

Beyond Heuristic DENDRAL

by

Bruce G. Buchanan

Edward A. Feigenbaum

Joshua Lederberg

I. INTRODUCTION

Theory formation in science embodies many elements of creativity which make it both an interesting and challenging task for artificial intelligence research. One of the goals of the Heuristic DENDRAL project has long been the study of processes underlying theory formation. This paper presents the first steps we have taken to achieve that goal, in a program called Meta-DENDRAL.

The Heuristic DENDRAL project has concentrated its efforts on the inductive analysis of empirical data for the formation of explanatory hypotheses. This is the type of inference task that calls for the use of a scientific theory by a performance program, but not for the formation of that theory. When we started on Heuristic DENDRAL we did not have the insight, understanding, and daring to tackle ab initio the problem of theory formation. But now we feel the time is ripe for us to

turn our attention to the problem of theory formation. Our understanding and our technical tools have matured along with the Heuristic DENDRAL program to the point where we now see clear ways to proceed.

As always, the proper choice of task environment is crucial, but for us the choice was absolutely clear. The theory formation task most accessible to us is the task of forming mass spectral theory because the Heuristic DENDRAL performance program uses such a theory to elucidate molecular structures of organic chemical molecules. Hence, the notion of building a level of programs "meta" to the DENDRAL performance program.

The theory of mass spectrometry contains numerous statements about the fragmentation patterns of chemical molecules upon electron impact in a mass spectrometer. Because this is a new science, the theory is still expanding rapidly. In other words, there exists a possibility that a theory formation program could discover genuine extensions of the theory.

The Heuristic DENDRAL system already contains an excellent mass spectral theory. We, therefore, have a clear idea of what a "correct answer" is like. DENDRAL's theory is represented in at least two different forms at present, so that we have a fair idea of the issues involved in representing mass spectral theory for a program. A theory language of notations, data structures, and

primitive concepts (with which we are intimately familiar because we developed it) is available. People who are expert in mass spectral theory are members of the DENDRAL research team. Many programs for manipulating mass spectral data have already been developed and are ready to be exploited as Meta-DENDRAL tools.

The goal of the Meta-DENDRAL program is to infer the theory that the performance program (Heuristic DENDRAL) needs in order to solve problems of mass spectral analysis. The following table attempts to sketch some differences between the programs at the performance level and the meta-level.

	Heuristic DENDRAL	Meta-DENDRAL
Input	The mass spectrum of a molecule whose structure is not known (except, of course, in our test cases).	A large number of recorded mass spectra and the associated (known) molecular structures.
Output	A molecular structure inferred from the data.	A set of cleavage and rearrangement rules constituting a subset of the theory of mass spectrometry.
Example	Uses alpha-carbon fragmentation theory rules in planning and in validation.	Discovers (and validates) alpha-carbon fragmentation rules in a space of possible patterns of cleavage. Uses set of primitive concepts but does not invent new primitives.

In our view, the continuity evident in this table reflects a continuity in the processes of inductive explanation in science. Moves toward meta-levels of scientific inference are moves toward encompassing broader data bases and constructing more general rules for describing regularities in the data.

Beyond this level of Meta-DENDRAL there are still higher levels. Not all theory formation is as simple as the program described here assumes it is. For example, the representation of chemical molecules and the list of basic mass spectral processes are both fixed for this program, yet these are concepts which a higher level program should be expected to discover. Also, there is no postulation of new theoretical entities in this program.

But, again, higher levels of theory formation certainly do include this process.

The task of theory formation certainly can be and has been discussed out of the context of any particular theory. However, writing a computer program to perform the general task is more difficult than working within the context of one particular scientific discipline. While it is not clear how science proceeds in general, it may be possible to describe in detail how the scientists in one particular discipline perform their work. From there, it is not a large step to writing the computer program. Thus this paper attacks the general problems of theory formation by discussing the problems of designing a computer program to formulate a theory in a specific branch of science.

The general strategy of Meta-DENDRAL is to reason from data to plausible generalizations and then to integrate the generalizations into a unified theory. The input to the Meta-DENDRAL system is a set of structure-mass spectrum pairs. It receives essentially the same data as a chemist might choose when he attempts to elucidate the processes underlying the mass spectrometric behavior of a class of molecules. When the pioneers of the field turn their attention to a class of chemical compounds whose mass spectrometric behavior is not well understood, they must collect mass spectra for a number of the compounds and look for generalizations in the data. The

generalizations have to be tested against new data and against the established theory. If new data provide counterexamples, the generalizations are changed. If the generalizations are not compatible with the old theory either the old theory or the generalizations are changed, depending on the seriousness of the discrepancy, the nature of the statements involved, the scientist's commitment to the old theory, and so on.

This paper is organized by the three main subproblems around which the program is also organized. The first is to explain each individual spectrum, given the molecular structure associated with it. That is, determine the processes (or alternative sets of processes) which account for the experimental data. The second subproblem is to generalize the results from each spectrum to all spectra. In other words, find the common processes and sets of processes which can explain several spectra. The last is to integrate the generalizations into the existing theory in such a way that the theory is consistent and economical. Within each of the three main sections, the subsections indicate further subproblems which the program must solve.

II. FIRST SUBPROBLEM: EXPLAINING EACH SPECTRUM

The so-called "method of hypothesis" in science is sometimes

proposed as the essence of scientific work. Restating it, in a deliberately imprecise way, the method is roughly to formulate a hypothesis to account for some of the observed data and make successively finer adjustments to it as more observations are made. Very little is said about the details of a scientist's intellectual processes as he goes through the method. Thinking of hypotheses, for example, is an unsystematic and mysterious task which must be elucidated before the method can be programmed. That is the task we have designated as the first subproblem.

The program starts with individual spectrum- structure pairs as separate from one another. It constructs alternative explanations for each spectrum and then considers the spectra all together. An explanation, for the program, as for the chemist, is a plausible account of the mass spectrometric mechanisms which produced the peaks in the spectrum. The explanation is something like a story of the molecule's adventures in the mass spectrometer: certain data points appear as a result of cleavage, others appear as a result of more complex processes. At this stage of development of the theory, the chemist's story does not account for every data point because of the complexities of the instrument and the vast amount of missing information about mass spectrometry.

A. REPRESENTATION

The well-known problem of choosing a representation for the statements of a scientific theory and the objects mentioned by the theory is common to all sciences. In computer science it is recognized as a crucial problem for the efficient solution (or for any solution) to every problem. Improper choices of representation can forestall the solution for both humans and machines. At this stage there are no computer programs which successfully choose the representation of objects in a problem domain. Therefore we, the designers of the Meta-DENDRAL system, have chosen representations with which we have some experience and for which programmed subroutines have already been written in the Heuristic DENDRAL performance system.

