A Hill-Climbing Approach to Machine Discovery

DONALD ROSE (DROSE@ICS.UCI.EDU)
PAT LANGLEY (LANGLEY@ICS.UCI.EDU)
Department of Information & Computer Science, University of California, Irvine CA 92717 USA

Abstract

To account for new observations, scientists must often change an initial theory (set of beliefs) that has become inconsistent due to this data. Automated discovery systems should also be able to revise their beliefs in such scenarios, in order to regain a consistent database. To this end we have constructed REVOLVER, a program that performs both discovery and belief revision. When the system makes erroneous inferences, it employs hill climbing to search for a new consistent set of beliefs. In this paper, we first discuss the goals and tasks addressed by REVOLVER. Next we describe the program's operation, using an example from chemistry to illustrate the system's representation, its basic rules and inference process, and its method for belief revision. We then evaluate the system, showing its generality (through its replication of discoveries in natural domains) and its robustness (through its ability to run efficiently in artificial domains). Finally, we discuss related work on belief revision and ideas for future work.

1. Introduction

What is the composition of matter? This problem has arisen repeatedly during the history of science. Scientists often responded by searching for models of substances in terms of more fundamental entities. Historically, they first considered models of observable substances in terms of other observable substances; we call these *observational* models. However, as new physical laws and facts were discovered, the focus shifted to models in which these substances were composed of more 'basic' entities at a lower unseen level; we call these *theoretical* models.

In this paper we describe REVOLVER, a system that formulates both types of models and replicates historical discoveries from domains such as chemistry and physics. The program inputs reactions relating groups of substances and uses heuristics to transform these premises into models. If premises lead to inconsistent beliefs, the system uses hill climbing to search the space of revised premises in order to resolve the errors. In the following pages we describe REVOLVER's basic inference and revision processes, then discuss its replication of scientific discoveries. After presenting experimental studies of the program's behavior as both the system and its domain are varied, we discuss related work in belief revision and consider options for future research.

2. The REVOLVER System

The system deals with two kinds of beliefs: *reactions* and *models*. Reactions, which are the system inputs (premises), represent relations between types of objects (e.g., the inputs and outputs of a chemical reaction). Examples from 18th century chemistry would be the observation that potassium and oxygen react to form caustic-potash and water ((1) K O $\rightarrow$ P W) and the observation that potassium and water react to form caustic-potash, hydrogen and water ((2) K W $\rightarrow$ P H W). Given premises such as these, REVOLVER uses heuristics to try to infer models (i.e., relations between substances and components). Since W appears on both sides of (2), it is reduced from the reaction, leaving (3) K $\rightarrow$ P H. When a substance is alone on one side of a reaction, the system infers that the opposite side contains its components; hence REVOLVER infers from (3) the model (4) K = P H. The program then substitutes K's components from (4) into (1) to get (5) P H O $\rightarrow$ P W. Next, P is reduced from both sides of (5), leaving (6) H O $\rightarrow$ W. Finally, the model (7) W = H O is inferred.

In the above example, the premises were consistent; the system reached a quiescent state without inferring any erroneous beliefs. However, the premises given to REVOLVER sometimes lead to reactions having either no inputs or no outputs. In such cases the program invokes its belief revision process; REVOLVER finds the premises that led to the bad inference, then considers revisions to them that would bring the database closer to consistency. After revising a premise, the system continues making new inferences and, if it detects new inconsistencies, revises premises again. This cycle continues until no more inferences can be made and no inconsistencies exist.

Continuing our earlier example, suppose the system receives new premise (8) P $\rightarrow$ K O. Substituting for K leads to (9) P $\rightarrow$ P H O and reducing P yields inconsistent reaction (10) nil $\rightarrow$ H O. Invoking belief revision, the program identifies premises (2) and (8) as the sources of the error, and must decide how to revise them in order to resolve the inconsistency. Since REVOLVER eliminates one substance from an inconsistency during each revision step, the system proposes six candidate revisions in our example:

R1: add H to (2)'s inputs; R2: add O to (2)'s inputs; R3: delete H from (2)'s outputs;
R4: add H to (8)'s inputs; R5: add O to (8)'s inputs; R6: delete O from (8)'s outputs.

