Abduction by Classification and Assembly[1]

John R. Josephson, B. Chandrasekaran,
Jack W. Smith Jr., and Michael C. Tanner

The Ohio State University

1. *Introduction*

We describe a general problem solving mechanism that is especially suited for performing a particular form of *abductive inference*, or best-explanation finding. A problem solver embodying this mechanism *synthesizes composite hypotheses*. It does so by by combining hypothesis parts as a means to the satisfaction of explanatory goals. In this way it is able to arrive at complex, integrated conclusions which are not pre-stored.

The intent is to present a computationally-feasible, task-specific problem solver for a particular information processing task which is nevertheless of very great generality. The task is that of synthesizing coherent composite explanatory hypotheses based upon a prestored, and possibly vast collection of hypothesis-generating "concepts". The authors' claim is nothing less than to have shown, in a new sense, and surpassing all other work in this area, how it is computationally possible for an agent to come to "know", based upon the evidence of the case.

The mechanism is described here functionally, and structurally; that is, the *why* and *what* of a computation are described, and algorithms are presented that show *how* the computations can be accomplished. Overall a classification machine is used for selecting plausible hypotheses, and a specialized means-ends machine is used for assembling a best explanation from the plausible hypotheses that are selected. The activity of the assembler is supervised by an overview critic, which uses the assembler to pointedly investigate the space of alternative composites. The result is an integrated knowledge-based problem solver functionally suited to its abstract information processing task.

Although the mechanism is an abstraction of the architecture of the Red-2 system, several other diagnostic AI systems realize it too, in varying degrees.

PSA 1986, Volume 1, pp. 458-470

https://doi.org/10.1086/psaprocbienmeetp.1986.1.193146 Published online by Cambridge University Press

1.1. Red

Red is a knowledge-based medical expert system for use in blood banks as a red-cell antibody identification consultant. (Josephson et al. 1984, Josephson et al. 1985, Smith et al. 1985). The system has now been through two working versions, and a third is under construction at the time of this writing. This paper presents an abstract description of the problem solving mechanism of Red-2, the second distinct version of the system. Thus Red-2 serves as a working proof of the realizability of the abstract design.

1.2. Abduction

Abduction or *inference to the best explanation* is a form of inference that follows a pattern something like this:

> D is a collection of data (facts, observations, givens),
> H explains D (would, if true, explain D),
> No other hypothesis explains D as well as H does.
> ---
> Therefore, H is correct.

The strength of an abductive conclusion will in general depend on several factors, including:

* how good H is by itself, independently of considering the alternatives,
* how decisively it surpasses the alternatives,
* how thorough the search was for alternative explanations, and
* pragmatic considerations, including
    * how strong the need is to come to a conclusion at all,
    * the costs of being wrong, and benefits of being right.

Abductions, as they are characterized here, go from data describing some thing or situation, to an explanatory hypothesis that best accounts for that data. Notice that calling an inference "abduction" carries with it the idea of its goal: a best explanation. In contrast classifying an inference as "deduction" carries instead the idea of a constraint that it satisfies: that it is valid or truth-preserving.

C. S. Peirce used the term "abduction" for a form of inference close to what we describe here (Peirce 1955 p. 150 ff). Gilbert Harman and others have written of "inference to the best explanation" for essentially the same pattern (Harman 1965. p. 88; Ennis 1968, p. 523; Josephson 1982); and Lycan calls it "the explanatory inference" (Lycan 1985). Sometimes a distinction has been made between an initial process of coming up with explanatorily useful hypothesis alternatives, and a subsequent process of critical acceptance where a decision is made as to which explanation is best. Often the term "abduction" has been reserved for the initial, hypothesis-originating stage (Peirce 1955). We use the term here for the whole process of inferring *from the data to the best explanation*.

Abductions appear to be ubiquitous in the un-selfconscious reasonings and perceivings of ordinary life, and in the more critically aware reasonings upon which scientific theories are built (Josephson 1982).

