

REPORT DOCUMENTATION PAGE

LEVEL



AD A 096802

DTIC FILE COPY

1. REPORT NUMBER HPP80-27		2. GOVT ACCESSION NO.	3. SECURITY CLASSIFICATION OF THIS REPORT LEVEL
4. TITLE (and Subtitle) The Plausible Mutation of DNA		5. TYPE OF REPORT & PERIOD COVERED Technical Rept.	
7. AUTHOR(s) Douglas B. Lenat		8. CONTRACT OR GRANT NUMBER(s) N00014-80-C-0609	
9. PERFORMING ORGANIZATION NAME AND ADDRESS Computer Science Department Stanford University Stanford, California 94305		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 116 Nov 80	
11. CONTROLLING OFFICE NAME AND ADDRESS Mathematical & Information Sciences Div. Office of Naval Research, 800 No. Quincy Street, Arlington, Va. 22217		12. REPORT DATE November 6, 1980	
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)   1232		13. NUMBER OF PAGES 31	
16. DISTRIBUTION STATEMENT (of this Report)  <b>DISTRIBUTION STATEMENT A</b> Approved for public release; Distribution Unlimited		15. SECURITY CLASS. (of this report)  15a. DECLASSIFICATION/DOWNGRADING SCHEDULE	
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)  Reports Distribution List, July 31, 1980			
18. SUPPLEMENTARY NOTES			
19. KEY WORDS (Continue on reverse side if necessary and identify by block number)  Heuristics, Automatic Programming, Evolution, Knowledge, Expert Systems.			
20. ABSTRACT (Continue on reverse side if necessary and identify by block number)			

DTIC ELECTE  
S MAR 25 1981 D

DD FORM 1 JAN 73 1473

EDITION OF 1 NOV 65 IS OBSOLETE. S/N 0102-LF-014-8601

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

471281 2 18 002

# The Plausible Mutation of DNA

Douglas B. Lenat<sup>1</sup>

November 6, 1980

Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By <i>Per Ltr. on file</i>	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
<i>A</i>	

Abstract

DNA may be regarded as a *program* for constructing and maintaining an organism. The field of Automatic Programming studies computer programs which synthesize new and different programs, or which modify and improve themselves. When DNA molecules do this, we call it Evolution. Biological research has to date identified several mechanisms which change DNA (substitution, insertion, deletion, translocation, inversion, recombination, segregation, transposition, etc.) Current theories assume the basic process of evolution to be *random* mutation (using these mechanisms) followed by natural selection. Early automatic programming systems were also built to work via this same *Random Generate and Test* process. But that mechanism failed, and we now recognize the reasons for that failure and the prescription for success. These results lead us to hypothesize that the generation of mutations may be highly non-random, that the dominant process of evolution in higher organisms is *Plausible-Generate and Test*. Long before our three billion line genetic *program* evolved randomly. Nature may have happened upon a more powerful method of *automatic programming*, such as heuristic search: the accretion and use of knowledge to guide the mutation process.

DTIC  
ELECTE  
S MAR 25 1981 D

<sup>1</sup> The author is an asst. professor of computer science at Stanford University, Stanford, Ca. 94305. D

## Introduction

Several biological mechanisms are known to result in altered DNA, mechanisms such as substitution, insertion, deletion, translocation, inversion, recombination, segregation, and transposition. In all known mechanisms, the major substratum for evolution can be said to be *random genetic events*.

Current theories assume the basic process of evolution to be random mutation (using these mechanisms) followed by a severe natural selection. The mutation may be fortuitous or neutral, the selection may be delayed for generations, but there is little doubt that the mechanisms are operating purely stochastically. If a new morphological structure, e.g., requires hundreds of local changes to the genome, then it is assumed that the large allelic variation of the population enables the coincidental combination of those changes.

DNA is a *program* for producing each protein required in the development and day-to-day maintainance of an organism. Viewed in this way, evolution is the process by which the DNA program is modified and extended. The branch of computer science which deals with computer programs altering and extending themselves, writing whole new programs, etc., is called *automatic programming*. This paper extends this superficial "Evolution as Automatic Programming" analogy. The results from twenty years worth of computer science experiments suggest a new hypothesis in biology.

The early (1958-70) researchers in automatic programming were confident that they could succeed by having programs randomly mutate into desired new ones. This hypothesis was simple, elegant, aesthetic -- and incorrect. The amount of time necessary to synthesize or modify a program was seen to increase *exponentially* with its length. Switching to a higher level language (the analogue of recombination and gene duplication) merely chipped away somewhat at the exponent, without muffling the combinatorial nature of the process. All the attempts to get programs to "evolve" failed miserably, casualties of the combinatorial explosion.

During the last decade, significant progress has been made in automatic programming, by providing such systems with great quantities of knowledge about programming in general and knowledge about the specific field in which the synthesized programs are supposed to operate. By employing this knowledge to constrain and guide them in their search, programs have finally begun to synthesize large new programs and modify themselves successfully. A study of the earlier "random mutation" work reveals that only after some such knowledge was added were the systems capable of successfully producing new programs or changes of more than a very few lines in length.

The key to the solution (using knowledge to guide the code synthesizer) appears quite simple in hindsight. How is the knowledge to be acquired? In the case of most automatic programming systems, it is provided by human experts. In the case of some programs, however, it is discovered automatically. The necessary machinery for learning from experience is not very complex: accumulate a corpus of empirical data and make simple inductive generalizations from it. The first requires some kind of memory, the second requires some kind of pattern-matching ability. Processes

similar to memory and matching are well-known to exist already (reliable information storage in nucleic acids, reliable matching of tRNA to mRNA at ribosomes). Certainly the needed processes (memory and pattern-matching) are orders of magnitude more elementary than, say, the functioning of the immune system and the central nervous system.

From this we are led to hypothesize that the generation of mutations may be highly non-random. Instead of "Random-Generate & Test", the dominant mechanism of evolution in higher organisms may be "*Plausible-Generate and Test*".

Suppose one were given five years to build a large computer program to forecast weather, and one knew little about programming *or* meteorology. Then it's clearly cost effective to take a couple years to develop some expertise in both fields. Similarly, while it is *possible* that nature evolved a three billion line program using only recombination, gene duplication, etc., it might be much more efficient to record and use knowledge: general knowledge about evolving and specific knowledge about the particular species itself and its genetic ancestry. In the past billion years, nature may have happened upon this more powerful method of "automatic programming": building up a body of knowledge to guide the mutation process.

How might this work? Some of the DNA records past states of the genome, and patterns in that record may be noticed and exploited. For example, consider cephalo-pelvic proportion (the relation between an infant's biparietal diameter and its mother's pelvic diameter.) If skull size of some species were to increase significantly, the females would have great difficulties giving birth, and the members of the population having such an increase would be selected against. The only exception is when the species' mean pelvic diameter simultaneously increases (fortuitously). Thus, looking back over a genetic history of a successful species, it would appear that increases in skull size are almost always accompanied (or immediately preceded) by increases in pelvic diameter. Once such a pattern is noticed, it can be used to guide future mutation: to encourage specific related groupings of mutations. When an increase in skull size is going to happen (a mutation occurs in the appropriate genes of the DNA in a germ line cell), a simultaneous increase in pelvic diameter should be made. A species would be better off if it could recognize and use such patterns -- such heuristics. In this case, the heuristic said "IF biparietal diameter is increasing, THEN increase the chance of pelvic diameter increasing." Many more heuristics are illustrated in Appendix 1.<sup>2</sup>

Consider a species capable of storing its genetic history, noticing empirical regularities in it, and using them to guide constellations of interrelated mutations in the future. Its rate of evolution might dwarf that of species which had to rely on fortuitous co-occurrences of random genetic events. Notice there is no inherent "direction" that such plausibility constraints are defining; rather, it is simply a

---

<sup>2</sup> A simple form of the cephalo-pelvic proportion heuristic could be implemented just by locating the genes determining these dimensions next to each other along the genome; thus the chance of a mutation affecting both simultaneously is great. This doesn't account for the same *direction* of change of both, nor can all of the heuristics present, e.g., in Appendix 1, be implemented by judicious placement of genes.

mechanism for avoiding what seems empirically to be deleterious and for seeking what seems empirically to be advantageous. Certainly there is nothing surprising in this; many creatures compile their experiences, in hindsight, into heuristic rules which guide their future *behavior*. This paper is suggesting that it may also be true of DNA.

Species whose evolution was guided by heuristics (compiled from the species' genetic history) would be better adapted at *evolving*. Their rate of evolution would be higher, and the fraction of offspring having a favorable co-occurrence of mutations would be elevated. Their DNA would be longer and largely unexpressed (containing much information which is historical and useful for inferring regularities in evolution but not needed for the maintenance of an adult organism). By also using this historical record for developmental functions, its integrity would be assured over many generations; ontogeny of such creatures would resemble a recapitulation of their phylogeny. The obvious hypothesis that this is leading to is that while evolution *began* as random generation, by now the evolution of most higher animals and plants may be under the guidance of a large corpus of heuristics, judgmental rules abstracting a billion years of experience into prescriptions and (much more rarely) proscriptions regulating and coordinating clusters of simultaneous mutations. See Appendix 1 for an example of a set of such rules and how they work together to design an improved organism. Random mutation would still be present, but in higher organisms its effect might be mere background noise.

### Lessons from Automatic Programming

We begin by sketching the "DNA as program" analogy. Information in the DNA molecule<sup>3</sup> is essentially in secondary storage analogous to magnetic tapes or disks; it must be swapped in to core -- i.e., copied from secondary storage into main memory -- (by mRNA), and brought to a processor (ribosome) to be run. The ribosome translating an mRNA into an amino acid sequence resembles a Turing machine [Minsky 67] reading along its input tape and writing out a new one. Feedback closes this loop (e.g., via production of repressor proteins) and raises the power of the mechanism to that of a universal Turing machine. The sophistication of the system is best displayed during the development of the fetus, when many delicate changes in gene expression must be coordinated. Only about a tenth of the four million genes in human DNA code for known proteins; the function of the other gene "subroutines" may include regulating pathways -- developmental, metabolic, and perhaps *evolutionary* ones.