Implementing any of these revisions eliminates one substance from the inconsistency (after the program makes the revision, deletes beliefs supported by that premise, and restarts inferencing).

However, REVOLVER only carries out one of its candidate revisions, using an *evaluation function* to make the selection. The program scores the premises considered for revision along several criteria, multiplies each score by a weight (indicating the priority given to each criterion), sums the weighted scores, then revises the premise having the lowest total score. Since the system does not retain the alternate premises generated during each revision step, REVOLVER is a hill-climbing system, relying on its evaluation function to guide search towards consistent beliefs without storing past states (sets of beliefs).

One heuristic for selecting revisions embodies a preference for revising premises that support the least number of beliefs. Referred to as *minimum mutilation* in (Quine & Ullian, 1978) and *conservatism* in (Harman, 1987), this heuristic enables one to change belief in a way that affects the rest of one's belief system in the least drastic manner. If only this heuristic is used, the system would prefer changing only (8), since that premise supports no models while (2) supports one. Hence, REVOLVER would narrow its choice to R4, R5 or R6 (either add H to (8)'s inputs, add O to (8)'s inputs, or delete O from (8)'s outputs).

However, REVOLVER employs additional measures to ensure that it selects the best possible revisions. Another important criterion is the *complexity* of the premises (i.e., the number of substances in a reaction). By appropriately setting this complexity weight, REVOLVER can prefer either the deletion of substances in premises (i.e., prefer parsimony), or the addition of substances (i.e., assume that substances were observed correctly, but that some substances may have been overlooked). Let us assume that the belief support measure has priority over the premise complexity measure, and that the latter is set to prefer addition of substances. In this case, (8) is still the preferred premise to revise (due to minimal belief support), leaving R4, R5 and R6. The first two are preferred since they involve adding substances. Suppose R4 (adding H to (8)'s inputs) is chosen over R5; this leads to revised premise (11) P H $\rightarrow$ K O. Substituting P and H for K in (11) and reducing twice yields (14) nil $\rightarrow$ O. REVOLVER proposes revisions again:

R7: add O to (2)'s inputs; R8: add O to (11)'s inputs; R9: delete O from (11)'s outputs.

Since belief support has priority over premise complexity, REVOLVER prefers revising (11); R8 is chosen over R9 since it involves addition of a substance. Implementing R8 results in (15) P H O $\rightarrow$ K O. Further inferencing leads to quiescence; the database is now consistent.

REVOLVER's evaluation function also incorporates a metric that prefers revising premises that have been changed less often that others. Using only the belief support and complexity preferences, there are cases where premises may be revised such that earlier changes are undone and previous inconsistent premises reappear in the system's database. In some cases, this cycle can continue indefinitely, since REVOLVER does not save previous beliefs to check for such cycling. A *minimum revision* criterion helps alleviate such cycling without the need to store previous memory states. This measure ensures that each premise leading to an inconsistency will eventually be revised in turn if the hill-climbing process begins cycling among the same sets of beliefs.

Finally, REVOLVER's evaluation function can be set to prefer revising only the number of substances present in a premise (i.e., assume premises are already correct in the type of substances present). When this *same-type assumption* is highly preferred, the program will usually add a substance to a premise only if it is already present, and delete that substance only if it is still present afterwards. For complex problems this assumption can be beneficial, since it constrains the space of beliefs the system must search in order to find consistency.

3. Evaluating the System

Having described the system and its behavior on a specific example, we now evaluate REVOLVER in more general terms. First, we show that the program can replicate discoveries involving both observational and theoretical beliefs, and do so in multiple domains (e.g., chemistry and particle physics). Second, we examine the behavior of the system when both its evaluation function and its domain are varied.