A common view is that diagnostic reasoning in general is abductive in character (Charniak and McDermott 1985; Pople 1973; Reggia 1985, p. 484). The idea is that the task of a diagnostic reasoner is to come up with a best explanation for the symptoms, i.e., those findings for the case which show abnormal values. The explanatory hypotheses appropriate for diagnosis are *malfunction hypotheses*, typically disease hypotheses for physicians, and broken-part hypotheses for mechanical systems.

The characteristic reasoning processes of fictional detectives has also been characterized as abduction (Sebeok and Umiker-Sebeok 1983). It has even been alleged that there are a minimum of 217 abductions to be found in the Sherlock Holmes canon (Truzzi 1983). It is arguable that abduction is an epistemologically fundamental form of non-deductive reasoning (Harman 1965, Josephson 1982, Chap. 3).

1.3. Need for Efficient Hypothesis Assembly

In some problem situations abduction can be accomplished by a relatively simple classification or hypothesize-and-match mechanism. If the number of potentially applicable hypotheses is small, if each one can be specified in advance, and if only one can be correct for a particular case, then each stored hypothesis can be matched against the data, with the quality of the matchings determining the winning hypothesis. But if the number of potentially applicable hypotheses is large, and if more than one can be correct at the same time, then the combinatorics of the situation will not permit us to have one pre-established pattern for each possible conclusion. One main alternative seems to be to actively construct the abductive conclusion as a combination of sub-hypotheses which are either abductive conclusions themselves, or are the products of some selection mechanism working from pre-established patterns.

Up to 2^n > different combined conclusions are made available by assembling from a space of n possible hypotheses. Thus a very large space of possible conclusions can result from a relatively small space of primitive categories. For example Red-2 has 54 most-detailed hypothesis parts, giving rise to more than 10^{16} potential conclusions. (Many of these, however, would not be internally consistent, and so could never be produced by the system. Eliminating inconsistent conclusions still leaves more than 10^{12} possible conclusions.) We will describe a mechanism that can efficiently pick out the best combination, even from so large a space.

2. The Mechanism

In this paper we will concentrate on presenting an abstract and functional description of the mechanism, with enough detail about the algorithms to make it clear how they work. An analysis of the computational complexity of the hypothesis assembly algorithm shows that modest assumptions about the domain suffice to make the algorithm tractable (Allemang et al. 1986). An evaluation of Red-2's performance is reported in (Smith et al. 1986), and shows that the system almost always produces clinically acceptable answers, even in complex cases.

https://doi.org/10.1086/psaprocbienmeetp.1986.1.193146 Published online by Cambridge University Press

2.1. Task and Subtasks

Suppose that we intend to build a computer program to capture expertise at a certain abductive task. That is, our program is to take, as input, data of a certain type; and produce, as output, best explanations for a well-defined subset of the input data. Suppose that we are given a large number of potentially applicable hypotheses "concepts" or "frames" to base the system on; and that more than one concept can correctly apply at the same time.

Notice that this is precisely the diagnostic situation a physician must face, where the pre-enumerated hypotheses correspond to known and named diseases, and where multiple diseases are common, especially among the very sick people seen at major hospitals, and among those with unobvious ailments. Notice too that this is (an aspect of) the situation faced by *any intelligent knowledge-using agent* facing a complex, changing world, armed primarily with "concepts" of what is possible, and having the goal of trying to "understand" some part of its experience by forming a "good" composite hypothesis.

Suppose further that interactions of various sorts between the pre-enumerated hypotheses can occur, making it unsatisfactory to just match each separately to the case and accept all those above a certain threshold of confidence. (We will discuss hypothesis interaction in more detail in a later section.)

One way to organize a system for this sort of task, and indeed the organization we are proposing, is to set up separate problem-solving structures for the distinct subtasks of:

* coming up with a relatively small number of "plausible" hypotheses from the much larger number of prestored patterns,
* building a "best" composite hypothesis using these plausible hypotheses as available parts,
* testing and improving the "goodness" of the composite.