The current stock of mutation methods is presumed to be adequate to account for the evolution of all DNA programs. That is, random changes occur in the sequence, manifest themselves as mutated progeny, and are judged by Natural Selection for

---

<sup>3</sup> Each nucleotide contains two bits of information, since there are four possible bases it could contain. Three nucleotides in a row form an instruction or codon. A codon contains 6 bits of information, so there are 64 possible instructions. The task of the program is to assemble a sequence of amino acids (a protein), and each codon specifies what the next amino acid should be, or else says STOP.

fitness. The DNA program for even such a complex organism as the interested reader is assumed to have developed by such a random generate & test progression.

We in Artificial Intelligence (AI) now know only too well the weakness of doing automatic programming by random changes of (and random additions of new) program instructions. Certainly it *can* be done, but it is extremely slow. It would be more accurate to say that AI researchers *today* have that intuition (that the combinatorics of the situation are deadly) but when the *first* AI researchers tackled this problem they didn't have the benefit of hindsight, of experience with searching. They quite naively but reasonably assumed that if you wanted to tell a program *what* to do, without telling it precisely *how*, then you'd have to employ some kind of random program generator, and follow it up with a test to see if the program was the desired one. As R. M. Friedberg (then at IBM) said in 1958,

*"Environment dictates what problems must be dealt with, but not how to deal with them... It is difficult to see a way of telling it what without telling it how, except by allowing it to try out procedures at random or according to some unintelligent system and informing it constantly whether or not it is doing what we wish."*

That is, computer scientists' intuitions then were precisely in agreement with biologists': the adequacy of random generate & test. Over the last twenty years, several painful research experiences have changed those computer science intuitions; we now examine some of those experiments.

The first effort along these lines was Friedberg's. His program searches through the space of all machine-language programs containing 64 instructions. It replaces each instruction in turn, looking for a local maximum of performance, and then repeats this procedure over and over again, a hundred times a second on an IBM704. The "environment" in this case is a specification of the desired behavior of the target program, the one which we want to have automatically synthesized. In each generation, the mutant programs whose behavior most closely resembles that of the target survive. This gradual approach to competence is termed *hill-climbing*, because it is akin to trying to find the top of a hill by taking a few steps in all directions, finding the one which got you the highest, moving there, taking a few steps in all directions, etc., etc.

When the target program was a couple instructions long (e.g., adding two 1-bit numbers), it took hundreds of thousands of generations to evolve such a program. When the target program was longer, say 5 or 6 lines long, it rarely had appeared even after millions of generations. But the immense number of generations required was not the biggest surprise:

To his shock, Friedberg found no stable islands in the search, that gradual hill-climbing was no better than generating an *entire program* from scratch each time. He built a system which tried completely new computer programs every "generation", which simply put together a new, random sequence of machine language instructions, ignoring its "parents'" design completely no matter how close their behavior was to that of the desired target program. This random program generator out-performed his gradual hill-climbing program-evolver every time.

What is the problem with hill-climbing? The most devastating phenomena are the frequent *local maxima* upon which a hill-climber gets trapped. He comes to the top of a small hill, and *any* small step, regardless of direction, will take him downwards, hence he stays where he is. There may be a much higher hill nearby, but he would have to do down into a valley before he could start up that next hill, so he never finds it. Human mountain climbers may have their vision obscured by false peaks, but their knowledge of mountaineering guides them onwards. They may have no certain information about the true mountaintop, yet be able to break out of local maxima using their past mountaineering experiences to generalize from (i.e., they are using empirical induction, not teleology.)

Almost all machine language programs are local maxima: to convert a useful one into another useful one requires altering many machine instructions simultaneously. The only way that Friedberg was ever able to get any successes out of the program-evolver was by building in some heuristic rules to guide its search for new programs:

If a program fails, lower the chance of selecting a program with *any* instructions unchanged from this one.

If a program succeeds, reward all its component instructions, i.e. increase the chance of selecting a program with many of the same instructions in the same locations as this program.

Do local optimization of each instruction in turn

Partition a problem and deal with its parts in order of difficulty.

Prime the system by telling it which data bits are the input, and which are the output.

His final result [Friedberg 59] is that "HOMER [completely new programs] makes large-scale changes upon failure, and surpasses SAMSON [hill-climbing mutations]. THALES [incorporating the heuristics listed above], on the other hand, undertakes only small changes; but those changes made are likely to be in the right direction."

One trouble with machine language programs is that they are doubly unstable: a small change in their flow-chart may engender an enormous number of changes in which locations in memory contain which instructions; conversely, a small change in the contents of some core locations may dramatically change the function computed by the program. Maybe the right level to work at, then is that of *flowcharts*.

Fogel, Walsh, and Owens decided in 1966 to attempt something very much like this: their program roamed about in the space of finite state automata, using operations close to those that we would have for mutating flow charts: redirecting arrows, adding nodes, relabelling arcs, etc. Fogel defined "intelligence" to mean the ability of a finite state automaton to anticipate its environment, its predictive power. Each generation, his program would select a mechanism of mutation from the following table:

60%	Change one of the next-input predictions (arc labels)
35%	Redirect an arrow (change its terminus)
3%	Add a whole new state (node) to the machine
2%	Eliminate an entire state from the machine

Once it selected a mechanism, his program altered the finite state automaton in that way. Fogel let his automatic programming system run for five generations, keeping the three best offspring in each generation (the ones with the highest percentage of

correct predictions so far), then he had them each make another prediction, and he inputted one more symbol (the next event from his environment). The process then iterated. The details of this work are reported in [Foget *et al* 66].

In one experiment, Fogel fed in the sequence (101110011101)\*. But this is 8/12, or 75% 1's; hence a very good guess, a good machine, is one which always returns (predicts) "1". This is what was settled on, in fact; a local maximum, a local peak from which it was impossible to escape with only slight variations. A similar problem occurs when one tries to synthesize a program to predict whether a number is prime or not (it's always easiest and best to simply guess that the answer is No.)

Incremental approaches to competence didn't seem to be working, yet if Fogel allowed large simultaneous variations, he would have had even worse behavior. He says:

*"The efficiency of pure trial and error exploration is sharply reduced with an increase in the dimensionality of the domain being explored. As long as the investigator is interested only in a single aspect of his environment, random exploration may prove worthwhile, but as soon as he attempts to map a domain of more practical interest he encounters so many possibilities that only carefully-guided trial-and-error exploration is likely to prove profitable... In man's initial exploration of the unknown, the scientific method would have been a luxury; however, with the increased scope and depth of his inquiry, use of the scientific method becomes an absolute necessity."*

What, then, is the solution being proposed? Flowchart-modifying should be guided by knowledge: knowledge about how to design and carry out telling experiments rather than random modifications, and knowledge about whatever task domain the synthesized program is supposed to perform in.

Consider the case of writing a program to test a number for Prime-ness. One general piece of programming knowledge is that a program should begin with some initializations, enter a computational loop, and ultimately return some value. Any flowchart not having that structure can be immediately eliminated from consideration. A general piece of knowledge looks at the definition of prime numbers, sees that it specifies "...whose only divisors are 1 and N", and recognizes this as a constraint on the flowchart: the central loop should terminate early with a "not-prime" answer sometimes, and if the loop runs to completion then the answer should be "is-prime". A specific domain-dependent piece of knowledge is that there are many primes and many non-primes, so any flowchart which always returns Yes (or always returns No) is bound to be wrong.

By employing a collection of such pieces of knowledge, the space of allowable flowcharts shrinks dramatically in size. The chances of finding a successful flowchart are raised dramatically.

Arthur Samuel, working at about the same time as Fogel, wrote his famous checker-playing program [Samuel 67]. It was designed to get better and better over time, by gradually improving its scoring polynomial (a function that evaluated the overall worth of a checkerboard position from, say, Red's point of view.) Samuel found it important to add several heuristics to guide the mutation of his scoring polynomial:

The first term should always be: *a move exists* (= viability)  
 Let A & B compete, and kill the loser (= natural selection)  
 Recall your earlier predictions, and rate them in hindsight  
 Artificially lower the coefficients of new terms, to forestall wild initial fluctuations  
 Count a recent fluctuation more heavily than an old one  
 Never have more than 16 terms (of the 38 known terms) at a time  
 After each series of games, drop the term with the lowest (magnitude) coefficient  
 It's worth risking introducing a *few* of the 38x38 cross-terms  
 AFTER 1967: Separate polynomials for opening, midgame, and endgame  
 AFTER 1967: Group parameters into clusters (signature types)

My own research in Automatic Programming recapitulated much the same error [Green *et al* 74]. I began in 1972 with a program called PW1, which had a few templates or schemata for recursive LISP functions, and which had a set of 10-20 functions it could plug in for each function mentioned in the schema. One of the templates was:

```
F(x) =df [λ (x) IF f1(x)=b1 THEN f2(x)
           ELSE f5(f3(First-element-of(x)),
                 f4(All-but-1st-element-of(x)))]
```

The program picked a random instantiation and mutated it until its input/output behavior agreed with the example I/O pairs which comprised the specification of the desired program. For instance, suppose the desired target program was one which found the smallest element of a list of numbers *x*. The user would type in a few I/O pairs as examples, such as

```
Input (1 3 5 0 8), Output 0
Input (9 8 7 6), Output 6
Input (1 3 5 7), Output 1
```

PW1 randomly chose a function to substitute for *f1*, *f2*, etc., and then ran the resulting program for *F* on each of the Inputs above. The values actually returned by the *F* function were compared with the stated target Output values, and if *F* was not yet in agreement its definition (choice of *f1*, *f2*,...) was mutated randomly. In the above schema, agreement with the list of input/output examples could be achieved when *f1* was instantiated as All-but-1st-element-of, *b1* as EmptyList, *f2* as First-element-of, *f3* as the Identity function, *f4* as *F* itself, and *f5* as Smaller (a function which returns the smaller of its two arguments). The final definition of *F* is then a definition of the function Smallest-element-of.