3.1 REVOLVER on Historical Domains

By the 1950s, physicists had discovered a large set of subatomic particles called hadrons, leading to a search for more fundamental entities that could serve as their building blocks. These scientists constructed observational models of hadrons in terms of other hadrons, just as chemists had formed models of observed substances (compounds) in terms of other observed substances (elements). We have already seen a simple example of how REVOLVER constructs such observational models in chemistry; an interesting characteristic of the system is that the same methods (with the addition of a few domain-specific heuristics) can be applied to similar domains, such as particle physics. For this new task, we gave REVOLVER relations between hadrons (e.g., the proton p and the neutron n) and associated properties called *quantum numbers* (e.g., electric charge, strangeness, and isospin). The amount of each property exhibited by a particle is denoted by the number of occurrences of $+$, σ, and I in the right side of that particle's premise reaction. The $+$ represents an electric charge of 1, the σ represents a strangeness of 1, and the I represents an isospin of 1/2. Similarly, $\mp$, $\bar{\sigma}$ and $\bar{I}$ represent −1, −1 and −1/2, respectively.[1] The complete set of initial hadron reactions, corresponding to knowledge used historically,[2] are:

$$\Xi^- \rightarrow \bar{\sigma}\,\bar{\sigma} \mp \bar{I},\ \Xi^0 \rightarrow \bar{\sigma}\,\bar{\sigma}\, I,\ \Sigma^- \rightarrow \bar{\sigma} \mp \bar{I}\,\bar{I},\ \Sigma^0 = \Lambda \rightarrow \bar{\sigma},\ \Sigma^+ \rightarrow \bar{\sigma} + I\, I,$$
$$n^- \rightarrow \mp \bar{I}\,\bar{I}\,\bar{I},\ n \rightarrow \bar{I},\ p \rightarrow +\, I,\ n^{++} \rightarrow +\, +\, I\, I\, I,\ \pi^+ \rightarrow +\, I\, I,\ K^0 \rightarrow \sigma\, \bar{I},\ K^+ \rightarrow \sigma + I$$

REVOLVER now tries to solve for σ, $+$, I, and their antiparticles in terms of hadrons. After inferring the models $\bar{\sigma} = \Lambda$, $\bar{I} = n$, $\sigma = \bar{\Lambda}$ and $I = \bar{n}$, REVOLVER substitutes these 'components' for $\bar{\sigma}$, σ, $\bar{I}$ and I in each reaction. The system then transforms $p \rightarrow + \, \bar{n}$ into $p\, n \rightarrow +$ and $\bar{p} \rightarrow \mp\, n$ into $\bar{p}\, \bar{n} \rightarrow \mp$,[3] which leads in turn to $+ = p\, n$ and $\mp = \bar{p}\, \bar{n}$. After annihilating particles and

[1] For instance, $+\,+$ indicates an electric charge of 2, $\bar{I}\,\bar{I}\,\bar{I}$ indicates an isospin of −3/2, and so on.

[2] Reactions for their antiparticles, which are not shown but were also part of the run, have opposite valued quantum numbers. The final models for antiparticles also have opposite values.

[3] A new rule the program uses in this domain states that a particle on one side of a reaction is equivalent to deleting it and placing its antiparticle on the other side. For example, $A\, \bar{b} \rightarrow c$ can be transformed into $A \rightarrow b\, c$ in order to isolate A, an observable particle for which a model is sought.

antiparticles,[4] the final set of models are:

$$\Xi^- = \Lambda\ \Lambda\ \bar{p},\ \Xi^0 = \Lambda\ \Lambda\ \bar{n},\ \Sigma^- = \Lambda\ \bar{p}\ n,\ \Sigma^0 = \Lambda = \Lambda,\ \Sigma^+ = \Lambda\ p\ \bar{n},$$
$$n^- = \bar{p}\ n\ n,\ n = n,\ p = p,\ n^{++} = p\ p\ \bar{n},\ \pi^+ = p\ \bar{n},\ K^0 = \bar{\Lambda}\ n,\ K^+ = \bar{\Lambda}\ p$$