We will see that this decomposition provides a good way of controlling the potentially explosive combinatorics of the problem.

2.2. The Major Modules and Their Functions

The overall function of the abduction machine can be described as that of producing a "best explanation" for a given set of data. A side effect is that information is made available about where there are alternative ways of explaining things, and this information is useful for critically assessing the goodness of that best explanation.

The major modules are:

* a classification machine for selecting plausible hypotheses,
* a specialized means-ends machine for assembling a subset of the plausible hypotheses into a "best" composite explanation, and

* an overview critic (described here algorithmically) which uses the means-ends assembler, first to produce a tentative initial composite, then repeatedly to explore the space of alternative composites, and then possibly again to rebuild a final "best explanation" after the pointed investigation of alternative explanations. This overview critic also does some problem solving to guarantee that the composite it finally produces is parsimonious, i.e., has no explanatorily superfluous parts.

2.3. The Classification Machine

Taking the MDX (Chandrasekaran et al. 1979, Chandrasekaran and Mittal 1983) system as it's point of departure, the classifier is implemented as a taxonomic hierarchy of hypothesis specialists. Each specialist in the hierarchy specializes in a single "concept". When invoked it will match that concept to the details of the case, either ruling it out of further consideration, or else producing a hypothesis that has an associated symbolic likelihood, and offers to explain certain of the findings of the case.

The hierarchy organizes the specialists from most general at the top, to most specific at the tip nodes. The hypothesis selection activity proceeds in a top-down, more-general-to-more-refined manner, taking advantage of the search pruning effect that comes from ruling out whole subtrees of hypotheses by ruling out at high levels of generality. This top-down, prune-or-pursue control regime, associated with MDX-like diagnostic systems has been called "establish-refine". It can in principle proceed in parallel, matching of two sub-concepts being typically independent of each other, although we have only made serial implementations of establish-refine up to this time. By efficiently pruning the search for plausible hypotheses, establish-refine is a significant contributor to taming the combinatorics of the problem space. It makes it efficient and practical to search a very large space of stored hypotheses for just those that plausibly apply to the case.

2.4. Plausible Hypotheses

Each hypotheses that is considered and cannot be ruled out is matched against the data of the case to produce a description of which parts of the data it can explain (or contribute to explaining), and how plausible it is under the circumstances. Thus each plausible hypothesis delivered by the classifier comes with:

* a description, particularized to the case, of which findings it offers to explain,
* a symbolic plausibility value representing a prima facie estimate of likelihood for the hypothesis.

Each plausible hypothesis has its own consistent little story to tell, and to contribute to the larger story representing the abductive conclusion.

2.5. Hypothesis Interactions

Hypothesis interactions are considered to be of two general types, each with its own kind of significance for the problem-solving:

https://doi.org/10.1086/psaprocbienmeetp.1986.1.193146 Published online by Cambridge University Press

* *explanatory interactions*, i.e., due to overlapping in what the hypotheses can account for, and
* *substantive interactions* of mutual support and incompatibility, e.g., resulting from causal or logical relations.

For example two disease hypotheses might offer to explain the same findings without being especially compatible or incompatible causally, logically, or definitionally. On the other hand hypotheses might be mutually exclusive (e.g., because they represent distinct sub-types of the same disease), or mutually supportive (e.g., because they are causally associated). The Internist system did not make a clear distinction between hypotheses which are competitors because they are both capable of explaining the same findings in the case (thus not both needed), and those that are competitors because they are mutually exclusive. (Pople 1977. p. 1030). Internist was only concerned with the former type. In general the elements of a diagnostic differential need to be exhaustive of the possibilities, so that at least one must be correct (which one can be discovered by exclusion), but they need not be mutually exclusive.