It was natural to use these representations since the meta-program itself will not only interface with the Heuristic DENDRAL performance program, but is built up from the LISP functions of the performance program. Except for some supervisory routines, all of the routines used for solving the first subproblem (reasoning from data to plausible individual explanations) had already been written for use in predicting mass spectra. For example, the routine to break a bond is used in both the meta-program and the predictor phase of the Heuristic DENDRAL system. These common features will be described in more detail below.

Specifically, for this program, the input data are chemical structures paired with their experimental mass spectra:

((Structure-1 . Spectrum-1) ... (Structure-n . Spectrum-n)).

The representation of chemical structures is just the DENDRAL representation used in the Heuristic DENDRAL system. It has been described in detail elsewhere <1,2,4,7>: essentially it is a linear string which uniquely encodes the graph structure of the molecule. The mass spectra, also, are represented in the same way as for the Heuristic DENDRAL performance system. Each spectrum is a LISP list of dotted x-y pairs, where the x-points are masses of fragments and the y-points are the relative abundances of fragments of those masses.

Internally, a molecular structure is stored as a list of named atoms in the structure (e.g., C1, O2, C3) and the connections among the atoms. By putting the connections of the atoms on their property lists (along with other properties such as the type of atom (carbon, oxygen, etc.) and the number of hydrogens), manipulating chemical structures in LISP becomes a matter of changing values on property lists. This representation has been used in the mass spectrum predictor, for which the functions for manipulating structures were also developed.

The representation of the statements in the program's mass spectrometry theory is also the one used in the predictor phase of the Heuristic DENDRAL system. It is discussed in detail in

Section IV of this paper.

B. SEARCH

It is not clear what a scientist does when he "casts about" for a good hypothesis. Intuition, genius, insight, creativity and other faculties have been invoked to explain how a scientist arrives at the hypothesis which he later rejects or comes to believe or modifies in light of new observations. From an information processing point of view it makes sense to view the hypothesis formation problem as a problem of searching a space of possible hypotheses for the most plausible ones. This presupposes a generator of the search space which, admittedly, remains undiscovered for most scientific problems.

In the Heuristic DENDRAL performance system the "legal move generator" is the DENDRAL algorithm for constructing a complete and irredundant set of molecular models from any specified collection of chemical atoms. Each complete molecular structure is the terminus of some branch in the tree generated by the program. Intermediate nodes in the tree represent structures partially determined but with some chemical atoms as yet unallocated to the emerging structure. The primitive concepts of this generator are chemical atoms and bonds of molecules (or nodes and edges of graphs) and valence of atoms (number of allowable edges emanating from any node).

The problem of finding sets of mass spectrometric processes to explain each spectrum is also conceived as a heuristic search problem. As such, the structure of the computer program is easily described. In broad terms, the program contains (1) a generator of the search space, (2) heuristics for pruning the tree, and (3) evaluation criteria for guiding the search. Except for problems inherent in the task, then, the problems of such a program are reasonably well understood. These three main components of the heuristic search program are considered one at a time in the immediate discussion.

1. GENERATOR

Writing a computer program which solves a scientific reasoning problem is facilitated by seeing the problem as one of heuristic search. This is as true of the meta-program which reasons from collections of data to generalizations as for the performance system which reasons from one set of data to an explanation. For this reason we have called the process of induction "a process of efficient selection from the domain of all possible structures."<17>

For this part of the Meta-DENDRAL system, the generator is a procedure for systematically breaking apart chemical molecules to represent all possible processes occurring in a mass

spectrometer. In addition to single cleavages, the generator must be capable of producing all possible pairs of cleavages, all possible triples, and so forth. And, for each cleavage or set of cleavages it must be able to reproduce the result of atoms or groups of atoms migrating from one fragment to another. For example, after the single break labeled (a) in Figure 1 below, subsequent cleavage (b) may also occur. The result of (a) + (b) is the simple fragment CH₃.

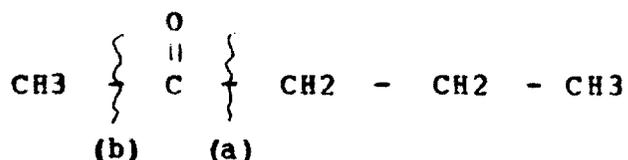


FIGURE 1

Or, for the same molecule cleavage (c) may be followed by migration of one hydrogen atom from the gamma position (marked with an asterisk) to the oxygen, as shown in Figure 2:

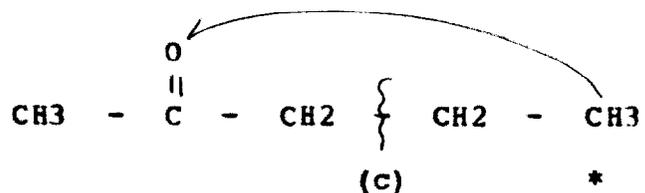


FIGURE 2

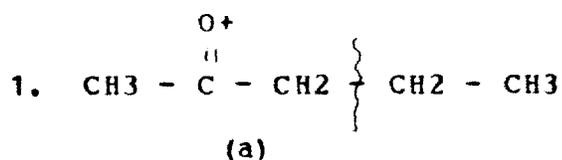
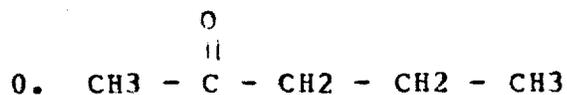
The resulting fragment, then, contains three carbon atoms, six

hydrogens and an oxygen (C₃H₆O). Since no such fragment results from simple cleavage or any combination of simple cleavages, it is necessary to postulate cleavage plus hydrogen migration as the origin of the mass spectral peak corresponding to this fragment and peaks corresponding to similar fragments for the mass spectra of other molecules. Historically, this double process was postulated for just these reasons and was subsequently confirmed experimentally.

The generator of the search space will postulate these processes as possible explanations of the mass spectral peaks at masses 15 (CH₃) and 58 (C₃H₆O) for this particular molecule. But it will also postulate the simple cleavage (b) in Figure 1 as the explanation of the peak at mass 15. And for the peak at mass 58 from the process in Figure 2 it will postulate the alternative migration of a hydrogen atom from the position labeled beta. From the generator's point of view these processes are at least as good as the more or less accurate processes shown in Figures 1 and 2.

