The Processes of Scientific Discovery: The Strategy of Experimentation

Deepak Kulkarni and Herbert A. Simon

Departments of Computer Science

and Psychology

Carnegie-Mellon University

Running Head: Kulkarni and Simon

This research was supported in part by the Defense Advanced Research Projects Agency under Contract F-33615-84-K-1520, and in part by the Personnel and Training Programs, Psychological Sciences Division, Office of the Naval Research, under Contract No. N00014-86-K-0768. Reproduction in whole or in part is permitted for any purpose of the United States Government. Approved for public release; distribution unlimited.

We are deeply indebted to Professor Frederic L. Holmes of Yale University, whose research on Hans Krebs' discovery of the ornithine cycle provided the basic data upon which we have drawn for this research, and who has provided valuable comments on drafts of our paper. Thanks are also due to Kenneth Shaffner, whose suggestions have greatly improved the presentation of this paper. Among others whom we would like to thank for comments on the manuscript or for information about the chemistry of our problem are John Modell, Chen Ho, Elaine Kant, David Hackney, David Steier, Yumi Iwasaki, Craig Knoblock, Natarajan Ganesh, Uday Shenoy and Raju Ramanujan.

Correspondence and requests for reprints should be sent to Herbert A. Simon, Department of Psychology, Carnegie-Mellon University, 5000 Forbes Avenue, Pittsburgh, Pa 15213.

Abstract

Hans Krebs' discovery, in 1932, of the urea cycle was a major event in biochemistry. This paper describes a program, KEKADA, which models the heuristics Hans Krebs used in this discovery. KEKADA reacts to surprises, formulates explanations and carries out experiments in the same manner as the evidence in the form of laboratory notebooks and interviews indicates Hans Krebs did. Further we answer a number of questions about the nature of the heuristics used by Krebs, in particular: how domain-specific are the heuristics? To what extent are they idiosyncratic to Krebs? To what extent do they represent general strategies of problem solving search?

The relative generality of KEKADA allows us to view the control structure of KEKADA and its domain-independent heuristics as a model of scientific experimentation that should apply over a broad domain.

This paper is part of a program of research aimed at studying the processes of scientific discovery by constructing computer programs that are capable of making discoveries and that simulate, at a grosser or finer level of approximation, the paths that have been followed by distinguished scientists on their roads to important discoveries. Predecessors to this paper include the work of Buchanan and others on Meta-DENDRAL (Buchanan & Feigenbaum, 1978),of Lenat on AM (Davis & Lenat, 1980), of Friedland (1979) on MOLGEN and of Langley, Simon, Bradshaw, and Zytkow (1987) on BACON and related programs .

Since scientific discovery involves a whole array of activities -- designing and performing experiments, inferring theories from data, modifying theories, inventing instruments, and many others -- any single inquiry will necessarily focus on some special aspects of the whole process. The research on BACON, for example, was concerned mainly with the ways in which theories could be generated from empirical data, with little or no help from theory. The question of where the data came from was left largely unanswered. The processes of designing experiments and programs of observation were not investigated.

The present paper represents a first investigation of some of the domains left unexplored by the previous research. It was made possible by the existence of a detailed historical study of a particular scientific discovery: Hans Krebs' elucidation of the chemical pathways for synthesis of urea in the liver (Holmes, 1980). That study traces in detail the sequence of experiments carried out by Krebs and Henseleit between July 1931 and April 1932, the strategies that determined the experimental program, and the gradual emergence of a theory of the urea synthesis pathway from the experimental data in combination with previous literature on the problem.

The discovery of the ornithine cycle was the *first* demonstration of the existence of a *cycle* in the metabolic biochemistry and it marked "a new stage in the development of biochemical thought" (Fruton, 1972).

Holmes' reconstruction of this discovery from published papers, laboratory notebooks, and interviews with Krebs, provides a magnificent body of data for developing and testing theories of many aspects of the scientific discovery process.

The system, KEKADA[1] , which we have built does not, of course, capture the full detail of the actual historical process; but it does represent a serious attempt to describe both the knowledge and the heuristics that Krebs used in his research. In addition to domain knowledge and special experimental techniques, domain-independent methods played a significant role in this discovery. By extracting these general discovery heuristics from the problem-specific knowledge of KEKADA, we can derive from the system a number of domain-independent methods of discovery which may be used in the future to create a more general discovery system.

