for example, the sheer amount of sequential computation should not exceed human capabilities). From a CS viewpoint that emphasises algorithms and automation, the value of such psycho-historical work includes: (1) uncovering new scientific tasks to focus on; (2) contributing new concepts with which to reason about and describe scientific inference; and (3) elucidating the heuristics that scientists follow, especially on highly qualitative tasks that appear to be open-ended or ill-structured. I will illustrate these types of contributions with the following examples.

(1) The book *Scientific Discovery* by Langley et al. (1987) is a cognitive-science study of historical discovery tasks, taken mostly from the early history of chemistry. The authors pointed out that some tasks of model building in particle physics are analogous to the tasks that are modelled in their book. Kocabas (1991a) followed up on their suggestion and reported the first program that modelled the discovery of conservation laws in particle physics. In turn, I followed up on Kocabas by reporting a more systematic approach to the same task (Valdés-Pérez, 1994a). In parallel, I published an account for physicist audience (Valdés-Pérez and Erdman, 1994), to gain external validation of reality and significance, as has been recommended above.

(2) An early paper by Lea and Simon (1974) is often cited as an account of the dual-space nature of rule induction, which is a laboratory-psychology task that is somewhat analogous to the scientists' task of inducing general rules from experimental data. The notion of dual-space search, which has been elaborated by further cognitive-science studies (e.g., Klahr and

Dunbar, 1988), provides a basic concept with which to describe the diversity of reasoning tasks that confront the scientist. This concept has not shown up in the CS-oriented literature that focuses on inferential tasks involving a single problem space, but the concept has proved valuable for describing more open-ended tasks.

A second example of the use of psychological concepts within machine discovery can be drawn from my own work. A recent article (Valdés-Pérez, 1995d) reports a general pattern within several recent human/computer discoveries, and makes use of the Newell/Simon concept of a representation (Newell and Simon, 1972) to express this pattern: each of the discoveries involved a change in the representation of a scientific task.

(3) There have been many projects that deeply studied historical cases of discovery to elucidate the heuristics and problem spaces that scientists may have followed. For example, A. Gordon et al. (1994) studied the discovery of structural models for solutions in the history of chemistry. Tweney (1990) examined in detail Michael Faraday's notebooks, Holmes (1980) did likewise for Hans Krebs's discoveries in biochemistry, Kulkarni and Simon (1988) built a cognitive model of Krebs's urea-cycle discovery based on Holmes's work, and Thagard and Nowak (1990) studied the evolving conceptual structures during the development of plate tectonics in geology; many other similar projects could be cited. More recently, Dunbar (1995) has studied the reasoning strategies employed by working molecular biologists after spending many months at several of their research labs.

There seems to be much potential for such work to influence directly the design of computational actors in science practice. An obstacle to be overcome is that the discovery episodes addressed by these projects are often rather diffuse and special, so that extracting a promising task to automate, e.g., one that recurs in scientific practice, is not easy. Earlier, Newell explored in depth some of these difficulties, in the context of making use of G Polya's heuristics for mathematical discovery within an AI system (Newell, 1983).

7 Conclusion

This essay has examined briefly the history of research on the inferential problems that arise in science and mathematics, emphasising the viewpoint of computer science (i.e. automation) and also problems of creative scientific reasoning rather than the modelling of nature. I have also analysed the relation of CS-oriented discovery research to peripheral fields that either share an interest in the phenomena of science or share some of the methods of approach. The important issue of evaluating research quality has been tied to the fundamental problems of the field, such as extending the realm of computing within scientific practice at its most creative levels.

I close this essay with some brief recommendations. First, workers in this area should become familiar with a wide variety of intellectual tools, including the basic tools of the qualitative theory of AI and heuristic search, combinatorial algorithms and applied mathematics. This is the best equipment for perceiving new scope for computational scientific discovery. An acquantance with peripheral fields such as the psychology and philosophy of science (e.g., Salman, 1966) will have significant heuristic value as well. A willingness to penetrate deeply into the content of scientific problems is important.