The simple function schema above can be instantiated in many ways, to yield definitions of Largest-element-of, Smallest-element-of, Length, Has-odd-length, Reverse, Contains-repeated-elements, Sort, and (unfortunately) millions of others. The first attempts had to be halted after hours of computer time had been extended fruitlessly seeking a valid definition of Smallest-element.

My first intuition was to fix this by having the definition gradually evolve. To this end, several mutations were made simultaneously by the system, and the one which had I/O results most closely matching the user-provided examples was chosen as the survivor in the next generation. Surprisingly to me at the time, this was not noticeably better than the original, completely random generation scheme.

PW1 did eventually synthesize several short target programs, but only after I adopted the method of supplying it some frequency hints (e.g., First-element is the most likely function to try for  $r_1$  in the schema), some applicability constraints, and a few simple ways in which to look directly at the I/O pairs in constraining which functions to try (e.g., IF the Outputs are always members of the Input lists, THEN  $r_5$  must be a function whose output is always one of its inputs).

My last automatic programming effort [Lenat 75] was the PUP6 program. It took a high-level specification of the desired behavior of a program (a dialogue in a tiny subset of English) and synthesized a target program meeting those specs. PUP6 was able to write a simple classificatory concept formation program (similar to [Winston 70]), an airline reservation system, and a grammatical inference program. It managed this by drawing upon a huge body of information about programming, concept formation, and inference.

Very recently, impressive synthesized programs have been produced from Cordell Green *et al.*'s PSI system [Barstow 79]. Their automatic programming system is guided by *hundreds* of rules about programming in general and about the task domain of the target program (the one being synthesized) in particular.

A similar solution was found to the problems of knowledge engineering, of building large expert systems for tasks such as medical diagnosis [Feigenbaum 77], mineral exploration, and mathematical theory formation [Lenat 79]. In each case, the program contains many (typically several hundred) heuristic rules, which guide its behavior, which suggest plausible moves for the programs to follow and implausible ones for them to avoid.

All our experiences in AI research have led us to believe that for automatic programming, the answer lies in *knowledge*: add a collection of expert rules which will guide code synthesis and transformation. Each rule is a kind of compiled search, a bit of condensed hindsight. While far from complete or foolproof, they are nevertheless far superior to blind changes in program instructions (Friedberg) or flowcharts (Fogel) or even mutation of duplicated program chunks (Lenat).

#### Idea #1: Add heuristics to DNA

Finally, we are ready to turn to the biological analogue of this idea. Just as automatic programming taught us to guide program synthesis and transformation by heuristic rules, so it might be cost-effective for evolution to be guided by heuristic rules. Appendix 1 presents a small example of a body of heuristic rules which are general and plausible, and which work together efficaciously to guide the evolution of a simulated organism.

Can we extend the DNA *qua* program analogy by somehow adding knowledge to the DNA, knowledge about which kinds of mutations are plausible, which kinds have been tried unsuccessfully, what combinations have and have not performed well in the past, etc.? That is, can we imagine what it might mean to turn DNA's random mutant generator into a plausible move generator? If there is a way to encode such

knowledge, such heuristic guidance rules, then we might expect that an organism with that kind of compiled hindsight would evolve in much more regular, rapid a fashion. The "test" would still be natural selection, but instead of blind generation the DNA would be conducting (and recording) plausible experiments.

What would such heuristics "look like"; i.e., how might they be "implemented" in the DNA program? Almost surely they would be written in the alphabet of bases, but their interpretation might not be as codons for proteins (in which case their expression would have to be suppressed.) At times of reproduction, however, they would specify allowable (and prevent other) changes to be made in the new copy. That is, they would sanction certain complex copying "errors" (e.g., statically by inserting noncoding sequences, or dynamically by interfering with the repair polymerases) and prevent others (e.g., via site-specific repair enzymes). The IF- parts of such "IF...THEN..." heuristics could be almost completely specified by position (proximity to genes to which the heuristic wishes to refer), and the start of such a heuristic would have to be signalled by some special sequence of bases (much like parentheses in Lisp). Each heuristic could have some demarcated domain or scope. Thus, "use a repressor/anti-repressor mechanism rather than an induction mechanism" might hold true for a patch of DNA which synthesized the organism's most important enzymes, and it would be easy to specify the scope by placement along the genome. So-called mutation "hot-spots" are a unary example of this kind of heuristic; heuristics taking more than one "argument" would of course be much more powerful, just as the site-specific mutators are more powerful than a global increase in the overall mutation rate could ever be. The "THEN..." part of a heuristic could direct gene rearrangement, duplication, placement of mutators and intervening sequences, etc.

Perhaps more likely would be for each heuristic to code for a very rarely-expressed protein. The heuristic could code for (or regulate) an enzyme which reentered the nucleus, "matched" against some number of patterns in the DNA, bound itself to those regions (the "IF" part), and thereby increased the chance of a certain type of mutation occurring at those regions (the "THEN" part). Such an enzyme might be produced in such small quantities, and with such small frequency, that it would be unlikely to be noticed in most cases. Its effects would be felt only if it affected germ line cells, and it might only be expressed in them. A final possibility is that it would be expressed only during embryogenesis, that each neonate's germ cells' DNA has *already* been altered, thus determining (to within sexual recombination and random mutation) the spectrum of changes which it might potentially pass along to its offspring.

#### Idea #2: They may already be there

Nature might already have become as good at programming in the last billion years as we have in the last forty. DNA might have *already* evolved from random generate & test into an expert program (expert at mutating itself in plausible ways). Since the heuristics deal with DNA subsequences, and they themselves are also DNA subsequences, they (or at least *some* of them) might be able to modify, enlarge, improve themselves and each other. That is, by now the heuristics themselves may

be developing under heuristic guidance: rules which encapsulate a billion years of experience at devising and changing and using heuristics.

What I conjecture is that Nature (i.e., natural selection) began with primitive organisms and a random-mutation scheme for improving them. By this weak method (random generation, followed by stringent testing), the first primitive heuristics accidentally came into being. They immediately overshadowed the less efficient random-mutation mechanism, much as oxidation dominated fermentation once it evolved.

Each heuristic proposes a plausible change (call it  $\Delta$ ) in the DNA. The progeny which incorporate  $\Delta$  (call them  $\Pi^\Delta$ ) also get a new heuristic indicating that that kind of change has been made and is good. This might be as simple as adding one new noncoding sequence inside that mutated gene. It might be as complex as producing a whole new mutated gene and keeping the old one around as a pseudogene. The progeny  $\Pi$  which do not incorporate  $\Delta$  get no such heuristic. If  $\Pi^\Delta$  is viable, then the new heuristic it contains will have proven to be correct. "False" heuristics die out with the organisms that contain them.

Consider a very simple example. Here is a mechanism which embodies the heuristic "If a gene has mutated successfully several times in the recent past, then increase its chance of mutating in the next generation, and conversely". All we need to posit is that somehow a short noncoding sequence -- we'll call it an *asterisk* -- is added to a gene each time it mutates. To see how this operates, consider human DNA: any genes which have several such asterisks testify that they have been mutated successfully, advantageously, many times in the past; genes with few or no asterisks suggest that modifying them has always led to detrimental changes in the offspring. All we need now do is propose some mechanism (e.g., stereochemical) whereby genes with many asterisks are more likely to be mutated, duplicated, etc., than genes with few or none. Since the asterisks provide no specific benefits to the individual, they will gradually be lost over time, so that when a gene no longer *should* be mutated, its asterisk count will slowly decline over several generations. Whether or not it was ever actually adopted, the power of this simple mechanism is clear.

As the species evolves, so do the heuristics. One big lesson from the AM program [Lenat 77] was the *need* for new heuristics to evolve continuously. Otherwise, as animals got more and more sophisticated, they would begin to evolve more and more slowly (random mutations, or those guided by a fixed set of heuristics, would become less and less frequently beneficial to the complex organism, less frequently able even to form part of a new stable subassembly, as Simon suggests).

Using a higher level language like gene duplication, rearrangement, and recombination, instead of sequence mutation, would give only a constant factor of improvement (i.e., as if we did automatic programming by random changes in LISP programs instead of in assembly language programs), and this constant must fight against the rapidly decreasing number of organisms born each year as one ascends the evolutionary ladder. Thus we expect a phylogenetic increase in the number of heuristics, the sophistication of those heuristics, and the relative proportion of DNA devoted to heuristics.

Heuristics condense past history into judgmental rules. They are kernels of knowledge which, if only they'd been present earlier, would have gotten us to our present state much faster. A heuristic prescribes some action which is appropriate in a given kind of situation, or proscribes one which is dangerously inappropriate. They are *useful* because the world is continuous: if several features of the current situation are similar to some earlier one, then the set of actions which are -- and are not -- appropriate will probably also be similar. Thus it is cost-effective to compile experiences into heuristics, and to then use the heuristics for guidance. Even if the environment is rapidly changing, some useful heuristics may be extractable, so long as there are some regularities to those changes -- to the environment. Physics equations are no less useful just because the world is constantly changing -- if anything, they are *more* useful than they would be in a static world where abstraction would be a luxury. So it is with bioheuristics for evolution: by embodying a deep enough model of the past, the heuristics can cope with a diversity of future problems.