Thinking-aloud protocols have been used extensively as a tool for obtaining insights into psychological processes in problem solving. They have even been used for studying some learning and discovery tasks (Anzai & Simon, 1979; Shrager & Klahr, 1986). The focus of this research was to study discoveries that occur in experimental sciences. Since the research leading to such discoveries sometimes spans months or years, it is not practical to gather continuous protocols of the process. Thus we must seek other sources for insights into the processes: for example, scientists' recollections, published papers on the discovery, and accounts from diaries and laboratory notes.

1. **Accounts by recollection.** The discovery is recounted by the discoverer from his recollections. This is a very common source of information about discoveries, much of it contained in scientists' autobiographies.

2. **Accounts from published papers.** Another easily available source of information about a discovery is the papers which the scientist has published in the course of discovery.

3. **Accounts from diaries and laboratory notes.** The course of discovery is reconstructed from notes and diaries of the discoverer. Gaps in the diaries may be filled in by retrospective recollections of the discoverer during his lifetime. Holmes' reconstruction of Krebs' discovery was based on Krebs' laboratory notebooks, supplemented by interviews.

Given the known fallibilities of human memory, accounts by recollection, though by far the most common, are also the least reliable. There are likely to be errors of both omission and inclusion, the likelihood increasing with the gap in years between the time the work was done and the time when the

[1]The system is named KEKADA for two reasons. KEKADA is a Hindi synonym for the German word Krebs. Thus we named the system after Hans Krebs, the great biochemist. Secondly, KEKADA means a crab in English. The process of scientific discovery is analogous to a crab crawling slowly to a destination.

recollections were recorded. Kekule first reported publicly his famous anecdotes about the imagery he used in discovering the benzene ring some 29 years after the event. How much probative weight can we place on such recollections?

Technical papers on the discovery are written at a time when memory of it is fresher than in the case of a scientist recollecting after 30 years. But generally the papers explain and justify a discovery and rarely describe how the scientist made it. Besides technical papers are written not on a daily basis, but after a major piece of work is completed. In the absence of better sources they are sometimes used to get clues about psychological processes. For example, Friedland (1979) used published papers and interviews as a source of information for understanding how people design experiments. On the basis of this information, in 1979 he constructed MOLGEN, a system that designs experiments in molecular biology.

In most experimental sciences it is customary for scientists to record the details of their experimental activity on a daily basis in a laboratory notebook or log. Logs may be bareboned, or they may contain reasons for carrying out an experiment, observations, and conclusions drawn from the data. Experiments would seldom be omitted. Some scientists also note in their notebooks when new ideas occur to them and how their thoughts and plans were influenced by them. Since the log entries are usually made daily, when the investigator has no knowledge of the discovery that will later emerge, the accounts are not influenced by the future results.

In relatively theoretical sciences, scientists would do much deep thinking about the domain which may not be reflected in the logs and thus the account from logs may have major gaps. On the contrary in a domain that has a relatively shallow theory, the scientist may not rule out possibilities without actually carrying out experiments and the reasoning behind an experiment would be easy to guess. In such cases an account from logs can provide a very close, if not complete, picture of the thinking that lead to the discovery.

Holmes' reconstruction, based on laboratory notebooks and retrospective interviews falls in the second category. First of all, the domain of biochemistry in the 1930s had a relatively shallow theory. In addition "Having had less than a year of systematic training in chemistry, Krebs did not possess the extensive knowledge of the properties and reactions of organic compounds necessary to reason

deeply about the metabolic steps that would be most likely, on theoretical grounds, to take place. He could only follow every plausible suggestion he came across." (Holmes, 1986) Ironically, his lack of expert knowledge of organic reactions freed Krebs from some of the biases built into the conceptual frameworks within which contemporary biochemists operated and thus conferred on him some real benefits. (Holmes, 1986) Consideration of these factors in the context of a specific domain makes it plausible that Holmes' reconstruction is a close description of how Krebs attacked the problem and thought about it. It therefore follows that it should be possible to create a good theory based on such data.