Each hadron is now composed of other hadrons. These models, proposed historically by S. Sakata, were collectively called the *composite model of hadrons* (Nambu, 1985). Models for p, n and Λ (and their antiparticles) are not present because they have no components; they are the building blocks of the other particles, a conclusion that matches Sakata's theory. [5]

While we have now seen REVOLVER construct observational models both with and without the need for belief revision, the system can also formulate theoretical models in which components are unobservable substances. After inputting assumptions about which 'smaller' (unseen) components make up each substance, REVOLVER repeatedly revises these assumed models until they are consistent. Thus, although belief revision is the exception when constructing observational models, it is the norm when constructing theoretical models.

To illustrate this process, let us examine REVOLVER's formulation of theoretical models in both chemistry and physics.[6] Suppose we give the program the simplest initial atomic models for hydrogen and oxygen (H = h and O = o), along with the simplest atomic model for water that matches the observational model inferred earlier (i.e., W = h o, where the hydrogen and oxygen atoms here correspond to the observed substance types in the earlier model W = H O). Once the water reaction H H O → W W is asserted,[7] substitution leads to h h o → h o h o, which in turn leads to nil → o. Invoking belief revision, REVOLVER generates O = o o as the revised oxygen model. Substitution into the water reaction yields h h o o → h o h o, then nil → nil after reduction. Quiescence has been reached; hence the models O = o o, H = h and W = h o are consistent and successfully summarize the data (i.e., the water reaction). Although these models differ from the modern view, they were in fact proposed by the 19th century chemist Dalton.

Now suppose the system encounters (the simplest) initial models for nitrogen (N = n) and ammonia (A = n h), along with the ammonia reaction H H H N → A A. Further revisions now become necessary, and REVOLVER's search for consistency continues. The main snapshots are:

W = h o, H = h, O = o o, nil → nil, N = n, A = n h, h h h n → n h n h
W = h o, H = h, O = o o, nil → nil, N = n n, A = n h, h h h n n → n h n h
W = h o, H = h, O = o o, nil → nil, N = n n, A = n h h, h h h n n → n h h n h h
W = h o, H = h h, O = o o, h h → nil, N = n n, A = n h h, h h h h h h n n → n h h n h h
W = h o, H = h h, O = o o, h h → nil, N = n n, A = n h h h, h h h h h h n n → n h h h n h h h
W = h h o, H = h h, O = o o, nil → nil, N = n n, A = n h h h, nil → nil

At this point the system has found consistent models, since the water reaction and the ammonia reaction have both been reduced to nil → nil. These models match those proposed by the chemist Avogadro, and are still held today.

However, this process of finding theoretical models is not specific to chemistry alone; REVOLVER can apply this process in particle physics as well. The system can take initial assumptions about

[4] A second new rule states that particles cancel out their antiparticles when on the same reaction side.

[5] In general, the system can rediscover many kinds of observational models, and replicates the examples of observational model construction in chemistry discussed in (Rose & Langley, 1986) and (Zytkow & Simon, 1986).

[6] The program uses the same-type assumption in both domains; initial models are assumed already correct in the type of substances present, but the number of occurrences in beliefs is allowed to change.