The following types of hypothesis interaction can be accommodated and treated appropriately by the mechanism we are describing. Appropriate handling for all of them has all been implemented and tested:

* A and B are mutually compatible, and represent explanatory alternatives where their explanatory capabilities overlap.
* Hypothesis A is a subhypothesis of B (i.e., a more detailed refinement).
* A and B are mutually incompatible.
* A and B cooperate additively where they overlap in what they can account for.

Another form of hypothesis interaction that it is easy for the mechanism to handle is where one hypothesis, if it is used as part of a best explanation, suggests that a particular other one be used also. This has only been partly implemented, and has not been tested, but there seems to be no special problem in doing so.

Yet another form of interaction is where one hypothesis, if it is accepted, raises explanatory questions of its own that are resolved by appeal to another hypothesis. For example a medical diagnosis machine might hypothesize the presence of a certain pathophysiological state to explain certain symptoms, and then hypothesize some more remote cause to account for the pathophysiological state. The tummy ache is explained by the presence of the ulcers, and the ulcers is in turn explained by the anxiety neurosis. Unfortunately our present domain does not call for exploring this kind of hypothesis interaction, yet it appears to be a relatively straightforward matter to assemble composite hypotheses along these lines. This will be discussed in a later section. We look forward to exploring this dimension of hypothesis assembly in some future domain.

2.6. The Hypothesis Assembler

A mechanism for hypotheses assembly is used which is reminiscent of the means-ends regime of GPS (Newell and Simon 1963). It detects differences between the goal state (everything explained) and the present state (the working hypothesis does not explain everything), and extracts a salient difference (a most significant unexplained finding). It uses this unexplained finding to select a hypothesis part to integrate into the growing working composite.

We begin by describing a basic hypothesis assembler, capable only of treating one type of hypothesis interaction. Then we will describe how it can be enhanced to appropriately treat the other types of interaction.

2.6.1. The Basic Assembler

The basic assembler treats only hypotheses that are mutually compatible and that represent explanatory alternatives where their explanatory capabilities overlap. A set of findings is given, the object is to assemble an explanation for them, and to do so in a manner that respects the plausibilities of the candidate parts. Note that the findings to be explained are in general a proper subset of all of the findings of the case. We might try to explain the patient's symptom, but we won't try to explain his age.

The assembler works by using the plausibilities to guide a means-ends search whose goal is a complete explanation for the set of findings.

Procedure:

> * Loop until there is nothing left to explain, or nothing left that can be explained.
> *
>
> > ° Focus attention on an unexplained finding (initially the whole set is unexplained). If domain knowledge is available to point out the most significant unexplained finding, then well and good; but if not, then the choice can be made at random.
> > ° Pick the most plausible hypothesis that explains that finding. If no plausible explanation for it can be found, then note the finding as unexplainable and loop again, else continue. If more than one explanation for the finding is maximally plausible, then if knowledge is available to guide the choice, use it, and if none is available choose at random.
> > ° Add the chosen finding into the unstructured set of hypotheses that constitutes the growing composite hypothesis.
> > ° Compute what the composite can now explain.
> > ° Compare what can be explained to what needs to be explained overall and determine the unexplained remainder.
>
> * End loop.

The basic assembler produces a composite hypothesis which is as complete as possible. Since it uses the most plausible explanatory hypothesis at each choice point, the composite hypothesis is sense

https://doi.org/10.1086/psaprocbienmeetp.1986.1.193146 Published online by Cambridge University Press

maximally plausible as well, or nearly so. (The conditions under which this process produces an optimally plausible composite have been investigated to some extent, and will possibly form the subject of a future paper.)

It is easy and computationally inexpensive to rid the composite of explanatorily superfluous parts: check through the parts in order of least plausible to most plausible; for each part compute the explanatory capabilities with the part removed; and check to see if there is any loss.