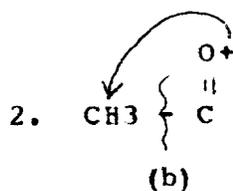
Another device is used by mass spectroscopists to explain their preference for the double cleavage (a) + (b) over the simple cleavage (b) for describing the mechanism resulting in the peak at mass 15 (CH₃) for the molecule shown. Mass spectroscopists often appeal to the localization of the positive charge in the charged molecule to explain why one peak appears in a spectrum

but another does not. It is known that only the charged fragments are recorded by the mass spectrometer. In this particular case (Figure 1) the positive charge initially resides on the C=O group of the molecule. For this reason, the simple cleavage (b) will produce a peak corresponding to the fragment C₄H₇O, the right-hand fragment resulting from that break. But since there is no positive charge on the CH₃ fragment after (b) above, the peak at mass 15 in the spectrum could not have resulted from this process. Because of the particular stability of carbon monoxide, however, the experts are willing to postulate that cleavage (a) is followed by loss of the uncharged C=O piece from cleavage (b), thus leaving a positive charge on the only other atom in the original fragment. The generator program must also manipulate charges then. It does this particular piece of reasoning in several steps, as shown in the steps of Figure 3, not just by the double process described above.

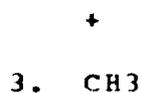


Place the plus charge on the O of C=O.*

Make cleavage (a) leaving left-hand side.



Migrate the charge from C=O to CH₂.



Final result, mass=15, after cleavage (b).

 *The correct notation from a chemist's point of view is O⁺. instead of O⁺ -- the dot representing a free electron on the oxygen. The Meta-DENDRAL program ignores this subtlety, although the Heuristic DENDRAL program gets this straight.

FIGURE 3

The primitive mechanisms of the generator are charge localization, cleavage, and group migration (where a group can be a positive charge, a single atom, or a set of connected atoms). The generator is a procedure for producing all possible charged fragments, not just all possible fragments, in other words. Putting these mechanisms together in all possible ways leads to an extremely large space of possible explanations for the peaks in the mass spectrum of a molecule. The pruning heuristics discussed in the next section alleviate that problem somewhat; for the moment let us turn to the actual design of the generator.

There are clearly several alternative ways of putting the three primitive mechanisms together into a generator of charged fragments. For example, chemists consider the initial localization of the charge before the possibilities for simple cleavage. But an alternative is to consider the possibility of cleavages first and then ask about the location of the positive charge. The generator actually programmed allows for changing the canons of generation to experiment with these alternatives, but only the set of canons described below has actually been used thus far.

At the first level of branching in the tree all possible single cleavages are performed on the original molecular structure resulting in all possible primary fragments. At the next level, the positive charge is assigned to all possible atoms in the

fragments. (Switching these two steps gives the same results and is closer to the conceptualization used by the chemist; it results in a less efficient program, however.) Starting with level 3 the procedure for generating successive levels is recursive: For each charged fragment at level n ($n > 2$) produce the charged fragments resulting from (i) cleavage of each bond in the fragment and (ii) migration of each group from its origin to each other atom in the fragment, where 'group' currently means 'positive charge or hydrogen atom'.

In general a group can be any radical in the fragment, and the generator can handle the general case. Limiting group migrations is done by having the program consult a list of items it will consider as candidates for group migrations. That list now contains only '+' and 'H', which reflects both the designers' and the chemists' preference for simple migrations. The generation tree is schematized in Figure 4.