In this study, we use Holmes' reconstruction, based on laboratory notebooks and retrospective interviews, as our source of insight into the process that led to the discovery of the ornithine cycle for the synthesis of urea. Using this reconstruction, we have built a computer program. KEKADA, that placed in the situation in which Krebs began his work, simulates this discovery. In the next section, we will summarize Holmes' account. Then we will describe the heuristics employed by KEKADA for the simulation. In a third section, we will report the behavior of KEKADA when placed in the situation in which Krebs began his research, and we will compare the actual history with the simulation.

1. The Ornithine Cycle

We paraphrase here (with his kind permission) Holmes' (1980) account of the discovery of the ornithine cycle. The direct quotations are from Holmes' paper. The discovery, in 1932, of this chemical pathway was of major importance to biochemistry. The problem that Krebs attacked, to discover how urea was synthesized in living mammals from the decomposition products of proteins, had been investigated extensively for many years with very limited success. The methods used in Krebs' discovery, and the general nature of the catalytic process discovered, served as prototypes for much subsequent research and theory on metabolic phenomena.

1.1. Background of the Discovery

Early in the 19th Century, urea had been synthesized in the laboratory, and knowledge of its composition and the synthesis paths led to certain hypotheses as to how it might be synthesized *in vivo*. Feeding experiments with animals showed that adding glycine or leucine to the diet increases the secretion of urea, and led to the conclusion that these amino acids were the intermediates between protein and urea. Similar feeding experiments later showed that ammonium salts added to the diet would also increase the output of urea.

By the use of isolated perfused livers, it was then shown that ammonium salts, leucine, tyrosine, and aspartic acid increase the formation of urea, and it was concluded that the liver produces urea from amino acids and ammonia. Experimental difficulties with perfusion methods left the question of the actual mechanism undecided -- it appeared to be "impossible to prove experimentally which of the several theories of the reaction mechanism derived from test tube processes was the one that occurred physiologically" (Holmes, 1980).

Attempts to get around the limitations of the perfusion experiments by attempting to synthesize urea with tissue extracts also failed to obtain conclusive results, supporting the opinion of Loffler that "urea formation in the surviving liver is bound up with the integrity of the cell structure" (Loffler, 1920). This was the situation that prevailed, in 1931, when Krebs began his research on this topic.

1.2. Course of Krebs' Research

The account of Krebs' research can be divided conveniently into three major segments: the first from July 26, 1931 to November 15, when the effects of ornithine were first noticed; the second from November 15 until about January 14, 1932, when evidence indicated that the effect was quite specific to ornithine; the third from January 14 to April 13, when Krebs was sufficiently convinced that he had discovered the synthesis mechanism to send off a paper for publication. Thus, the critical phenomenon that led to the solution of the problem was detected after about three and a half months of work, while interpreting the new phenomenon and testing the theory required another five months.

1. **The ornithine effect.** Krebs began with the idea of using the tissue-slice method, a technique he had acquired in Otto Warburg's laboratory, to study urea synthesis. He tested the efficacy of various amino acids in producing urea, with generally negative results. When he carried

out the experiment with ornithine (one of the less common amino acids) and ammonia, unexpectedly large amounts of urea were produced. He then focused on the ornithine effect.

2. **Determination of scope.** Krebs next followed a standard strategy: if a given compound exerts a particular action, check whether derivatives of that compound have a similar action. Thus, he carried out tests on some ornithine derivatives and substances similar to ornithine. But none of these substances had effects comparable to ornithine.

3. **Discovery of reaction path.** New apparatus that he obtained at this time enabled him to determine that the nitrogen in the urea produced was comparable in quantity to the nitrogen in the ammonia consumed. He concluded that the ammonia, not the amino acids, was the source of the nitrogen. Krebs now sought to elucidate the mechanisms of the ornithine effect. It occurred to him that the (known) arginine reaction, by which arginine is converted to ornithine and urea, might be related to the ornithine effect. Concluding from the quantitative data that the ornithine could only be a catalyst, he inferred that ornithine with ammonia produces arginine, which in turn produces urea and ornithine. Later experiments indicated that citrulline was an intermediate substance between ornithine and arginine.

We must now spell out the details of Krebs' experiments and reasoning somewhat more fully, still following closely the account of Holmes.