Second, the primary questions that reviewers of new scientific discovery programs should pose is derived from the comparative advantage of computer scientists that was discussed above: (1) what new scope for computing in science is convincingly demonstrated? and (2) is there (some) generality *within science*? By no means should these be the sole questions; instead, the intent is to contrast them with a frequent alternative: what new AI method is proposed? Of course, the ideally best work would accomplish everything: demonstrate new scope, apply to a broad range of science, enable a new scientific discovery, incorporate an unprecedented method, etc.

Third, describing a new program is not the only way to advance the (computer) science of scientific discovery. There is a need for theory, either qualitative or more formal, that will enable general concepts and propositions concerning the role of computing in science.

Finally, the research program that this essay has developed will hasten the interesting and desirable disciplinary transformations that were foreseen a decade ago by Allen Newell (Bobrow and Hayes, 1985):

> We should, by the way, be prepared for some radical, and perhaps surprising, transformations of the disciplinary structure of science (technology included) as information processing pervades it. In particular, as we become more aware of the detailed information processes that go on in doing science, the sciences will find themselves increasingly taking a metaposition, in which doing science (observing, experimenting, theorizing, testing, archiving, ...) will involve understanding these information processes, and building systems that do the object-level science. Then the boundaries between the enterprise of science as a whole (the acquisition and organization of the knowledge of the world) and AI (the understanding of how knowledge is acquired and organized will become increasingly fuzzy).

References

Bobrow, DG and Hayes, PJ, 1985, "Artificial intelligence – where are we?" *Artificial Intelligence* **25** (3) 375–415.

Buchanan, B, Smith, D, White, W, Gritter, R, Feigenbaum, E, Lederberg, J and Djerassi, C, 1976, "Applications of artificial intelligence for chemical inference. 22. automatic rule formation in mass spectrometry by means of the meta-dendral program" *Journal of the American Chemical Society* **98** (20) 6168–6178.

Carbonell, JG, 1992, "Machine learning: A maturing field" *Machine Learning* **9** 5–7.

Chung, F, 1988, "The average distance and the independence number" *Journal of Graph Theory* **12** (2) 229–235.

Crevier, D, 1993, *AI: The Tumultuous History of the Search for Artificial Intelligence*, Basic Books.

Dunbar, K, 1995, "How scientists really reason: Scientific reasoning in real-world laboratories." In: *The Nature of Insight*, 365–395, MIT Press.

Gelernter, H, 1963, "Realization of a geometry-theorem proving machine." In: EA Feigenbaum and J Feldman, editors, *Computers and Thought*, McGraw-Hill.

Gordon, A, Edwards, P, Sleeman, D and Kodratoff, Y, 1994, "Scientific discovery in a space of structural models: An example from the history of solution chemistry." In: *Proceedings of the 16th Conference of the Cognitive Science Society* 381–386.

Hearn, AC, 1990, *Future Directions for Research in Symbolic Computation*, Siam, Philadelphia.

Hendrickson, J and Toczko, A, 1989, "SYNGEN program for synthesis design: basic computing techniques" *Journal of Chemical Information and Computer Sciences* **29** (3) 137–145.

Holmes, FL, 1980, "Hans Krebs and the discovery of the ornithine cycle." In: *Proc of 63rd Annual Meeting of the Federation of American Societies for Experimental Biology 39*, Symposium on Aspects of the History of Biochemistry.

Klahr, D and Dunbar, K, 1988, "Dual space search during scientific reasoning" *Cognitive Science* **12** 1–48.

Kocabas, S, 1991a, "Conflict resolution as discovery in particle physics" *Machine Learning* **6** (3), May, 277–309.

Kocabas, S, 1991b, "Computational models of scientific discovery", *Knowledge Engineering Review* **6** (4) 259–305.

Kulkarni, D and Simon, H, 1988, "The processes of scientific discovery: The strategy of experimentation" *Cognitive Science* **12** 139–175.