[7] Multiple occurrences of symbols indicate that scientists (e.g., 19th century chemists) eventually begin considering amounts (not just types) of substances in their reasoning. In the chemical domain, capital letters represent molecules (the number of occurrences based on Gay-Lussac's law of combining volumes); lower case letters denote atoms.

theoretical hadron models, stated in terms of smaller (hypothetical) particles called quarks, and revise these premises into a consistent set. In this new example, suppose the initial hadron models contain the correct types of quarks (called u, d, $\bar{u}$ and $\bar{d}$)[8] and REVOLVER inputs only hadrons that do not exhibit strangeness (n^-, n, p, n^{++}, π^- and π^+). In this case, the main snapshots of the inference and revision process are as follows:

$$n^- = d,\ n = u\ d,\ p = u\ d,\ n^{++} = u,\ \pi^+ = u\ \bar{d},\ \pi^- = \bar{u}\ d,\ n \rightarrow p\ \pi^-,\ n^{++} \rightarrow p\ \pi^+,\ n^- \rightarrow n\ \pi^-$$
$$n^- = d,\ n = u\ d,\ p = u\ d,\ n^{++} = u,\ \pi^+ = u\ \bar{d},\ \pi^- = \bar{u}\ d,\ n \rightarrow p\ \pi^-,\ u \rightarrow u\ d\ u\ \bar{d},\ n^- \rightarrow n\ \pi^-$$
$$n^- = d,\ n = u\ d,\ p = u\ d,\ n^{++} = u\ u,\ \pi^+ = u\ \bar{d},\ \pi^- = \bar{u}\ d,\ n \rightarrow p\ \pi^-,\ \text{nil} \rightarrow \text{nil},\ d \rightarrow u\ d\ \bar{u}\ d$$
$$n^- = d\ d,\ n = u\ d,\ p = u\ d,\ n^{++} = u\ u,\ \pi^+ = u\ \bar{d},\ \pi^- = \bar{u}\ d,\ u\ d \rightarrow u\ d\ \bar{u}\ d,\ \text{nil} \rightarrow \text{nil},\ \text{nil} \rightarrow \text{nil}$$
$$n^- = d\ d,\ n = u\ d\ d,\ p = u\ d,\ n^{++} = u\ u,\ \pi^+ = u\ \bar{d},\ \pi^- = \bar{u}\ d,\ u\ d\ d \rightarrow u\ d\ \bar{u}\ d,\ \text{nil} \rightarrow \text{nil},\ \text{nil} \rightarrow d$$
$$n^- = d\ d,\ n = u\ d\ d,\ p = u\ u\ d,\ n^{++} = u\ u,\ \pi^+ = u\ \bar{d},\ \pi^- = \bar{u}\ d,\ \text{nil} \rightarrow \text{nil},\ \text{nil} \rightarrow u,\ \text{nil} \rightarrow d$$
$$n^- = d\ d\ d,\ n = u\ d\ d,\ p = u\ u\ d,\ n^{++} = u\ u,\ \pi^+ = u\ \bar{d},\ \pi^- = \bar{u}\ d,\ \text{nil} \rightarrow \text{nil},\ \text{nil} \rightarrow u,\ \text{nil} \rightarrow \text{nil}$$
$$n^- = d\ d\ d,\ n = u\ d\ d,\ p = u\ u\ d,\ n^{++} = u\ u\ u,\ \pi^+ = u\ \bar{d},\ \pi^- = \bar{u}\ d,\ \text{nil} \rightarrow \text{nil},\ \text{nil} \rightarrow \text{nil},\ \text{nil} \rightarrow \text{nil}$$

The final result is a consistent set of models, matching the early quark models held by physicists for the hadrons shown in this example.[9] In summary, REVOLVER embodies an approach to belief revision and discovery that illustrates the generality of the hill-climbing metaphor (Langley, Gennari & Iba, 1987). As long as domain knowledge can be expressed as relations among objects, and these entities are assumed to be conserved across these relations, REVOLVER can be used to construct models of some subset of objects in terms of other objects.

3.2 Experiments with Artificial Domains

We have seen that during belief revision REVOLVER may generate several candidate revisions but that it selects only one for implementation. Although this approach saves memory, it could potentially lead the system to take unreasonable time to reach a solution. To study this issue we ran two experiments, determining whether the time needed to arrive at a consistent set of models increases dramatically as REVOLVER's evaluation function is changed or as the system addresses more complex problems.