Note that this interpretation of Ockham's Razor has clear epistemic virtues. Logically the composite hypothesis is a conjunction of little hypotheses; so, if we remove one of the conjuncts the resulting hypothesis is distinctly more likely to be true, since it makes fewer commitments. Superfluous hypothesis parts make factual commitments, expose themselves to falsity, with no compensating gain in explanatory power. Thus the sense of parsimony we propose here is such *that the more parsimonious hypothesis is more likely to be true.*

By this assembly process we arrive at a composite hypothesis which is *as complete as possible, maximally plausible* (or nearly), and *parsimonious.*

Since the assembly process added monotonically to a growing hypothesis, with incrementally growing explanatory power, and with no backtracking, the process is computationally very inexpensive. In general the greatest computational expense will be in checking through the available hypotheses to determine which one is the most plausible way to explain the finding of attention. But the classifier will collaborate to reduce the alternatives to a relatively small number, and one pass through the set will suffice. The whole process of assembly is computationally very efficient.

2.6.2. *Extensions and Elaborations*

Extensions can be made to the basic assembler to handle the other types of hypothesis interaction we have mentioned.

If hypotheses in the space come with subtype relationships, as they normally would with a hierarchical classifier, the assembler can preferentially pursue the goal of explanatory completeness and secondarily pursue the goal of refining the constituent hypotheses down to the level of most detail. This is unimplemented in the present version of Red, which builds up its composite hypothesis at the level of the most refined hypotheses. But there seems to be no special difficulty with this strategy - there is more than one way to go about the refinement process - and we expect to implement it in the next version.

A more difficult problem is in devising a strategy for when some of the hypotheses in the space are mutually incompatible. (We assume the ability to determine whether any given pair of hypotheses is compatible or not.) One thing we can do is to maintain the consistency of the growing hypothesis as we go along. If a finding is encountered whose only available maximally plausible explainers are incompatible with something in the growing hypothesis, then we add one of these incompatible hypothesis to the growing hypothesis, removing from it any parts inconsistent with the new one. (If we remove parts from the growing hypothesis we introduce the danger of an infinite loop, but fortunately this can be dealt with fairly readily.) The basic idea is that the finding must be explained, even if that forces a serious revision of the growing hypothesis. This is all implemented in Red-2.

This scheme for dealing with incompatible hypotheses seems to be little more than a clever trick for getting the job done. It is weak on its use of knowledge, and is not very adaptive or opportunistic. It endangers the computational feasibility of the hypothesis assembly process by threatening to force a search through the potentially large space of all possible consistent combinations (though it does search in such a way that it favors hypothesis parts of higher plausibility and greater explanatory indispensability, which makes it pretty clever.)

If hypotheses can cooperate additively where they overlap in what they can explain, all we need to do is to suitably incorporate this knowledge into the methods for computing what a composite hypothesis can explain. This too is already implemented in the present version of Red.

In order to handle the kind of hypothesis interaction where one hypothesis suggests the use of another, as for example if there is available knowledge of a statistical association, we can give extra plausibility credit to the suggested hypothesis if the hypothesis making the suggestion is already part of the growing composite. The availability of a way to grow the hypothesis preferentially along lines of statistical association provides a rudimentary ability for it to grow along causal lines as well. This has only been partly implemented, and has not been tested, but there seems to be no special problem in doing so.

A more interesting ability to grow along causal lines results if we permit one hypothesis, if it is accepted into the growing hypothesis, to raise explanatory needs of its own. For example, a newly added hypothesis can be posted as a kind of higher-level finding which needs to be explained in its turn by the growing assembly. Thus at the same time that the newly added hypothesis succeeds in explaining some of the findings, it introduces a "loose end". This provides a way in which the growing hypothesis can move from hypotheses close to the findings of the case, and towards more and more remote causes of those findings. This has not been implemented; Red's domain problem does not seem to call for this ability.