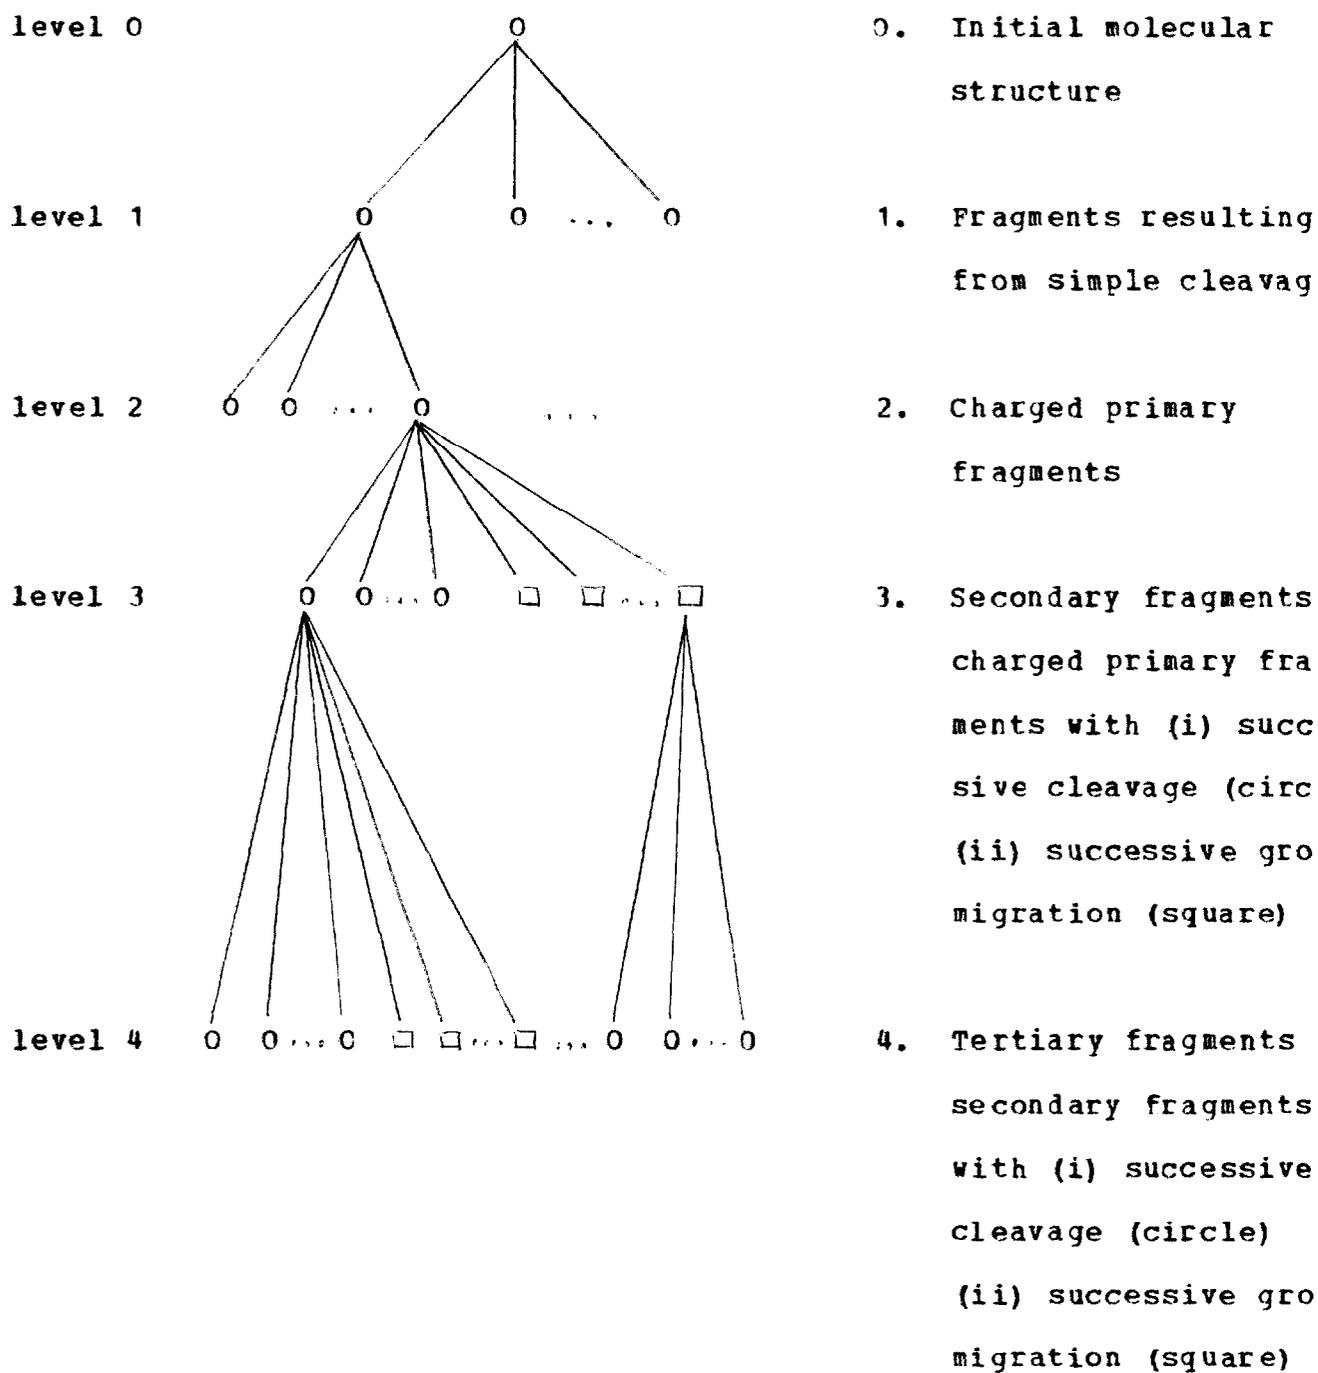


FIGURE 4

It is desirable to avoid the possibility of successive levels of some branch in this tree containing all the same atoms and differing only in the migration of atoms or of the charge. In this case duplications would result and the tree would be infinite. Thus the generating algorithm applies the group migration mechanism only to nodes in the tree which resulted from cleavage.* The generating process terminates when there are no more bonds to break in a fragment.

*In principle it is desirable to allow fragments resulting from n group migrations and only a single cleavage. Thus the algorithm is constructed to allow up to n group migrations after a cleavage. Currently n=1, so the procedure works as described in Figure 4.

2. PRUNING HEURISTICS

The size of the search tree may become quite large, for there are many combinations of mechanisms to consider. For example, if only bond cleavages are considered, the number of possible explanations for data points in the spectrum of a molecule is the number of ways of breaking combinations of the bonds in the molecule. The number of combinations of b bonds in a molecule

is $2^{*b} - 1$. This number represents a lower bound on the size of the search space, since secondary processes such as group migrations will be considered in conjunction with bond cleavages. As with any large search tree, then, pruning heuristics are essential for reducing the search.

Three simple pruning techniques are currently used by the program. (1) Since the result of breaking a pair of bonds (or n bonds) is independent of the order in which the bonds are broken, allow only one occurrence of each bond set; (2) Since mass spectrometric processes tend to follow favorable pathways, prune any branch in the tree which is no longer favorable, as evidenced by failure of a fragment's mass to appear in the mass spectrum; (3) Limit the number of allowable group migrations after each cleavage.

The first pruning technique is hardly a heuristic, in the sense of a risky technique, but its pruning effectiveness should be obvious. Duplications of nodes in the search space are unnecessary in this case and can be avoided by removing a bond from consideration after all possible results of breaking it have been explored. The second technique does carry an element of risk, because mass spectrometry theory includes no guarantee that every fragment in a decomposition pathway will produce a peak in the mass spectrum. In fact, the pruning can only be done after a complete cycle of cleavage plus migration because these processes

occur together in the mass spectrometer -- without the appearance of the intermediate fragments. The third technique also is truly heuristic since there are no theoretical reasons why group migrations might not occur in complex and exotic patterns between cleavages. It has been mentioned earlier that the current limit is one group migration after any cleavage. Migrations of two and three hydrogens are known in mass spectrometry, so these will not now be found by the program. Also, limiting allowable migrations to hydrogen atoms and the plus charge is somewhat risky since there are known cases of migration of larger groups, such as methyl (CH₃). The bias of mass spectroscopists toward simple mechanisms, however, leads us to believe that they would place little faith in exotic mechanisms as explanations of mass spectral peaks, at least not without other corroborating evidence.

3. EVALUATION

The search tree can be pruned effectively using the techniques just mentioned, with a substantial but not breath-taking reduction in search. Evaluation of alternative mechanisms is still necessary, either during generation or after it is completed, in order to distinguish the highly attractive explanatory mechanisms from those which are merely possible. Evaluation routines applied during generation can also be used to

limit search if threshold values can be established with confidence. This has not been done because of our lack of experience with the program so far. But even when the mechanisms are evaluated individually during generation, the program must be prepared to evaluate sets of alternative mechanisms later to find consistent and simple explanations. Evaluation of individual mechanisms and of sets of mechanisms will be considered separately in the following discussion, although it is clear even to the designers that these evaluation criteria are not independent.

We would be very happy to have the program rediscover just those explanatory mechanisms which expert mass spectroscopists have postulated, and even happier if the program postulated a set which mass spectroscopists accepted as a viable alternative to their own. So far, however, our evaluation criteria have not been strong enough to pick out mechanisms which experts think are plausible, while giving a low ranking to the others.

Without building in the biases of experts toward their current theory it is difficult to evaluate mechanisms at all. The concept of cleavage adjacent to the C=O (alpha-cleavage) in molecules like that of Figure 1 is well-established in the mass spectrometry literature. Without putting in criteria relating score to distance from non-carbon atoms, it is difficult to discover the attractiveness of alpha-cleavage. The criteria

must, it seems, be independent of any particular theory; they must be criteria for evaluating the worth of generalizations in any theoretical domain.