To understand REVOLVER's sensitivity to the weights in its evaluation function, we ran the system on examples where these weights were systematically varied. In particular, we looked at how changing the preference for minimum revision, as well as changing the initial models given to REVOLVER, affects the number of revisions required to reach a solution. Three domains have been studied. The runs analyzed in Figure 1 are from chemistry data (two reactions and five initial models), particle physics data (three reactions and six models) and artificially created data (six reactions and ten models). In each case, the program started with the simplest (theoretical) models possible given correct substance types, and used the same-type assumption during belief revision. Figure 1 shows that the system is not sensitive to variations in the revision weight except at very low values, and that this result does not change with different sets of initial models.

Our second experiment examined how problem complexity affects REVOLVER's ability to find solutions. We varied the data given to the program along two dimensions to determine the effect on the number of revisions needed to reach consistent beliefs. The first dimension was the *distance* of initial premises from a set of models known to be consistent. Based on these models, we gave REVOLVER new sets of models that differed from the 'correct' set by an increasing number of substances. For each initial model distance d we ran the system 30 times, each time deleting

[8] In this new domain, the same-type assumption holds once again.

[9] Similar results hold for the other hadrons, and can be replicated by asserting the appropriate initial models and particle reactions.

d components at random from the known-consistent models (within the limits of the same-type assumption) and using these as its premise models. The second varied dimension was the amount of *interdependence* among premise beliefs; interdependence of level n means that each object (substance) appears in n premises. This dimension is important because increasing interdependence increases the chance of a cascading effect, in which one premise revision causes another revision that causes still another and so on. For each initial model distance, three levels of premise interdependence were used, ten runs for each level.

Figure 2 shows that the number of revisions, averaged over each set of ten runs, increased in a nearly linear fashion as initial model distance was increased. In addition, this linear relationship was not greatly affected by increasing interdependence of premise reactions. However, search time did increase for some individual runs, during which REVOLVER could not find consistent models without overriding the same-type assumption. In short, the number of steps REVOLVER usually needs to find consistent beliefs does not seem to increase unreasonably with harder problems; hence the time tradeoff for using memory-efficient hill climbing seems to be within acceptable bounds.

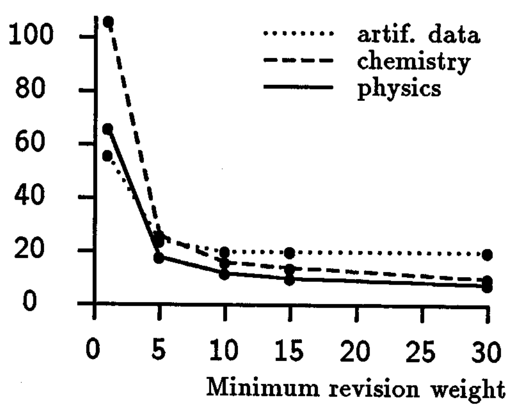

Figure 1. Search as a function of revision weight for three domains.

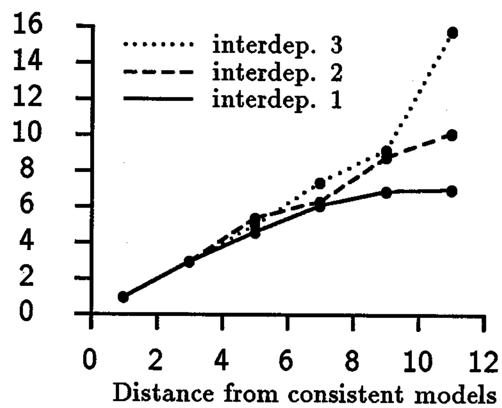

Figure 2. Effect of distance from consistent models and premise interdependence on search.