2.7. The Overview Critic

Procedure:

* The assembler is invoked to produce a tentative best explanation.
* Explanatorily superfluous parts are removed.

https://doi.org/10.1086/psaprocbienmeetp.1986.1.193146 Published online by Cambridge University Press

* The assembler is invoked repeatedly as necessary to assess which of the hypotheses in the composite are indispensable. A hypothesis is judged indispensable if removing it from a composite which is a complete explanation, leaves behind a composite which cannot then be assembled to completion without reintroducing the removed one. It follows that a hypothesis is indispensable if and only if something that it explains *has no other plausible explanation*.
* The non-indispensable parts of the composite are removed, and the assembler is invoked again to rebuild from the core of indispensables back to a complete explanation.
* Explanatorily superfluous parts are removed.

At the end of this process the composite hypothesis explains as much as possible, is maximally plausible (or nearly), is parsimonious, and has been built up by going from a core of hypotheses which are most certain.

At this stage the best explanation has been inferred, or at least A best explanation has been inferred, there being no a *priori* guarantee that a best explanation is unique.

3. Extensions and Elaborations

The degree of intimacy between the classifier, the overview critic, and the assembler, is an unresolved research issue which we are actively exploring. In Red-2 the classifier runs first, producing a set of plausible hypotheses, and then is followed by the critic, which uses the assembler to produce the best explanation. In the near future we anticipate a version where the classifier and the overview critic run concurrently, with the critic using its perspective on the progress of the problem solving, to help guide the search for plausible hypotheses. More distantly we envision a version where lots of little hypothesis assemblers, and also maybe bits of overview criticism, are distributed over a conceptual structure that makes local abductions, producing little assembled best explanations. By solving subproblems the little abductors serve the needs of larger abductors, and make it possible to assemble hypotheses from parts which are themselves assemblies.

4. Summary

We have described how best explanations can be inferred by a mechanism which tames the combinatorics of very large spaces of explanatory hypothesis. Structured conclusions can be arrived at whose parts are connected by relationships of type-subtype, statistical association, and explainer-explained. An instance of this machine exists which exercises some of the capabilities we attribute to the abstract machine, and that gets correct answers in complicated situations. (Smith et al. 1986).

A computational description has been given to the functional architecture of a possible mind, or rather, of a certain dimension or slice of a possible mind. The kind of synthesis of explanatory hypotheses we describe here is a *generic task of higher intelligence*. It must be accomplished somehow by any intelligent, knowledge-using agent that comes to "know" by calling upon "concepts", attaching them to situations or objects, and using the resulting little hypotheses as materials to form composite "best explanations". The task is general, but specific. There are a limited number of functional architectures that could accomplish it, especially when account must be taken of the constraints imposed by limited knowledge, limited time, and limited computational resources. There are even fewer architectures that are *especially suited* to the task, and one of them has been described.

Note

[1]This work has been supported in various stages by NSF Grant MCS-8305032, and NIH Grant R01 LM 04298 from the National Library of Medicine. Dr. Jack W. Smith, Jr. is supported by NLM Career Development Award K04 LM00083.

Praise is due Tom Bylander for showing that the task of producing a consistent composite is NP complete. Thanks are also due to Tom Bylander for his helpful comments on a previous draft, to Bill Punch and Dean Allemang for their discussions and encouragement of the approach to abduction, and to Jon Sticklen for arguing until things were better justified. Also thanks to the members of the recent graduate seminar at Ohio State on diagnostic reasoning for their helpful comments, and to two anonymous reviewers of an earlier draft, who, by their failure to understand what was being presented, pointed the way to an improved explication.

References

Allemang D.; Tanner, M.; Bylander T.; and Josephson, J. (1986). On the Computational Complexity of Hypothesis Assembly. Technical Report, Laboratory for Artificial Intelligence Research, Department of Computer and Information Science, The Ohio State University, The Ohio State University, LAIR, CIS Dept., CAE Bldg., Columbus, Ohio, 43210-1277, 1986. Submitted for publication.