Langley, P, 1986, "Editorial: On machine learning" *Machine Learning* **1** 5–10.

Langley, P and Michalski, RS, 1986, "Editorial: Machine learning and discovery" *Machine Learning* **1** 363–366.

Langley, P, Simon, H, Bradshaw, G and Zytkow, J. 1987, *Scientific Discovery: Computational Explorations of the Creative Processes*, MIT Press.

Lea, G and Simon, HA, 1974, "Problem solving and rule induction: A unifield view." In: LW Gregg, editor, *Knowledge and Cognition*, Lawrence Erlbaum.

Lindsay, R, Buchanan, B, Feigenbaum, E and Lederberg, J, 1980, *Applications of Artificial Intelligence for Organic Chemistry: The Dendral Project*, McGraw-Hill.

Lindsay, R, Buchanan, B, Feigenbaum, E amd Lederberg, J, 1993, "DENDRAL: a case study of the first expert system for scientific hypothesis formation" *Artificial Intelligence* **61** (2) June, 209–261.

Mackenzie, D, 1995, "The automation of proof: a historical and sociological exploration" *IEEE Annals of the History of Computing* **17** (3) 7–29.

Newell, A, 1983, "The heuristic of George Polya and its relation to artificial intelligence." In: R Groner, M Groner and W Bischof, editors, *Methods of Heuristics*, Lawrence Erlbaum.

Newell, A and Simon HA, 1972, *Human Problem Solving*, Prentice-Hall.

Newell, A and Simon, HA, 1976, "Computer science as empirical inquiry: Symbols and search" *Communications of the ACM* **19** 111–126.

Oliver, JE, 1991, *The Incomplete Guide to the Art of Discovery*, Columbia University Press.

Salmon, W, 1966, *The Foundations of Scientific Inference*, University of Pittsburgh Press. In: J Shrager and P Langley, editors, *Computational Models of Scientific Discovery and Theory formation,* Morgan Kaufmann.

Simon, HA, 1987, "Why should machines learn" *Machine Learning* **1** 25–38.

Spirtes, P, Glymour, C and Scheines, R, 1993, *Causation, Prediction, and Search*, Springer-Verlag.

Stevenson, D, 1994, "Science, computational science, and computer science: At a crossroads" *Communications of the ACM* **37** (12) December, 85–96.

Thagard, P and Nowak, G, 1990, "The conceptual structure of the geological revolution." In: J Shrager and P Langley, editors, *Computational Models of Scientific Discovery and Theory Formation* 27–172, Morgan Kaufman.

Tweney, RD, 1990, "Five questions for computationalists." In: J Shrager and P Langley, editors, *Computational Models of Scientific Discovery and Theory Formation* 471–484, Morgan Kaufmann.

Valdés-Perez, RE, 1994a, "Algebraic reasoning about reactions: Discovery of conserved properties in particle physics" *Machine Learning* **17** (1) 47–68.

Valdes-Perez, RE, 1994b, "Human/computer interactive elucidation of reaction mechanisms: Application to catalyzed hydrogenolysis of ethane" *Catalysis Letters* **28** (1) September, 79–87.

Valdés-Perez, RE, 1995a, "Generic tasks of scientific discovery." In: *Working Notes of the Spring Symposium on Systematic Methods of Scientific Discovery*, AAAI Technical Reports.

Valdés-Perez, RE, 1995b, "Machine discovery in chemistry: New results" *Artificial Intelligence* **74** (1) March, 191–201.

Valdés-Perez, RE, 1995c, "Machine discovery praxis" *Foundations of Science* **1** (2) 219–224.

Valdés-Perez, RE, 1995d, "Some recent human/computer discoveries in science and what accounts for them" *AI Magazine* **16** (3) Fall, 37–44.

Valdés-Perez, RE and Erdmann, M, 1994, "Systematic induction and parsimony of phenomenological conservation laws" *Computer Physics Communications* **83** (2–3) 171–180.

Zytkow, J, 1993, *Machine Learning* **12** (1). Special double issue on machine discovery.