Until the Eurisko program was conceived [Lenat 77], this would have been the end of the story. We would guess that new heuristics evolve randomly, and in the rare cases that they are improvements, they get perpetuated by the progeny which have them. Thanks to Eurisko, we see that since the heuristics are represented just like any other DNA, they can work on themselves as well: they can suggest plausible (and/or warn of classes of implausible) changes to make in both (i) the DNA which synthesizes proteins, and (ii) the DNA which serves as heuristics.

### Idea #3: Heuristics drive -- and are preserved by -- embryogenesis

Now we come to perhaps an even more radical speculation than the previous two. Why aren't the heuristics lost rather quickly? After all, in a few generations, some small error is bound to creep in, and would probably negate the heuristic. Yet the individual wouldn't be any less fit, only the rate of evolution of the progeny would suffer, hence he would pass this defect along. By now, e.g., we might expect that most of the traces of how we evolved would have been obliterated from our DNA, even if they had been originally stored there somehow. One answer would be if the heuristics form (part of) the *developmental* program of the individual; if an important one is lost, then the embryo will not develop viably. This accounts for the old saw about Ontogeny recapitulating Phylogeny.

H. A. Simon said ten years ago that DNA was a recipe for producing an organism, not a blueprint, that human embryogenesis was the following of a program, not a diagram of a finished product. I'm adding that that program is a production system, that it's built out of heuristic rules, like "If an organism's body shape is X, then a tail should be added for stability". Another rule firing later triggers the elimination of the tail, when it's no longer needed.

In general, the rules will be ordered by the time they evolved, earliest ones first. Sometimes, as we all know who work with production systems, a later rule will fire a bit early, and may change the world in such a way that some of the intermediate rules will never be relevant; i.e., several intermediate steps may get skipped from

time to time.

Note also that the rules being fired are ones which have accumulated throughout history, rules for producing a *baby* of each successive species. Thus the changes one sees during embryogenesis should resemble the sequence of neonatal fish, neonatal lemurs, etc. in our ancestry, rather than resembling *adults* from those species.

A final point worth mentioning is that modifications to very old, fundamental heuristics are much more likely to be detrimental than modifications of recent ones. Thus it is usually the tail end of the program for development which is modified, the rules which fire last during embryogenesis which get changed and added to.

This is a symbiotic relationship: the heuristics enable embryogenesis to take place without some horrendously complicated central control, and in return they become indispensable. Their other function, besides development, is to guide mutation in the future: the additions to the development program will not be random, but will be heavily skewed by what is already present in that program, toward mutations which are plausible ones to try next -- where plausibility is judged by knowledge accumulated across millions of generations of experience.

#### Biological Phenomena Accounted For

The central hypothesis of this paper is that heuristics may *somehow* already be guiding evolution of higher organisms. Specific mechanisms for effecting this process have intentionally been omitted; a few vague possibilities have been hinted at. Nevertheless, several biological phenomena can be accounted for using this hypothesis. They are briefly listed here. Certainly one can hypothesize some alternate explanations of every one of them; definitive experiments must be designed and carried out to test the theory.

1. The rapid evolution of very complex organisms, organs, behavior patterns, etc. For example, some computations show that the evolution of man in general and his brain in particular was much more rapid than one could expect from random applications of the known mechanisms of molecular evolution. This is perhaps the most important kind of evidence, for it argues loudly for the need for heuristic exploration instead of random trial and error; unfortunately, it is the most controversial type of evidence.

2. The rate of evolution is not slower for complex organisms than for simpler ones. Not only is the absolute amount of time it took to evolve, say, the human eye surprisingly brief, but the rate at which complex creatures evolve seems to be, if anything, *higher* than the rate at which simple ones do. Random generation processes are usually characterized by local maxima, by slowing down of the rate of improvement as the complexity of the product increases. By contrast, heuristic search procedures speed up as more and more heuristics are added. Examples of so-called orthogenesis could be accounted for.

3. The nonuniformities in the rate of evolution. Consistency, constancy, regularity are attributes of stochastic processes. Uniformity is demanded by unguided randomness, not by intelligent heuristic search. For example, some proteins evolve at rates ten times as slow as others, yet the rate of evolution is almost constant for proteins within certain classes. As [Wilson *et al.* 77] say: "It has been hard to understand why the rate is steady within a given class. As explanations involving natural selection did not seem satisfactory, some workers proposed a non-darwinian explanation... of the evolutionary clock..." Another type of nonuniformity discussed by [Patterson 78] is that "the adult size of members of species in many groups of animals does not vary gradually, but in jumps, the ratio between the size of one species and another being 1:2, or 1:4, or 1:8. In primates, for example, the ratios are 1:8:64:512, rising in eightfold steps." Heuristic search programs generally do *not* exhibit smooth, gradual progress, but rather more the nonuniform kinds of behaviors cited above.

4. The biological function of much of the unexpressed DNA in higher organisms: Much of this is used to store the records of the species' genetic evolution; some may be used to store condensations or abstractions of that history, e.g. in the form of very rarely expressed sequences which produce enzymes that selectively mutate the genome. Of course, there is so much unexpressed DNA that there may be several other independent mechanisms which generate and preserve such sequences.

5. The fraction of non-coding DNA increases phylogenetically. We expect that the percentage of DNA which codes for heuristics rather than proteins would increase with the complexity and sophistication of the organism. Man should have more heuristics than chickens, which should have more than *E. coli*. This isn't because we're "better", just because our DNA program is longer and more involved; if our ability to adapt is to be anywhere near as good as bacteria's, we must compensate for our unwieldy program size and generation time by employing powerful judgmental rules, heuristics which put each generation to maximum use.

6. The phenomenon that relearning a beneficial mutation is much quicker than initial learning, and the intermediate state of the de-learned DNA is slightly larger than the original length. Our theory would predict that the initial act of the learning causes a new heuristic to form. Even after the mutation is forced to be un-learned, the heuristic which summarizes that experience remains. Thus, the genome is slightly longer, the increase is not merely a duplicate of the old gene though it may be closely related to it, and the relearning rate is elevated. The evidence may also be adequately explained by positing a simple duplication of genes [Schimke 80].

7. The C-value problem (some close species differ by a factor of 20 in their amounts of DNA.) This phenomenon has already been evinced by Eurisko, a program designed to evolve new heuristics. What happens is that one of the new heuristics is bad, and it generates large quantities of new genetic material before it is recognized as bad (by other heuristics) and turned off. In Eurisko, one such heuristic was "It's worth composing every pair of operations now known, to form new operations, some of which might be very powerful". This initiated an exponential explosion in the number of operations defined in each successive generation. In nature, this would mean that the length of the genome might increase very rapidly over a small number

of generations, with no apparent benefit to the individuals or the species. When the bad heuristic is deactivated, the increase halts, but it may not be easy to track down all the useless by-products produced by that heuristic. Slowly, over much much longer time scales, the extraneous material may be excised in the usual garbage-collection manner, through accidental deletions which turn out to be just as viable. To summarize: a defective heuristic can quickly (over a few generations) cause a massive amount of extraneous genetic material to be synthesized.

8. The large morphological advances of some species (like Man) compared with others (like chimps and even more dramatically frogs), even though at the DNA sequence level they both advanced an equal number of base mutations. As Wilson, Carlson & White note, the speed at which an organism morphologically evolves seems totally unrelated to the rate at which his individual proteins evolve: "In spite of having evolved at an unusually high organismal rate, the human lineage does not appear to have undergone accelerated sequence evolution... This result raises doubts about the relevance of sequence evolution to the evolution of organisms". Our theory accounts for this by simply noting that heuristic search is powerful, and its efficacy is directly related to the number and quality of the heuristics available. Programs with more heuristics can get more done in  $N$  cpu cycles (in a given fixed amount of computer time). The rate of evolution should depend more upon the number and quality of heuristics than upon the raw number of changes in the DNA molecule which occur. That is, a huge program can be improved more by adding a few good heuristics than by allotting a few more cpu cycles.

9. The molecular basis for ontogeny recapitulating phylogeny. Insect larvae resemble adult forms of lower articulate animals more than they resemble their own parents; embryonic jellyfish look more like polyps than like adult jellyfish; as they develop, human embryos resemble microorganisms, fish, reptiles, and finally earlier mammals [Gould 77]. Our explanation is that during embryogenesis, the fetus develops not via an algorithm (an explicit, fixed procedure), but via an extremely efficient set of heuristics for guidance, heuristics which *implicitly* encode the blueprint for the final neonate. One of them might say "If you see the organism in state  $x$ , then gills are a good improvement"; another might fire much later, after several other developments have been made: "If the organism is in state  $y$ , then gills are no longer needed". We are therefore postulating that the DNA contains not a blueprint for the finished product, but rather a description (compiled into heuristics) of the changes that were made over the eons in the DNA, changes which led to the evolution of our species. We are saying that ontogeny is really recapitulating phylogony in each individual embryo. Hence evolution and development are really the same process (being guided by heuristic rules) operating over very different time scales. As the organism develops, the heuristics get relatively weaker and weaker, the rate of morphological change declines to a point where it is called something else (development into adulthood), then to a point where it is not even noticed (adulthood), and finally perhaps is interpreted as senescence. Note we predict that an individual's DNA will change slowly but continuously over its lifetime, and that across species such changes should increase phlogenetically.

10. The stages one passes through in ontogeny are more like the neonatal states of ancestral species than like the adult states of those ancestors. Note that this is a

second-order effect related to (9) above. The heuristic rules at any time are collectively a program for producing an infant of a given species. The earlier rules which talk about what is appropriate when an adult state of species X is being attained never fire. The human fetus cannot much more resemble an adult lemur than can a fetal lemur resemble an adult lemur.