Chandrasekaran, B.; Gomez, F.; Mittal, S.; and Smith, J.W. (1979). "An Approach to Medical Diagnosis Based on Conceptual Structures." In Proceedings of the Sixth International Joint Conference on Artificial Intelligence - Tokyo 1979. Pages 134-142.

------------------ and Mittal, S. (1983). "Conceptual Representation of Medical Knowledge for Diagnosis by Computer: MDX and Related Systems." In Advances in Computers. Edited by M. Yovits. New York: Academic Press. Pages 217-293.

Charniak, Eugene and McDermott, Drew. (1985). Introduction to Artificial Intelligence. Reading, MA: Addison Wesley.

Ennis, R. (1968). "Enumerative Induction and Best Explanation." The Journal of Philosophy LXV: 523-529.

Harman, Gilbert. (1965). "The Inference to the Best Explanation." The Philosophical Review LXXIV: 88-95.

Josephson, John R. (1982). Explanation and Induction. Unpublished Ph.D. Dissertation, The Ohio State University. Xerox University Microfilms Publication Number DE082-22107.

------------------; Chandrasekaran, B.; and Smith, J.W. (1984). "Assembling the Best Explanation." In Proceedings of the IEEE Workshop on Principles of Knowledge-Based Systems. IEEE Computer Society, Denver, Colorado, December 3-4, 1984. Pages 185-190. (A revised version by the same title is now available as a technical report.)

------------------; Tanner, Michael C.; Smith, Jack W.; Svirbely, John R.; and Straum, Pat. (1985). "Red: Integrating Generic Tasks to Identify Red-Cell Antibodies." In Proceedings of the Expert Systems in Government Symposium. Edited by Kamal Karna. Washington, DC: IEEE Computer Society Press. Pages 524-531.

Lycan, William G. (1985). "Epistemic Value." Synthese 64: 137-164.

Newell, A. and Simon, H.A. (1963). "GPS, A Program That Simulates Human Thought." In Computers and Thought. New York: McGraw-Hill. Pages 279-293.

Peirce, C.S. (1955). "Abduction and Induction." In Philosophical Writings of Peirce. New York: Dover. Pages 150-156.

Pople, H. (1973). "On the Mechanization of Abductive Logic." In Proceedings of the Third International Joint Conference on Artificial Intelligence - Stanford 1973. Pages 147-152.

--------. (1977). "The Formation of Composite Hypotheses in Diagnostic Problem Solving: An Exercise in Synthetic Reasoning." In Proceedings of the Fifth International Joint Conference on Artificial Intelligence - Cambridge, MA 1977. Pages 1030-1037.

Reggia, J. (1985). "Abductive Inference." In Proceedings of the Expert Systems in Government Symposium. Edited by Kamal Karna. Washington, DC: IEEE Computer Society Press. Pages 484-489.

Sebeok, Thomas A. and Umiker-Sebeok, Jean. (1983). You Know My Method. A Juxtaposition of Charles S. Peirce and Sherlock Holmes. Bloomington, Indiana: Indiana University Press.

Smith, Jack W.; Svirbely, John R.; Evans, Charles A.; Strohm, Pat; Josephson, John R.; and Tanner, Mike. (1985). "Red: A Red-Cell Antibody Identification Expert Module." The Journal of Medical Systems 9(3): 121-138.

--------------; Josephson J.; Tanner, M.; Svirbely, J.; Strohm, P. (1986). "Problem Solving in Red Cell Antibody Identification: Red's Performance on 20 Cases." Technical Report, Laboratory for Artificial Intelligence Research, Department of Computer and Information Science, The Ohio State University, The Ohio State University, LAIR, CIS Dept., CAE Bldg., Columbus, Ohio, 43210-1277, 1986. Submitted for publication.

Truzzi, Marcello. (1983). "Sherlock Holmes: Applied Social Psychologist." In The Sign of Three: Dupin, Holmes, Peirce. Edited by T.A. Sebeok and E. Umberto. Bloomington, Indiana: Indiana University Press. Chapter Two.