The program's evaluation routine presently contains only one a priori principle. In an attempt to measure the simplicity of the statements describing mechanisms, the program counts the number of primitive mechanisms necessary to explain a peak. Thus when there are alternative explanations of the same data point, the program chooses the simplest one, that is, the one with the fewest steps. Simple cleavage is preferred to cleavage plus migration plus cleavage, for example.

It is not desirable to use this ranking principle for pruning during generation, however, because there are cases where fragments can only come from complex pathways. For example, a simple mechanism and a complex mechanism may both explain the same peak, but a daughter peak can only be derived from the fragment along the more complex branch of the tree. The reason for this is simple: although different fragments may have the same mass (and thus explain the same peak), daughters of those fragments will not always have identical masses. So, the program must save the complex pathways as long as they may lead to fragments whose mass is not the same as the mass of fragments arrived at by a simple pathway.

The result of the generation process as described so far, with pruning and evaluation, is a set of candidate mass spectrometric processes for each structure which provides alternative explanations for data points in the associated mass spectrum. For instance, the program breaks the molecular structure shown in Figure ~~6~~⁵ at individual bonds or pairs of bonds to give the following information (atoms in the structure are numbered from left to right):

MASS EXPLAINED	PROCESS
103	cleavage: C2-C1 cleavage: C6-C7
89	cleavage: s3-C2 cleavage: C5-C6
75	cleavage: C4-C5
61	cleavage: S3-C4
60	cleavage: C4-C5 & C2-C1
57	cleavage: C4-S3
46	cleavage: S3-C4 & C2-C1
43	cleavage: C5-C4
42	cleavage: C4-C7 & C4-S3
29	cleavage: C2-S3
28	cleavage: C5-C6 & C4-S3

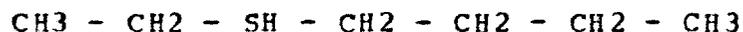


FIGURE ~~6~~⁵

In this example, the program used no migrations or charge localization information, for purposes of simplicity. The program explored all simple cleavages and found peaks corresponding to every resulting fragment but two.* For each of the successful fragments, the program broke each of the remaining bonds. From all the secondary breaks considered, the resulting fragments corresponded to only four additional peaks in the spectrum. So these four branches of the search tree were each expanded by one more simple cleavage. None of the tertiary fragments were found in the spectrum so the program terminated. It explained 11 data points in this way, the unfragmented molecule accounted for another one, and the program left 24 data points unexplained. Eight of these are explainable by hydrogen migration in conjunction with one or two cleavages, so just this additional capability adds substantial power to the program.

*The CH3 fragment was produced twice but there was no peak at mass 15 in the spectrum.

III. SECOND SUBPROBLEM: GENERALIZING TO ALL STRUCTURES

The method of hypothesis, mentioned earlier as a vague description of scientific work, suggests that a plausible hypothesis can be successively modified in light of new experience to bring a scientist closer and closer to satisfactory

explanations of data. Apart from the problem of formulating a starting hypothesis discussed above and the problem of terminating the procedure, it is not at all clear how the adjustments are to be made nor how to select the new experiences so as to make the procedure relatively efficient, or at least workable. These are well-known problems in the methodology of science. In other terms, the problem of successive modifications can be viewed as a problem of generalizing a hypothesis from one set of observations to a larger set.

The Meta-DENDRAL system does not yet contain a programmed procedure which solves this second main subproblem. The design of the procedure will be described only briefly in this section for this reason. The task is to generalize the explanations for each of the molecules into a consistent and simple set of explanations for the whole set of molecules.

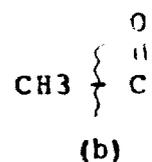
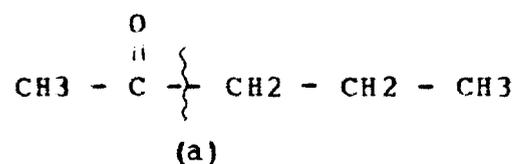
The input to this phase of the program is the set of explanatory mechanisms for each of the molecules in the original class. The output is one set of mechanisms which can explain every peak explained by the individual sets of explanations. If the program is successful, the output set of explanations will be a unified "theory" of the mass spectrometric behavior of all the molecules in the class. In operational terms, this means, at least, that the final set of explanations will be smaller than the union of all the individual sets of explanations.

A. REPRESENTATION

Perhaps the most pressing problem is to represent the candidate explanations in such a way that they can be compared and unified, without building in concepts which would beg the theory-formation question. The bonds in each molecule are now named only by the arbitrary numbering of atoms, the C1-C2 bond, for example. This is a convenient representation, for bonds in a single molecule but does not allow comparisons of bonds between molecules. What does the C1-C2 bond in one molecule have in common with the C2-C3 bond in another?

An attractive representation of a bond is a description of its environment: the types of atoms it links; the number of hydrogens, oxygens, nitrogens and so forth on each of the linked atoms; the types of bonds emanating from the linked atoms (single, double, triple); and so on for successive atoms and bonds away from the bond in question. This is attractive because it is easy to give the program just the right information for efficient comparisons. But this is the danger, too, for omitting "superfluous" information gives the program much too great a head start on the problem. It might discover what we want it to discover--the old principles-- but it would never discover anything new.

The most neutral representation for bonds we have found which still allows comparison is the pair of graph structures of the pieces of the molecule joined by the bond. Breaks (a) and (b) of Figure 1, then become:



Group migrations present a difficult problem. The most attractive representation now seems to be representing each migration as "before and after" graph structure pairs.

Graphs carry all the information about bonds and migrations; their main drawback is that they carry too much information. The program which looks for common features in subgraphs must be able to limit the amount of information it considers relevant. Some heuristics will be necessary--to limit the size of the environment that will be considered, for example. Again, at this point, there is a danger of telling the program too much about

what we want it to discover, or do not want it to discover.

B. SEARCH

As in any learning problem there will need to be considerable readjustment of the learned generalizations as new data are considered. It should also be expected that learning rates will be highly dependent upon the order of presentation of the data.

It is anticipated that the program will search for generalizations in roughly the same way as our mass spectroscopist friends say they do. The program will first choose one member from the class of compounds which is judged to be "typical", but "simple". By that, a chemist means a structure which is not so small as to present an inaccurate picture of the bond environments, and not so large and highly substituted as to bring in too many special problems. The program should allow these working definitions to be easily changed.