11. So-called parallel evolution. Before speciation, a body of more or less general heuristics has evolved. After the species divide, they may differ physiologically yet share the same heuristics. Thus their future evolution may seem surprisingly parallel. Parallel evolution is no doubt due to several species being forced to cope with the same gross environmental change; having some common heuristics increases the likelihood of their finding the same solution.

12. The ABC result (mutation rate per gram of DNA is not constant, but rather is proportional to the lengths of the DNA molecules making up the sample) [Abrahamson *et al.* 73]. The explanation here is simply that mutations are mediated by the heuristics, whose relative number increases in proportion to DNA length (roughly). One random change in a part of the DNA which is a heuristic can be expected to have a more dramatic influence than a random mutation somewhere in a coding region.

### Experiments to Test the Theory

A simple prediction is that interfering with regions corresponding to heuristics will affect the viability of mutant offspring. This may be one of the first experiments to perform, due to its general simplicity.

A more convincing experiment would be any one of the following form: Cause an organism to learn (adapt to) X, then to Y; Cause the same kind of organism to learn Y and then X. If the second learning is faster than the first in both cases, the organism somehow has learned a little bit about "learning to learn" -- i.e., it has gained or improved a heuristic. Some kind of "memory" is implied, hence it should be easier to cause a species to de-evolve than to evolve further.

As another type of experiment, raise mice in a very cold (hot) environment for several generations, allowing natural selection to take its course. Both my theory and Darwin's would predict that gradually the mice will be born better and better adapted to that temperature. Now, for the next several generations, turn natural selection off: i.e., keep *all* the mice alive. My theory would predict a kind of hysteresis effect: the mice will continue to be born with better and better cold adaption; Darwin would disagree. In other words, the mutations produced will be skewed toward those which work together to enable life in an extremely frigid temperature range. Biochemical changes in the environment of the cell trigger heuristics which take appropriate action, which trigger collections of coordinated mutations.

Unfortunately, this type of experiment is unpleasantly close to the "learned adaptation" brand, but our predictions are different: If the mice are cold and shiver a lot, the "learned adaptationists" would predict offspring which shivered more, whereas we predict offspring better suited to living in cold climates, hence shivering less.

This is a negative, rather than positive, feedback situation, a homeostatic counteracting to any environmental pressure that can be sensed at the molecular level (e.g., a decrease in overall food supply, quality of air, amount of available calcium, etc.) Even if there is no channel directly linking the external environment to the cellular environment, it is possible for the DNA to indirectly build up a model of what that environment must be like: When a mutation is made, say to aid in cold adaptation, an extra assertion is added to the DNA at the same time, namely that the climate is growing colder. If that offspring survives, then (by natural selection) it is likely that his mutation was useful and hence that his assertion about the climate was correct; see Appendix 1. This is why in the experiment above it was necessary to raise the mice in a cold climate for several generations under strong natural selection, before letting all the mice survive during the subsequent generations.

To test the hypothesis that individual development and evolution are linked, one might perform the following sort of experiment. One kind of H1 histone gene is active in the early life of an embryo; later, another kind is expressed. The experiment is to see whether the former is due to an H1 without an intervening sequence, and the latter is due to an H1 with an intervening sequence. The two types of H1s are both present in our DNA, and the latter evolved later. This would then confirm that a mechanism which evolved later is used later in the development of each embryo.

To test the hypothesis that intervening sequences are used as tags for "recently mutated successfully" messages, one could do the following experiments: (i) Test whether genes which are known to be highly conserved in evolution (e.g., for the Krebs cycle enzymes) have relatively few intervening sequences; (ii) Test whether the mean density of intervening sequences increases phylogenetically; (iii) Test whether genes known to be recently altered have a higher incidence of intervening sequences; (iv) Artificially introduce new intervening sequences into a gene and see if its mutation rate rises.

If indeed there is a universal scheme for encoding heuristics, then they may be usable across species boundaries. Even partially cracking the heuristics' "code" (which may involve positional referents and straight history, as well as domain-independent heuristics), one could try to transfer some of the heuristics from an advanced organism into a primitive one and observe their effects on the rate and direction of mutation. Nature may of course be doing this already: viruses keeping species informed of "big discoveries" such as endoskeletons across species boundaries. The biggest improvements might come about by transferring the meta-heuristics (those heuristics which deal with other heuristics, rather than with structural DNA).

The foremost problem, of course, is cracking the "heuristic code". What is the mechanism of the heuristics' functioning? Faith in unity and simplicity can both

guide our investigations and buoy our spirits with the hope that the answer is not a convoluted one. Perhaps one can look at the changes when a heuristic is transferred to various organisms, and induce what it says. How close are the analogues between programming and genetics? If the heuristics truly are IF/THEN type rules, what is the interpreter? Is the "IF" part partially or totally specified by position? Is the "THEN" part partially or totally a *history* of what the last (last few? all past?) modifications were? Are there different types of heuristics? Do some types correspond to data structures, some to plausibility rules which refer to those data structures, and others to interpreters? Are the numbers right -- i.e., is there still some missing mechanism to account for the rapid evolution the fossil record demands?

Even if it turns out that Nature has not yet hit upon the mechanism of heuristic search, there is still idea #1: design heuristics for plausible and implausible mutations, for recordkeeping, for dealing with (synthesizing, modifying, evaluating) other heuristics. They will have to be non-coding sequences, there will have to be an interpretation mechanism for obeying them at reproduction-time. Using extant techniques (e.g., plasmids), one could synthesize such sequences and insert them into DNA and study the results.

### Conclusion

Our central hypotheses are:

1. DNA has evolved into an expert program, i.e., one with heuristics for suggesting which (families of) mutations are plausible and implausible. This process began as neodarwinistic "random generate and test", but that process is not a fixed point: Evolution itself has evolved by now into a better process, one guided by past experiences, a "plausible generate and test".
2. Since the individual is viable today, his lineage is largely a series of successes; occasionally, often indirectly, knowledge of failures can be present as well. Plausible move suggesters are thus more frequent than implausible move pruners.
3. Such bioheuristics depend upon -- nay, they *embody* -- knowledge of the evolutionary history of the genome. As a species evolves viably, its body of heuristics is gradually altered (by adding new ones and modifying old ones) to capture the additional history, to compile the new hindsight.
4. Most of the library of heuristics are kept as unexpressed DNA, though it may be that expression does occur briefly, during development. This both ensures the preservation of the heuristics intact, and causes development to resemble a reenactment of the evolution of the species.
5. Since such heuristics are necessarily encoded into the DNA sequence, they can refer to (and operate on) themselves, in addition to referring to the other parts of the DNA (the structural, protein-encoding DNA). While the first heuristics originated fortuitously, the learning of new heuristics is itself by now probably under strict

heuristic control.

6. Thus the heuristics gradually grow in such a way as to better and better reflect the structure of the outer environment: the pressures, the common modes of flux, the interrelations between components. The species becomes better and better adapted to *evolving* in a complex, changing environment. The "plausibility" with which mutations are skewed increases exponentially, and this precisely counterbalances the natural deleterious effects of the combinatorial explosion, the exponential growth in the amount of time it takes to improve a program of a given length. In short, the growing "intelligence" of the mutation process is just strong enough to match the *need* for such sophistication.

These are radical hypotheses, and this paper has justified them primarily by analogy to the need for heuristics to guide automatic program synthesis. Appeals to analogy are not uncommon in molecular genetics: Enzyme induction mechanisms were debated in terms of locks & keys, templates & forms, and other real-world images. Adaptors were conceived as analogues of electrical wire or pipe adaptors. The analogy of restriction enzyme action to text editing has proven fruitful. Of course analogy is not proof nor foolproof. The purpose of the paper has been to suggest a potentially significant hypothesis for future investigation by biologists.

## Appendix 1: Examples of a set of heuristics guiding evolution

### 1.1. Guiding the simultaneous adjustment of many parameters

On the following pages is a small collection of 37 heuristics capable of guiding evolution. Thousands more would be needed for any quantitative study, but these will suffice to illustrate qualitatively how the guidance works. Each heuristic is a small, plausible, independent piece of knowledge, a generalization from past experience. Some of them are related by specialization (e.g., the first rule below is a generalization of the following three).

For simplicity, we have divided the heuristics into two classes: (i) 11 Declarative assertions ("oxygen consumption declines during sleep") and (ii) 26 Procedural IF/THEN rules which inspect the set of extant assertions occasionally "match" some of them, subsequently "fire", and result in new assertions being made.

Initially, assume that all the rules and assertions below not labelled "NEW" are present. Some rules will be relevant immediately (e.g., Rule 5), and most will require the presence of some kind of assertion before they are relevant (e.g., Rule 2). Although we have arranged them in an order related to their firing order, it is important to realize that the power of such a "rule-based" representation of knowledge lies in its *lack* of need to be ordered. One can simply add a new, general piece of knowledge -- a new rule -- to the set of existing rules, and since each rule has an IF-part the new rule should fire when (and only when) it truly is relevant to the current situation. The rules below are assumed to fire for a while, and eventually no rule in the rule set is relevant. By that time, 24 new assertions will exist which specify changes to make in the progeny (e.g., "the neck should be longer"). If such a process went on in germ line cells, and if such assertions did affect development, then the offspring would incorporate such changes.

1. IF some parameterized aspect of the world has shifted,  
THEN redesign some progeny to be better adapted to surviving if that aspect shifts even farther (with the assertion that it *is* continuing to shift in the same direction) and design a few to be *less* so (with the assertion that it's shifting back again)
2. IF the climate appears to be getting warmer,  
THEN with probability 90%: assert that progeny must be redesigned to be better adapted to heat (also: each offspring must have a new, built-in assertion that the climate is getting still warmer) and otherwise (probability 10%) assert that progeny must become better adapted to cold (also: give each offspring the assertion that it's becoming cooler again)
3. IF the climate appears to be getting colder,  
THEN with probability 90%: assert that progeny must be redesigned to be better adapted to cold (also: each offspring must have a new, built-in assertion that the climate is getting still colder) and otherwise (probability 10%) assert that progeny must become better adapted to heat (also: give each offspring the assertion that it's becoming warmer again)
4. IF the level of a nutrient, vitamin, desirable mineral, etc. is very low,  
THEN redesign some progeny to use less of it, and some to acquire more of it.