4. Discussion

REVOLVER subsumes its predecessor, the STAHLp [10] system (Rose & Langley, 1986). In addition to employing a more complex evaluation function (the previous system only used the conservatism criterion), REVOLVER makes revisions that deal with only one substance at a time, whereas STAHLp considers *all combinations* of revisions that resolve the entire inconsistency at once. Thus, the number of revisions STAHLp generates increases exponentially with the number of substances in each inconsistency, whereas this number increases linearly in REVOLVER. While REVOLVER often requires more revision steps than STAHLp to completely resolve an inconsistency, our experiments indicate that the number of steps increases linearly with initial distance from consistent models. Thus, although STAHLp's individual revision steps might cover more of the search space than could REVOLVER, the latter system usually requires less total computational effort.

We have seen that REVOLVER does not retain alternate revised premises, using its evaluation function to hill climb towards a solution. Assumption-based truth maintenance systems, or ATMS (de Kleer, 1984), embody a different philosophy. If used on the problems addressed by REVOLVER, an ATMS would save all premises ever acquired by the system and infer all possible consequences of

[10]This program was itself influenced by Zytkow and Simon's (1986) STAHL program.

those initial beliefs. Another approach taken by justification-based systems, or TMS (Doyle, 1979), would also save all premises but would maintain only one consistent set of beliefs, deciding which beliefs should be active or inactive at any given time. Both methods have drawbacks. An ATMS explores all solution paths when only one or a few might be sufficient. A TMS explores only one solution path at a time, but it uses ad-hoc choice rather than heuristics to decide which beliefs it should activate or deactivate. In short, both methods require computational effort that can grow excessive as the number of alternate premises (and inferences from them) increases over time. In contrast, REVOLVER requires less computation for many problems because it finds only one solution (theory) without saving alternate solutions along the way, using heuristics to guide search towards optimal theories. A related benefit of our approach is that it deals with the generation and selection of plausible revisions, tasks which truth maintenance systems do not address.

We have seen that REVOLVER embodies a two part approach to machine discovery. The first part transforms initial reactions into models of substances; the second part revises faulty premises into a consistent set, hill climbing through the space of potential premises. The program can construct observational and theoretical models in domains such as chemistry and particle physics, and despite its limited memory the system usually formulates correct models in reasonable time. However, much work remains to be done. For example, we should further explore the program's sensitivity to different evaluation functions. In addition, more heuristics are needed to further enhance these functions (e.g., to avoid cycling and to enable the program to search larger belief spaces). In summary, REVOLVER's integrated approach to belief revision and discovery has been a successful beginning, but more progress lies ahead.

Acknowledgements

This research was supported by Contract N00014-84-K-0345 from the Information Sciences Division, Office of Naval Research. We thank David Aha, Wayne Iba, Randy Jones, Bernd Nordhausen and Jim Wogulis for discussions that helped enhance this paper.

References

de Kleer, J. (1984). Choices without backtracking. *Proceedings of the Fourth National Conference on Artificial Intelligence* (pp. 79–85). Austin, TX: Morgan Kaufmann.

Doyle, J. (1979). A truth maintenance system. *Artificial Intelligence, 12*, 231–272.

Harman, G., Bienkowski, M. A., Salem, K., & Pratt, I. (1987). Measuring change and coherence in evaluating potential change in view. *Proceedings of the Ninth Annual Conference of the Cognitive Science Society* (pp. 203–209). Seattle, WA: Lawrence Erlbaum Associates.

Langley, P., Gennari, J. H., & Iba, W. (1987). Hill-climbing theories of learning. *Proceedings of the Fourth International Workshop on Machine Learning* (pp. 312–323). Irvine, CA: Morgan Kaufmann.

Nambu, Y. (1985). *Quarks: Frontiers in elementary particle physics.* Philadelphia, PA: World Scientific.

Quine, W. V., & Ullian, J. S. (1978). *The web of belief* (2nd Ed.). New York, NY: Random House.

Rose, D., & Langley, P. (1986). Chemical discovery as belief revision. *Machine Learning, 1*, 423–451.

Zytkow, J. M., & Simon, H. A. (1986). A theory of historical discovery: The construction of componential models. *Machine Learning, 1*, 107–136.