The fragmentation and migration patterns already observed for this typical member of the class will then be postulated to be typical for all members of the class. The rest of the program must look at patterns for additional molecules and adjust the initial hypothesis accordingly.

Again it should be noted that this design is certainly not the only alternative. It happens to be the one which appears most promising, in part because of its closeness to the procedure actually used by experts. Thus, it is the design we are now actively considering.

IV. THIRD SUBPROBLEM: INTEGRATING NEW STATEMENTS INTO EXISTING THEORY

Supposing the method of hypothesis to work, the scientist's problem does not necessarily end with the satisfactory formulation of a general statement explaining all the observed data. If he is working in a discipline for which there is no existing theory, he might stop here. But it is rare to be out of any theoretical context. Typically, the hypotheses are formulated as extensions of some existing theory.

In the course of formulating and modifying hypotheses, the existing theory usually serves as a guide for pursuing some hypotheses and not others. However, for reasons of simplicity the Meta-DENDRAL system currently ignores the existing theory until the end of its chain of reasoning. We are assuming here

that the classes of problems given to the system--the data to be explained--cannot be solved using the existing theory and that they do not conflict with the theory. These simplifying assumptions mean that the theory formulated by the program can be built up stepwise, class by class. New statements which are formulated are assumed to be consistent with the theory. Then the main problem remaining is to fit the new statements in with the old ones in an economical way.

One of the reasons we have rewritten the DENDRAL system's mass spectrum predictor was to separate the mass spectrometry theory from the LISP functions it drives. Making changes to the theory, then, does not require reprogramming, in the usual sense. Consequently, writing a program which updates the theory no longer seems to be an insurmountable task.

The problems of integrating new statements into the old theory are independent of the source of those statements. In order to study these problems we have written a program which (a) accepts new rules from human chemists and (b) updates the theory table of the program. The program for doing (a), called the dialog program, is not central to this paper, thus this section will focus on the work to accomplish (b), updating the theory. The dialog program represents one way of obtaining new chemical information to enter into the system. The automated data-analysis and theory-formation program incorporated in

Meta-DENDRAL represents another approach. The reason for discussing the former is that it serves as a simulation of the latter. Thus the problems of putting new chemical information into the system can be discussed from the standpoint of the dialog program, with which we have more experience.

A. REPRESENTATION

The mass spectral theory is represented in the program as a table of situation-action rules (S-A rules), patterned after Waterman's table of heuristics for good poker play.*

*D.A. Waterman, "Generalization Learning Techniques for Automating the Learning of Heuristics". Artificial Intelligence, 1:121 (1970).

Situations are predicates which are either true or false of a particular molecule or molecular fragment. (This is equivalent to Waterman's representation of situations as state vectors, but seems more economical for this program, because of the large number of terms which would have to be included in the state vector.) Actions are sequences of primitive mass spectral processes constituting rewrite rules for transforming one structural fragment into another. In this system, an action can also be another S-A rule, allowing nesting of rules in a manner quite natural to the current textbook descriptions of mass spectral theory.

The chemist interacting with the dialog program in time-sharing has already been given the explanation of primitive concepts known to the system which is given in Appendix A. For example, the chemist is told that he can use the functions BREAKBOND and ADDH, and is given some explanation of them. Mostly, the names of the LISP functions which are available suggest the effects of the function to the chemist.

All new rules must be expressed in terms of these concepts. This is a reasonable constraint in the dialog phase of this program because we have been willing to add new primitives whenever a chemist found some rule he could not express otherwise. For example, the concept of n-cleavage was added as a primitive to allow expression of a gamma-cleavage rule (NCLEAVAGE (QUOTE GAMMA)). This constraint also insures that the output from the dialog phase and the output of the automatic generalization program will contain the same terms.

As an additional means of maintaining parallelism between the dialog program and the automatic generalization program, the dialog program writes S-A rules on the basis of information obtained from the chemist, but does not allow him to write his own rules. The dialog program prompts the chemist for the information it needs in order to write the rules, and allows no opportunity for him to deviate from its prompting format. For each S-A rule, accompanying LISP code is written from the

information supplied by the chemist to define new situations and new actions, unless the chemist has merely put together old situations and old actions in new ways.

B. INTEGRATION

The output from the dialog phase is the same as the anticipated output from the generalization program: a set of S-A rules with accompanying definitions of situations and actions. The only expected difference is that the chemist can assign names to situations and actions during the dialog which are meaningful to him, whereas the generalization program must create its own symbols.

Inserting a new rule into the theory table requires an initial scan of the theory to find out whether the new rule describes (1) an entirely new situation, (2) a situation which subsumes an existing one, or (3) a situation which can be subsumed under an existing one. For cases (2) and (3), where there is some overlap between the new rule and some existing one (assuming for the moment there is only one), the actions of the rules must also be considered in deciding where in the table the new rule must go. After deciding where the rule must be inserted, the program adds the definitions of the new situation name and action name to the system.

Case (1), in which the situation of the new rule does not overlap with any existing situation, makes the insertion of new rules a simple matter of appending them to the theory table. The names of the situation and action are then given to the system with definitions, and the updating is completed.

Cases (2) and (3) are much more general and common, for there will often be some overlap between new situations and existing ones. The program must decide whether the overlap is interesting or substantial enough to warrant the extra work of reshuffling all the rules in order to insert the new one. It may decide that the situation is essentially a new one, in spite of some pieces of the situation appearing elsewhere. In this case it treats the rules as a case (1) rule and merely appends it to the table. This saving in processing at update time carries a corresponding cost in duplication of some parts of the rule. But the results of executing the rules will be the same in either case, with very little difference in execution time.

It is desirable to have the program nest rules some of the time, however, for reasons of conceptual clarity. If a new and an old rule describe exactly the same situation, for example, it inflates the theory needlessly to separate them as (i) $S \rightarrow A1$, and (ii) $S \rightarrow A2$. What we want is one rule of the form $S \rightarrow A1 \& A2$. Duplication of situations represents an extreme case of overlap which is also easy for the program to handle, once it

recognizes the duplication. It merely puts the second action into the action list of the old situation. Care is necessary, of course, to maintain independence of the actions so that they can be executed in any order with the same result.

Work on updating the theory has progressed this far. New rules can be added if the overlap with situations of existing rules either is null or the program decides it is insignificant. Or, a new rule can be added if the new situation is a duplicate of an old one. For the other cases, the general forms of (2) and (3), there is both the difficulty of finding the appropriate place to insert the non-overlapping pieces of the situation and the non-overlapping pieces of the action, and the the difficulty of deciding which of several rules should be updated to incorporate the new one.