5. IF no assertion exists about whether the climate is getting warmer or colder,  
THEN randomly assert either one or the other.

A model of the external environment is used by these heuristics. How is such a model built up by the DNA molecule? We are *not* supposing that there is any direct sensing of temperature, humidity, etc. by the DNA. Rather, the heuristics guide the production of, say, two types of progeny: the first are slightly more cold-adapted, and the second more heat-adapted. The first has an assertion that the climate is getting colder, the second that the climate is getting warmer. Initially, they are produced in equal numbers (see Rule 5, above). If one group dominates, then its assertion about the climate is probably the correct one. After a few generations, if the deme is indeed entering a glacial age, the offspring will become skewed (in almost every single litter) toward more and more cold-adaptedness; each of these offspring will in turn add an extra "very" to the genetic hypothesis that it is growing very, very,... very cold out.

Let's examine how heuristics can work together to coordinate plausible mutations. Suppose that no assertion exists that says whether the climate is getting colder or warmer. In such a situation, the IF-part of Rule 5 would be true: we say that Rule 5 "triggers" and is now ready to "fire". We fire it by obeying the THEN-part of the rule: a new assertion is made and added to the set of assertions already in existence. There's a 50-50 chance for either of two assertions; suppose the following one is the one actually chosen:

6. (NEW) The climate is getting colder.

Once this assertion is made, the IF-part of Rule 3 is satisfied, so Rule 3 triggers. When it fires, it adds a new assertion to the data base. This time, there is a 90/10% skewing in favor of the following assertion:

7. (NEW) Progeny must be redesigned to be better adapted to cold.

ALSO: Each offspring will have Assertion 6 replaced by: "The climate is getting MUCH colder"

Several more of the existing IF/THEN rules may now trigger, bits of judgmental knowledge gleaned over time from vast experience with making adaptations to cold and heat. Rule 8 below is one such:

8. IF progeny must be designed to be better cold adapted,  
THEN assert that progeny must have better mechanisms to conserve heat.

When Rule 8 fires, it causes assertion 9 to be made. That (along with assertion 10, which we assume already exists) causes rule 11 to fire and synthesize assertion 12.

9. (NEW) Progeny must have better mechanisms to conserve heat.

10. Evaporation dissipates heat.

11. IF some quantity Q must be conserved, and it is being squandered by X,  
THEN reduce X in order to waste less Q (i.e., assert that X must be reduced).

12. (NEW) Evaporation must be reduced.

At this point, a rule very similar to #11 can fire, one which finds a way in which evaporation can be reduced:

13. IF some mechanism must be diminished, and it is facilitated by X,  
THEN reduce X.

14. Evaporation is facilitated by morphological structures with large surface areas.

15. (NEW) Morphological structures with large surface areas must be reduced.

16. Ears are a morphological structure with large surface area.

17. (NEW) The size of ears must be reduced.

This is certainly a useful assertion, a useful suggestion for one tiny change to make if the animal is to be better adapted to a cold environment. We assume that lower-level mechanisms can actually carry out such biophoric operations as reducing ear size, and will leave this branch of the process at this stage. Additional morphological modifications may be suggested which reduce evaporation (other responses to assertion 12), such as thickening body hair or fur, but we will not pursue them here. While rule 11 said to conserve heat by not *wasting* as much as the current organism does, there is a symmetric response to assertion 9, namely to *heighten* any existing conservation measures:

18. IF some quantity Q must be conserved, and it is being conserved already by X,  
THEN increase X in order to further preserve Q.

19. Sleep (and dormancy in general) conserves heat.

20. (NEW) Sleep (and dormancy in general) must be increased.

Again, we will not delve into mechanisms whereby the offspring will sleep and rest more than their parents did, but rather assume some low-level means to achieve these goals once they are articulated. Assertions 20 and 21 might now trigger Rule 22, resulting in Assertion 23. Eventually, using assertions 24 and 25, new assertions such as 26-28 might be made.

21. The animal is very vulnerable during sleep.

22. IF the animal is vulnerable during X, and X must be increased,  
THEN some additional protection should be provided or sought during X.

23. (NEW) Additional protection should be present during sleep.

24. Additional protection can be provided by increasing body armor.

25. Additional protection can be provided by seeking external shelters.

26. (NEW) During sleep, the progeny should seek safer shelter.

27. (NEW) The progeny should build a safer warren to dwell in.

28. (NEW) The progeny should have stiffer fur and tougher skin.

Let us return to assertion 9 and rule 18. Other ways to conserve heat may be known and therefore amplified:

29. A thick layer of fat under the skin diminishes heat transfer between organism and environment, hence conserves heat.

30. (NEW) Add (or increase) subcutaneous layer of fat.

Let's now jump all the way back to assertion 6. It triggers several rules, not just number 3. For instance:

31. IF the environment is growing colder,  
THEN the glucose level of the organism may drop.

32. (NEW) The glucose level may drop.

33. Glucose is a crucially needed substance.

Assertions 32 and 33 can trigger rule 4, which can either assert that the progeny must be redesigned to live on less glucose, or that they be better suited to acquiring it. Let's suppose that the latter is asserted. There are now several rules which are relevant to increasing glucose level.

34. (NEW) Progeny must be redesigned to increase glucose level somehow.

35. IF the total body size diminishes,  
THEN the levels of many substances may increase.

36. IF the intake of substances increases,  
THEN the levels of those substances may increase.

37. (NEW) The total body size of the progeny should be smaller.

38. (NEW) The glucose intake of the progeny should be increased.

We'll assume that primitive operations carry out assertion 37, although perhaps many more rules are fired, rules which adjust parameters to suit a smaller overall body size. Several rules suggest ways of increasing glucose intake:

39. IF locomotive muscles are increased,  
THEN glucose level may increase.

40. IF teeth and claws are increased in size and sharpness, and jaw muscles are increased,  
THEN glucose level may increase.

41. IF brain size is increased,  
THEN glucose level may increase.

42. IF neck size is increased,  
THEN glucose level may increase.

Of course the justifications for the rules is probably beyond the store of DNA's knowledge -- rule 39 is based on catching prey more efficiently, rule 40 on tearing and chewing more efficiently, rule 41 on outsmarting prey, rule 42 on reaching more vegetable matter, etc. Some or all of these rules will fire, causing several

morphological redesigns of the progeny. Suppose that rules 39, 40, and 41 fire. Then the following assertions would be produced:

- 43. (NEW) Increase size of locomotive muscles.
- 44. (NEW) Increase size and sharpness of claws and teeth, increase jaw muscles.
- 45. (NEW) Increase brain size.

These assertions will, of course, engender several more rule firings to compensate; for instance:

- 46. Teeth and other bones require calcium.
- 46. IF a morphological structure is being increased,  
THEN the level of substances it is based on should be increased somehow.
- 47. (NEW) Calcium level must be increased somehow (to support larger teeth).

Now a whole new subproblem is being worked on: how to balance the level of calcium in the body. If teeth are to be more numerous and larger, then there is a danger of the calcium level getting too low. One way to compensate is:

- 48. IF a morphological structure is reduced in size,  
THEN the level of substances it is based on will increase.
- 49. (NEW) Reduce bone size generally (to balance increasing teeth size).

Both 43 and 44 assert that muscles will be increased in size. They trigger 50, which asserts 51. 51 and 52 together trigger 53, which asserts 54:

- 50. IF muscle size increases,  
THEN lactic acid concentrations may peak at higher levels.
- 51. (NEW) Peak concentrations of lactic acid may increase.
- 52. Lactic acid is an undesirable by-product of useful reactions.
- 53. IF peak concentrations of undesirable substances will increase,  
THEN redesign progeny to metabolize them more rapidly.
- 54. (NEW) Progeny must metabolize lactic acid more effectively.

But suppose there is no recorded mechanism for metabolizing lactic acid. What can the system do? It can rely on very general -- but weak -- knowledge about metabolizing any substances. This might have scores of possible suggestions (suggestions for new enzymes, old ones to vary, old ones to increase, etc.), and one or more might be tried out. Not only would the progeny then get these new mechanisms, but they would also receive new assertions that those mechanisms were effective for metabolizing lactic acid. The progeny who survive presumably have more accurate assertions than those who perish.

We have only explored a tiny part of the network of changes which would be triggered by the innocuous assertion that the environment is getting colder. Already,

we have redesigned the species to be smaller, lighter-boned, have bigger and sharper teeth, larger jaw muscles, larger leg muscles, increased brain size, sleep more, seek safer burrows, have thicker and stiffer fur, have an added layer of subcutaneous fat, have smaller ears, have one of a set of possible mechanisms to metabolize lactic acid more effectively, etc. The changes along any one parameter might be tiny, but (i) they would all complement each other, some even compensating for imbalances introduced by others, and (ii) the total of all these changes might be a significant change in the ability of the organism to withstand colder environments.

All these changes work together; they could all be tried simultaneously in a single offspring. If the rules were sophisticated enough, the modifications might not be "hard-wired" in, but rather *canalized* to let the actual environment tune the degree to which they took effect. The offspring differs in perhaps thousands of small ways -- a constellation of related changes that mesh with each other, that accomplish some goals. These are not the teleological goals of creationists -- goals which were somehow placed in DNA long ago; rather, they are short-term goals proposed by the DNA itself, on the basis of its knowledge about evolution, the structure of the environment, and possibly some feedback on the changes occurring in that environment. As we showed in the first paragraph of this section, such feedback ("growing very cold out!") can be inferred indirectly by the DNA without the need to postulate any direct external sensing abilities. Many of the goals are proposed simply to counteract side-effects introduced by earlier proposed mutations.