V. CONCLUSION

The Meta-DENDRAL program described here is a vehicle for studying problems of theory formation in science. It is built upon the concepts and programmed routines already available in the Heuristic DENDRAL performance program, which uses a scientific theory to explain analytical data in organic chemistry. The Meta-DENDRAL system goes beyond the performance program, however,

in attempting to formulate the theory which the performance program will use.

The Meta-DENDRAL program works much like a chemist who is extending his theory of mass spectrometry by looking at collections of experimental results. The data, for both the chemist and program, are mass spectra and the associated molecular structures. By selecting some "typical" examples, first-order general hypotheses about the whole collection of data can be proposed. Then, by subsequent adjustments, the generalizations are modified to explain all the data. The new generalizations are then integrated into the existing corpus of theoretical statements in ways dictated by considerations of simplicity and personal preference.

The first version of the Meta-program, which is described here, suggests that the design is workable. But it accentuates the arbitrariness of our design decisions and raises the questions of what alternative designs would look like and how good they would be. It also raises a number of issues important to understanding scientific methodology in general. The design question is certainly one such issue. Others are questions concerning the criteria of acceptable generalizations, criteria of good scientific theories, and criteria for deciding on a set of primitive concepts for a theory. None of these general issues will be resolved satisfactorily in the context of this program.

Yet none can be resolved for this program without saying something about the general solutions.

APPENDIX A.

PRIMITIVE CONCEPTS OF MASS SPECTROMETRY

KNOWN TO THE DENDRAL PROGRAM

This list is taken from an outline given to chemists who define new mass spectrometry rules for the system. The functions at the front of the list are most primitive, those at the end are more complex, and in fact are built out of the simpler ones. The starred functions perform "housekeeping" on electrons and charge, without the chemist having to handle these details explicitly.

To the chemist this list serves as a reminder of the names and associated syntax of the "building blocks" available to him for defining new rules. To the present reader it is meant to illustrate the concepts already programmed into the system.

FUNCTION	ARGUMENTS	DESCRIPTION
HOUSEKEEPING FUNCTIONS:		
ADDCHARGE	atm	Assign a positive charge to atm.
ADDDOT	atm	Assign a free electron to atm.
IONIZE	atm	Assign a dot and a charge to atm.
PAIRELECTRONS	list;nolist	Look among the atoms of LIST for adjacent atoms with free electrons. Pair up the electrons to make an explicit bond unless the pair is named in NOLIST.
REMOVECHARGE	atm	Take away the positive charge from atm.
REMOVEDOT	atm	Remove the dot (if present) from atm.

FUNCTIONS FOR MANIPULATING STRUCTURE WITHOUT HOUSEKEEPING:

ADDH	atm	Put a hydrogen on atm.
CHANGEBOND	atm1;atm2;n	Add n (pos. or neg.) to the order of the atm1-atm2 bond.
JOINATOM	oldatm;atm;bond;atomtype;nodenum;	Bring atm into the structure -- attach to oldatm with bond order BOND. Give the atom type and node number specified.
REMOVEBOND	atm1;atm2	Remove the bond between atm1 and atm2.
REMOVEH	atm	Take a hydrogen off atm.

STRUCTURAL MANIPULATION FUNCTIONS WITH HOUSEKEEPING:

BREAKBOND	* atm1;atm2	Replace the atm1-atm2 bond with a pair of electrons. Try to pair any other free electron with one of the new free electrons.
BREAKRING	* atm1;atm2	Do the same as BREAKBOND when it is certain that the atm1-atm2 bond is in a ring.
ELIMINATEH	* atm	Eliminate a hydrogen from atm, leaving a free electron.
LOSEALPHARAD	* atm	Lose the largest radical alpha to atm.
LOSENEXTAD	* atm	Lose the largest radical adjacent to atm.
MAKERING	* atm1;atm2;bond	Join atm1 & atm2 with bond to form a ring.
MIGRATEH	* atm1;atm2	Move a hydrogen from atm1 to atm2, leaving a free electron on atm1 (unless atm1 = ANYATOM, in which case the H comes from atm1).
NCLEAVAGE	* n,pct	Break the nth bonds away from the heteroatoms in the molecule and assign intensity=pct*oldint/100. If n is 0 or (quote adjacent), the adjacent bonds are broken, 1=(quote alpha), 2=(quote beta), 3=(quote gamma).
NEWBOND	* atm1;atm2	Replace adjacent free electrons on atm1 and atm2 with an explicit bond.

FUNCTIONS FOR CARRYING OUT COMPLEX REARRANGEMENTS:

AMINERR	* atm1st	Migrate an H to node 2; Change bond between nodes 1 & 2 (+1); Break bond between nodes 1 & 2.
COELIM	* atm1st	Breakbond between 3 & 2.
CO2ELIM	* 1st	Breakbond between 4 & 3.
ELIM	* 1st	Remove an H from node 1; breakbond between nodes 1 & 2.
ELIM1	* 1st	Eliminate an H from node 2.
MRRFGT	* 1st	Change bond between nodes 1 & 2 (-1); Migrate an H from node 6 to node 2; Change bond between nodes 1 & 4 (+1); Break bond between nodes 4 & 5.

FUNCTIONS FOR HIGHER-LEVEL MASS SPECTROMETRIC PROCESSES:

ALLCLEAVAGES * nil

Break each bond in the molecule (MOLAT) and return a list of masses of fragments each conseed to intensities calculated MSTHEORY. (This is the default function to be called when no functional groups be identified.)

LOSENEXTRAD * atm
MAJORALPHACLEAVAGE * nil

Lose the largest radical next to atm. Consider all bonds alpha to the first in ATMLST and break the one which lose largest radical.

MCLAFFERTY * nil

Find which McLafferty Rearrangements will occur (single, dbl, two dbl) and do them.

MINORALPHACLEAVAGES * nil

Consider all bonds alpha to the first heteroatom in ATMLST and break each one except the major.

BIBLIOGRAPHY

(1) J. Lederberg, "DENDRAL-64 - A System for Computer Construction, Enumeration and Notation of Organic Molecules as Tree Structures and Cyclic Graphs", (technical reports to NASA, also available from the author and summarized in (12)).

(1a) Part I. Notational algorithm for tree structures (1964) CR.57029

(1b) Part II. Topology of cyclic graphs (1965) CR.68898

(1c) Part III. Complete chemical graphs; embedding rings in trees (1969)

(2) J. Lederberg, "Computation of Molecular Formulas for Mass Spectrometry", Holden-Day, Inc. (1964).

(3) J. Lederberg, "Topological Mapping of Organic Molecules", Proc. Nat. Acad. Sci., 53:1, January 1965, pp. 134-139.