A sophisticated model of the physical environment may have been accreted over many generations, many individuals, and many variables. By now a large knowledge base may exist about ecology, geology, glaciation, seasons, gravity, predation, symbiosis, causality, conservation, behavior, evolution, and knowledge itself. In a small number of generations, man has managed to invalidate many of these bits of knowledge, this model of the world. If the heuristics can trace this breakdown to the increasing size of our brains, they might take quick corrective action, preserving homeostasis and the validity of their knowledge base by drastically decreasing human brain size over just a few generations. While this is of course a fanciful tongue-in-cheek extreme case, it -- and the longer example above -- demonstrates the power, the coordination, that a body of heuristics could evince if it were guiding the process of evolution.

## 1.2. Guiding the discovery of new features and mechanisms

Earlier in this paper, and at length elsewhere [Lenat 79], we discuss the genesis of new concepts and the discovery of conjectures connecting them, under the guidance of a body of heuristics. General rules say:

55. IF "2" occurs in some mechanism, structure, or rule,  
THEN replace it by "3" or by "1"

56. IF the products of an operation or mechanism or rule are of the the same category as the objects it takes as input,  
THEN define and investigate the set of inputs which are transformed to themselves.

57. IF an operation or mechanism or rule takes a pair of substances as inputs,  
THEN see what happens when it operates on two identical substances.

Those apply to mathematics, geology, and politics as well as to biology. More specific heuristic rules can of course be stated:

58. IF predators seem to be getting rarer,  
THEN protect sensory apparatus less and also add an assertion that they are getting yet rarer.

59. IF two copies of a sensor are separated,  
THEN perception is slightly improved.

Coupled with some assertions, such as those below, the rules can guide the formation of plausible new structures and mechanisms, some of which may actually be advantageous to the individual.

60. Predators are becoming rarer.

61. Separation can exist over space or (rarely) over time.

Assertion 60 may trigger rule 58, which may release some constraints on how heavily the eyes, nose, etc. must be protected. This may eventually ripple out to shallower eye sockets. Because of 61, rule 59 might trigger and separate the eyes a bit (which is now explicitly allowed as safe, due to rule 58). An extreme of this might be moving nostrils down to the sides of the neck. Any hypothesis involving moving nostrils or eyes very far would soon run up against still-active constraints about the high cost of long optic nerves and nasal passages. However, assertion 61 does permit several other potential improvements to be made: the sensory separation can be effected by having separate individuals communicate across reasonably large distances (certainly large compared to the diameter of the skull). By comparing sensory data across distances, the existing sensory mechanisms can be made to yield better results. A second use of 61 is to have a single individual store a sensory impression, run to a different spot (in space or time), and compare the old image with the new one. In the case of temporal delays, this gives rise to motion detectors (similar to blink boxes used by astronomers to find planets and comets). In the case of spatial delays, this would demand a photographic memory, but would yield greatly improved parallax information. While this has all been carried out at a superficial level, the intent is to convince the reader of the utility of using heuristics for guiding discovery.

We could have strung together any few sentences, out of a vocabulary that included words like Duplicate, Move, Perturb, etc., but the density of good new ideas would have been exponentially less than the way we got them above, using heuristics to suggest plausible combinations and alterations.

## Appendix 2: Relevant Existing Knowledge

It may be instructive to record the "context" of this hypothesis; the knowledge (and misinformation) that led to its creation. Below, asterisks (\*) indicate "facts" that I believed before the idea was formed, but which (due to subsequent reading/discussion) I now feel are wrong/unknown. Pluses (+) indicate facts I have learned since the idea was formed.

### 1. Mendelism is accepted absolutely.

That is, we are completely determined by our genetic makeup; in particular, by our genetic materials *at birth* (\*). Changing said genetic materials (in our germ cells) will alter the genetic makeup -- and hence the "blueprints", the design -- of our offspring

### 2. Evolution in the strict Darwinian sense (i.e., via a series of random mutations) is incapable of accounting for the presence of Man on earth today.

Certainly, we do not dispute that natural selection operates; rather, we are skeptical of the quantitative plausibility of the origin of the species in so short a time. The order of magnitude of such a "pure hillclimbing" toward Man is estimated to be as large as  $10^{10000000}$  years! The mutation rate per gene per generation is around  $10^{-7}$  (+); almost all random mutations are deleterious, or at best neutral; there is a good chance that even an advantageous new allele will be lost (die out before fixation occurs) due to fluctuations in its frequency in the population as a whole. Bear in mind that natural selection does not tolerate much curvilinear development. I.e., a very complex system (like the double-negative repression-repression system for B-galactosidase) would have had to evolve in steps *each* of which was a non-negative improvement over the last one. Below are a few of the many additional doubts and riddles presented in articles in Duncan & Weston-Smith's *Encyclopedia of Ignorance*:

"How is it possible for future evolutionary flexibility to be preserved when the exigencies of survival apply strong immediate selection pressure? ... Is it simply chance that some species preserve evolutionary flexibility while others do not?... All of these questions suggest that natural selection is a subtle process and that a significant part of the genetic information may not be subject to short-term selection. How could such information be stored, and over what period of time is it effectively selected? There are aspects of the fossil record which suggest parallel evolution of species lines that have been long separate. Such convergent or parallel evolution does not have an easy explanation and also suggests long-term storage of genetic information. On a molecular level there are also suggestions of freedom from selection pressure, or longer periods of integration. For example, mammals contain enough DNA per cell to code for an excessive number of potential genes (though most of this DNA is surely something other than structural genes...) There is obviously a lot of DNA in the genome of higher organisms that we can not account for. This has been termed the C-value paradox. To add to the mystery, most of the single copy DNA in primates changes so rapidly in evolution that it is probably under little or no selection pressure. We do not know what unexpressed potentialities exist in all of this 'extra' DNA... [There have been] 1500-15000 significant changes incorporated, after selection, into human DNA in 15 million years. Are these few base substitutions incorporated in the DNA enough to be the source of variation for the last 15 million years of evolution? It seems unlikely unless they had just the right effect. We can think in terms of changes in the gene regulatory system that would affect the form or function of an organ. But how many base substitutions can have such effects? Amino acid substitutions in typical proteins -- no way. Even billions [of small biochemical changes] might not be enough."

--- The Sources of Variation in Evolution (by Roy J. Britten)

J.C. Lacey, A.J. Weber, and K.M. Pruitt say, in *The Edge of Evolution*, "The primary DNA information, although inside the cell, now represents part of the environment for selecting the super [meta-level] information." Cf. their citation of E. Zuckerkandl and L. Pauling's "Molecules as documents of evolutionary history", *J. Theor. Biol.*, 8, 357-66, 1965. Tomlin says,

"Evolution was an hypothesis which hardened into dogma before it had been thoroughly analysed... Even sophisticated Darwinians such as Konrad Lorenz assume without question that the origin and formation of species can be explained as a succession of fortuitous

variations and mutations passing through the mesh of selection. The oddity of this theory is partially concealed by its mode of presentation. [Our tools -- both external ones like rotary saws and internal ones like enzymes -- must have developed] thematically; they cannot have come into being by a series of mutations or mechanical faults of copying".

---- Fallacies of Evolutionary Theory (E.W.F. Tomlin)

"Suppose that at a time 200 million years ago, during the age of reptiles, some event had taken place which doubled the rate of gene mutation in all existing organisms... Would the present state have been reached in only 100 million years? Or would the rate of evolution have stayed much the same?... The short answer is that we do not know... A theory of evolution which cannot predict the effect of doubling one of the major parameters of the process leaves something to be desired."

---- The Limitations of Evolutionary Theory (John Maynard Smith)

Additionally, several quotes ( + ) from Dawson's *Modern Ideas of Evolution* (1890) remain potent: "Viewed rightly, the direct equilibration of the parts of animals and plants is so perfect and stable, and such great evils arise from the slightest disturbance of it by the selective agency of man, that it becomes one of the strongest arguments against [Origin of the Species]... When the stability of an organism is artificially altered by man in his attempts to establish new breeds, infertility and death of these varieties or breeds results; and if this happens under the fortuitous selection supposed to occur in nature, any considerable variation would result either in speedy return to the original type or in speedy extinction. In other words, so beautifully balanced is the organism, that an excess or deficiency in any of its parts, when artificially or accidentally introduced, soon proves fatal to its existence as a species: so that, unless nature is a vastly more skilful breeder and fancier than man, the production of new species by natural selection is an impossibility." (pp 41-42) "It is to be observed here that every species of animal or plant, of however low grade, consists of many co-ordinated parts in a condition of the nicest equilibrium. ANy change occurring which produces unequal or disproportionate development, as the experience of breeders of abnormal varieties of animals and plants abundantly proves, imperils the continued existence of the species. CHanges must, therefore, in order to be profitable, affect the part sof the organism simultaneously and symmetrically, and must be correleated with all the agencies in heaven and earth that act upon the complex organism and its several parts. The chances of this may well be compared to the casting of aces [on dice] a hundred times in succession, and are so infinitely small as to be incredible under any other supposition than that of intelligent design." (pp. 105-6) I would add only that the so-called intelligence need not be *external*; adequate design knowledge may by now exist within the genome. See Appendix 1.

"A further difficulty arises from our failure to find satisfactory examples of the almost infinite alleged connecting links which must have occurred in a gradual development. This, it may be said, proceeds from the imperfection of the record: but when we find abundance of examples of the young and old of many fossil species, and can trace them through their ordinary embryonic development, why should we not find examples of the links which bound the species together? An additional difficulty is caused by the fact that in most types we find a great number of kinds in their earlier geological history, and that they dwindle rather than increase as they onward... Objections of this kind appear to be fatal to the Darwinian idea of slow modifications, proceeding throughout geological time, and to throw us back on a doctrine of sudden appearance of new forms..." (p. 33) This is reminiscent of the competing theories of *geologic* evolution via rare cataclysms *versus* via gradual change; eventually, that conflict was resolved by each side realizing that much of of what the other was saying was necessarily correct. Dawson gives several examples of the sudden emergence of new species: "Palaeontology has... adduced the advent of the Cambrian trilobytes, of the Silurian cephalopods, of the Devonian fishes, of the Carboniferous batrachians, land snails, and myriapods, of the marsupial mammals of the Mesozoic and the placental mammals of the Eocene, and of the Paleozoic and modern floras, as illustrations of the sudden swarming in of forms of life over the world, in a manner indicating flows and ebbs of the creative action, inconsistent with Darwinian uniformity, and perhaps unfavourable to any form of evolution ordinarily held." (p. 50) "Many new forms appear to be introduced at one time and apparently suddenly, so that such groups as the ferns and club-mosses and mares' tails among plants, and at a later date the more perfect fruit-bearing trees, the coral animals, the lampshells, the crinoids, the amphibians, the reptiles, the higher mammals enter on the scene

abruptly and in large numbers. Thus the impression left on our minds by this grand procession of living beings in geological time is not that of a mere continuous flow..." (p. 93) "the five fingers and toes of man appear to descend to us unchanged from the amphibians or batrachians of the Carboniferous period. In this ancient age of the earth's geological history feet with five toes appear in numerous species of reptilians of various grades. They are preceded by no other vertebrates than fishes, and these have numerous fin-rays instead of toes. There are no properly transitional forms, either fossil or recent, the nearest pectoral fins to fore limbs being those of certain Devonian and Carboniferous fishes; but they fail to show the origins of fingers. How were the five-fingered limbs acquired in this abrupt way? Why were they five rather than any other number? Why, when once introduced, have they continued unchanged up to the present day?" (pp. 141-2) As William R. Shea comments, "Dawson also made much of the existence of perfect organs such as the eye among the marine fauna of the early paleozoic seas. He believed that the two types of eye encountered -- one composed of many lenses, as in the modern fly, the other a single lense, as in most mammals -- were so different that neither could have originated from the other. Since the eye is obviously useless except in its final, complete form, how could natural selection have functioned in those initial stages of its evolution, when the variations had no possible survival value? No single variation, indeed no single part being of any use without every other, it seemed irrelevant to appeal to the survival of the fittest." (p. xxi) Darwin had earlier worried about this, and eventually conquered his doubts. "Since there were gradations of eyes among different organisms, even though there was no evidence for gradation among the lineal descendants of any one species, Darwin saw 'no very great difficulty... in supposing natural selection to have converted the simplest optic nerve into the most complex and powerful instrument'. When evidence failed to materialize, he enjoined his readers not to lose faith in a theory that had served them so well in other instances." (p. xxi)

"[Darwinism] seems to enthrone chance or accident or necessity as Lord and Creator, and to reduce the universe to a mere drift, in which we are embarked as in a ship without captain, crew, rudder, or compass..." (p. 27) The idea here is the metaphor to expertise in sailing: even though the final destination be unknown, the chances of success and efficiency of the voyage can be increased by having and using *tools* and *expertise in sailing*. A compass is of use even if there is no known goal (e.g., to keep one from going in circles), and knowledge of tacking and knot-tying is always indispensable. Teleology is *not* being claimed.

### 3. *Natural selection is accepted completely.*

Survival of the fittest, in a harsh environment, is the sole criterion for judging improvement (we needn't consider the past few thousand years, during which civilization has warped that standard). Natural selection is omnipresent and severe. So, e.g., curvilinear progress is rarely tolerated. That is, when a mutation produces an inferior result, it won't survive long enough to combine with a meshing inferior mutation to yield an improved combination. Of course (+), neutral mutations abound, and pockets of mutants may remain isolated and safe for generations.

### 4. *Eurisko is assumed to be viable*

The idea is that a body of heuristics can guide a program in discovering new domain concepts, conjectures, *and* new heuristics. Into this category we bundle all the following:

Complex tasks call for expert programs

To construct an expert program, we must somehow put "expertise" into programs.

Heuristic if-then rules are a reasonable language in which to state such expertise.

Generate&Test alone is much too weak to give adequate performance in complex domains.

Heuristic rules can efficiently guide huge searches (e.g., in medical diagnosis tasks).

The above applies to open-ended searches for new ideas (as in AM).

The above applies to searches for new heuristics as well as new math concepts.

Thus, a body of heuristics can improve and expand "itself".

### 5. *DNA is viewable as a program, but some subroutines serve as-yet unknown purposes.*

The percentage of such "non-coding" segments increases as one ascends the evolutionary ladder (+) from prokaryotes to yeast to chicks to humans.

6. Thus "evolution" is akin to "automatic programming".

From the latter comes the need to add additional knowledge, both about programming and about the task domain in which the target program is going to perform. If I want a computer-naive person to write an immense accounting program, it is clearly cost-effective for me to send that person away to learn something about programming and about accounting, rather than immediately sitting them down at a terminal and instructing them to keep trying. This is the theme of the paper and is discussed in detail therein.

### References

Abrahamson, Seymour, Michael A. Bender, Alan D. Conger, and Sheldon Wolff, "Uniformity of Radiation-induced Mutation Rates among Different Species", *Nature*, 245:5246, 460-2, October 26, 1973.

Barstow, David, *Knowledge Based Program Construction*, Elsevier North Holland, 1979.

Bukhari, A. I., J.A. Shapiro, and S.L. Adhya, *DNA Insertions, Elements, Plasmids, and Episomes*, Cold Spring Harbor Laboratory, 1977.

Cavalli-Sforza, L. L., and W. F. Bodner, *The Genetics of Human Populations*, W. H. Freeman and Company, San Francisco, 1971.

Dawson, Sir J. William, *Modern Ideas of Evolution*, New York, Neale Watson Academic Publications, Inc., 1977. Original work was issued in 1890. Critical introduction written in 1977 by William R. Shea.

Dobzhansky, T., F. Ayala, G. Stebbins, and J. Valentine, *Evolution*, W. H. Freeman, San Francisco, 1977.

Duncan, Ronald, and Miranda Weston-Smith, *The Encyclopedia of Ignorance: Everything you ever wanted to know about the unknown*, Pergamon Press, New York, 1977, 205-411.

Feigenbaum, E. A., *The Art of Artificial Intelligence*. Proceedings of the Fifth International Joint Conference on Artificial Intelligence, MIT, Cambridge, 1977.

Fogel, L., A. Owens, and M. Walsh, *Artificial Intelligence Through Simulated Evolution*, John Wiley & Sons, Inc., N.Y., 1966.

Friedberg, R. M., "A Learning Machine: Part I", *IBM J. of Research and Development*, 2, 1, January, 1958.

Friedberg, R. M., B. Dunham, and J. H. North, "A Learning Machine: Part II", *IBM J. of Research and Development*, 3, 3, July, 1959.

Gould, S. J., *Ontogeny and Phylogeny*, Belknap Press, Harvard University, Cambridge, 1977.

Green, C. Cordell, R. Waldinger, D. Barstow, R. Elschlager, D. Lenat, B. McCune, D. Shaw, and L. Steinberg, *Progress Report on Program-Understanding Systems*, AIM-240, STAN-CS-74-444, AI Lab, Stanford, August, 1974.

Judson, H., *The Eighth Day of Creation*, Simon & Schuster, 1979.

Lenat, D. B., "BEINGS: Knowledge as Interacting Experts", Proc. Fourth International Joint Conference on Artificial Intelligence, Tbilisi, USSR, 1975.

Lenat, D. B., "The Ubiquity of Discovery", *J. Artificial Intelligence*, Dec., 1977.

Lenat, D. B., "On Automated Scientific Theory Formation: A Case Study Using the AM Program," in (Jean Hayes, Donald Michie, and L. I. Mikulich, eds.) *Machine Intelligence 9*, New York: Halstead Press, a division of John Wiley & Sons, 1979, pp. 251-283.

Minsky, M. L., *Computation: Finite and Infinite Machines*, Englewood Cliffs, N.J., 1967.

Ohno, S., *Evolution by Gene Duplication*, Springer Verlag, New York, 1970.

Patterson, C., *Evolution*, Cornell University Press, 1978.

Polya, G., *How to Solve It*, Princeton University Press, 1945.

Samuel, A., "Some Studies of Machine Learning in the Game of Checkers II," *IBM J. Research and Development*, 11, 6, November, 1967, pp. 610-617.

Schimke, R. T., "Gene Amplification and Drug Resistance", *Scientific American*, 243, 5, November, 1980.

Simon, H. A., *The Science of the Artificial*, MIT Press, 1969.

Wilson, A. C., S. S. Carlson, and T. J. White, "Biochemical Evolution", *Am. Rev. Biochem.*, 1977, 46:573-639.

Winston, P. H., *Learning Structural Descriptions from Examples*, Project MAC TR-231, MIT AI Lab, Cambridge, Mass., September, 1970.

### Acknowledgements

I wish to thank the following for useful comments they have made upon the ideas in this paper: J. Seely-Brown, D. Brutlag, L. Cavalli-Sforza, L. Darden, R. Davis, E. Feigenbaum, P. Friedland, R. Hayes-Roth, L. Kedes, J. Lederberg, M. Stefik, and D. Wallace. This does not, however, imply their approval or acceptance of the hypothesis.