(4) J. Lederberg, "Systematics of organic molecules, graph topology and Hamilton circuits. A general outline of the DENDRAL system." NASA CR-48899 (1965)

(5) J. Lederberg, "Hamilton Circuits of Convex Trivalent Polyhedra (up to 18 vertices), Am. Math. Monthly, May 1967.

(6) G. L. Sutherland, "DENDRAL - A Computer Program for Generating and Filtering Chemical Structures", Stanford Artificial Intelligence Project Memo No. 49, February 1967.

(7) J. Lederberg and E. A. Feigenbaum, "Mechanization of Inductive Inference in Organic Chemistry", in B. Kleinmuntz (ed) Formal Representations for Human Judgment, (Wiley, 1968) (also Stanford Artificial Intelligence Project Memo No. 54, August 1967).

(8) J. Lederberg, "Online computation of molecular formulas from mass number." NASA CR-94977 (1968)

(9) E. A. Feigenbaum and B. G. Buchanan, "Heuristic DENDRAL: A Program for Generating Explanatory Hypotheses in Organic Chemistry", in Proceedings, Hawaii International Conference on System Sciences, E. K. Kinariwala and F. F. Kuo (eds), University of Hawaii Press, 1968.

(10) B. G. Buchanan, G. L. Sutherland, and E. A. Feigenbaum, "Heuristic DENDRAL: A Program for Generating Explanatory Hypotheses in Organic Chemistry". In Machine Intelligence 4 (B. Meltzer and D. Michie, eds) Edinburgh University Press (1969), (also Stanford Artificial Intelligence Project Memo No. 62, July 1968).

(11) E. A. Feigenbaum, "Artificial Intelligence: Themes in the Second Decade". In Final Supplement to Proceedings of the IFIP68 International Congress, Edinburgh, August 1968 (also Stanford Artificial Intelligence Project Memo No. 67, August 1968).

(12) J. Lederberg, "Topology of Molecules", in The Mathematical Sciences - A Collection of Essays, (ed.) Committee on Support of Research in the Mathematical Sciences (COSRIMS), National Academy of Sciences - National Research Council, M.I.T. Press, (1969), pp. 37-51.

(13) G. Sutherland, "Heuristic DENDRAL: A Family of LISP Programs", to appear in D. Bobrow (ed), LISP Applications (also Stanford Artificial Intelligence Project Memo No. 80, March 1969).

(14) J. Lederberg, G. L. Sutherland, B. G. Buchanan, E. A. Feigenbaum, A. V. Robertson, A. M. Duffield, and C. Djerassi, "Applications of Artificial Intelligence for Chemical Inference I. The Number of Possible Organic Compounds: Acyclic Structures Containing C, H, O and N". Journal of the American Chemical Society, 91:11 (May 21, 1969).

(15) A. M. Duffield, A. V. Robertson, C. Djerassi, B. G. Buchanan, G. L. Sutherland, E. A. Feigenbaum, and J. Lederberg, "Application of Artificial Intelligence for Chemical Inference II.

Interpretation of Low Resolution Mass Spectra of Ketones".

Journal of the American Chemical Society, 91:11 (May 21, 1969).

(16) B. G. Buchanan, G. L. Sutherland, E. A. Feigenbaum, "Toward an Understanding of Information Processes of Scientific Inference in the Context of Organic Chemistry", in Machine Intelligence 5, (B. Meltzer and D. Michie, eds) Edinburgh University Press (1969), (also Stanford Artificial Intelligence Project Memo No. 99, September 1969).

(17) J. Lederberg, G. L. Sutherland, B. G. Buchanan, and E. A. Feigenbaum, "A Heuristic Program for Solving a Scientific Inference Problem: Summary of Motivation and Implementation", In Theoretical Approaches to Non-Numerical Problem Solving (R. Banerji and M.D. Mesarovic, eds.) Springer-Verlag, New York (1970). (Also Stanford Artificial Intelligence Project Memo No. 104, November 1969.)

(18) C. W. Churchman and B. G. Buchanan, "On the Design of Inductive Systems: Some Philosophical Problems". British Journal for the Philosophy of Science, 20 (1969), pp. 311-323.

(19) G. Schroll, A. M. Duffield, C. Djerassi, B. G. Buchanan, G. L. Sutherland, E. A. Feigenbaum, and J. Lederberg, "Application of Artificial Intelligence for Chemical Inference III. Aliphatic Ethers Diagnosed by Their Low Resolution Mass Spectra and NMR

Data". Journal of the American Chemical Society, 91:26 (December 17, 1969).

(20) A. Buchs, A. M. Duffield, G. Schroll, C. Djerassi, A. B. Delfino, B. G. Buchanan, G. L. Sutherland, E. A. Feigenbaum, and J. Lederberg, "Applications of Artificial Intelligence For Chemical Inference IV. Saturated Amines Diagnosed by Their Low Resolution Mass Spectra and Nuclear Magnetic Resonance Spectra", Journal of the American Chemical Society, 92:23 (Nov. 18, 1970). 1970.

(21) Y. M. Sheikh, A. Buchs, A. B. Delfino, B. G. Buchanan, G. L. Sutherland, E. A. Feigenbaum, and J. Lederberg, "Applications of Artificial Intelligence for Chemical Inference V. An Approach to the Computer Generation of Cyclic Structures. Differentiation Between All the Possible Isomeric Ketones of Composition C₆H₁₀O". Organic Mass Spectrometry (in press).

(22) A. Buchs, A. B. Delfino, A. M. Duffield, C. Djerassi, B. G. Buchanan, E. A. Feigenbaum, and J. Lederberg, "Applications of Artificial Intelligence for Chemical Inference VI. Approach to a General Method of Interpreting Low Resolution Mass Spectra with a Computer". Helvetica Chemica Acta, 53:6 (1970).

(23) E. A. Feigenbaum, B. G. Buchanan, and J. Lederberg, "On

Generality and Problem Solving: A Case Study Using the DENDRAL Program". In Machine Intelligence 6 (B. Meltzer and D. Michie, eds.) Edinburgh University Press (in press). (Also, Stanford Artificial Intelligence Project Memo No. 131.)

(24) B. G. Buchanan and T. E. Headrick, "Some Speculation about Artificial Intelligence and Legal Reasoning". Stanford Law Review, November, 1970. (Also, Stanford Artificial Intelligence Project Memo No. 123.)

(25) B. G. Buchanan, A. M. Duffield and A. V. Robertson, "An Application of Artificial Intelligence to the Interpretation of Mass Spectra. In Mass Spectrometry, 1970 (G.W. Milne, ed.) Wiley (in press).