1.2.1. The Ornithine Effect

In the laboratory of Otto Warburg, from 1926 to 1930, Krebs learned the method Warburg had developed of carrying out reactions on tissue slices instead on the organ itself. The tissue slice method is simple and fast compared with the perfusion method used previously. Krebs conceived the idea of using the tissue slice method for problems other than the study of cellular respiration, which had been the focus of Warburg's work. Since the method preserved many cells intact, metabolic processes might be observed that disappeared with tissue extracts. Warburg did not support Krebs' idea, perhaps because he thought that energy-absorbing reactions (as contrasted with oxidation reactions) would not go forward in tissue slices.

When Krebs got freedom to initiate a major research enterprise of his own, in 1931, he decided to begin experiments of the sort he had conceived. Urea synthesis was an obvious choice of a

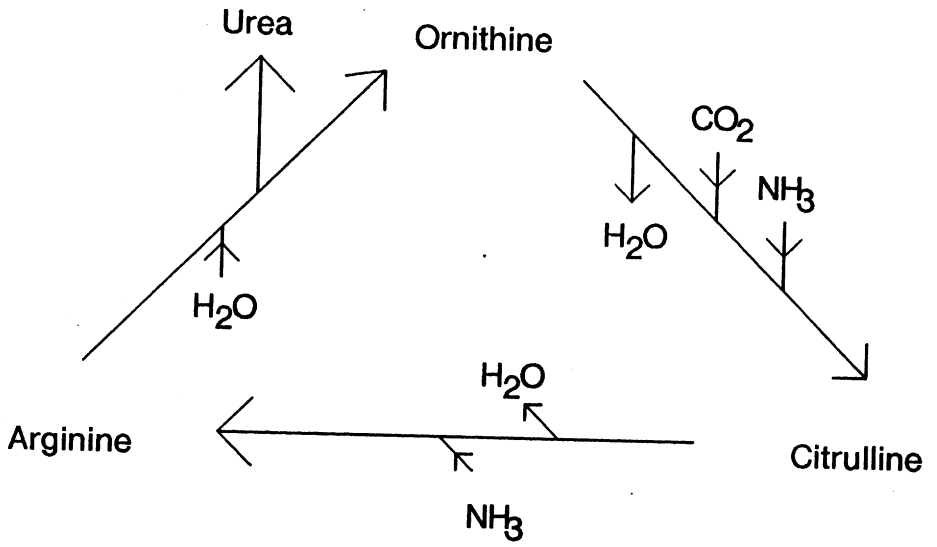


Figure 1. The Ornithine cycle

metabolic reaction that had received a great deal of attention. At the outset, he had no specific hypotheses about the reaction mechanism, but a number of more general questions: Is ammonia an obligatory intermediate; and how do rates of urea formation from various amino acids compare? These were not new questions, but Krebs thought that the tissue slice method would give him greater flexibility and more quantitative precision in seeking answers than did the methods used previously.

Krebs carried out his first experiment with alanine. The amount of urea produced in this experiment was much less than estimated according to the assumed equation of complete oxidation. Next he compared rates of urea formation from glycine, from alanine, and from ammonium chloride, in each case with glucose present in the medium. He found very little urea formation from glycine or alanine, but substantial amounts from ammonium chloride. He also noted that the rate of formation of urea from alanine declined in the presence of glucose. Therefore, Krebs concluded that the glucose inhibited the formation of ammonia from the amino acid. He apparently accepted the received view that ammonia was an essential intermediate product, and spent about four weeks characterizing the formation of urea from ammonia: checking the quantitative relations and the necessity of aerobic conditions, and testing the effects of changes in pH. He verified that the reactions proceeded only in liver tissue. All of this work was essentially a verification of known results.

From this point on, the work was carried on with the assistance of a new medical student, Henseleit. Krebs now turned back to determining the initial source of the urea nitrogen, which he presumed to be the amino acids. Testing alanine, phenylalanine, glycine, cysteine and cystine, he found they all produced urea at lower rates than did ammonium chloride. He also included other substances that might contribute amino groups that would be oxidized to ammonia, with the same result. Similar negative results were obtained in comparisons of ammonium chloride alone and in combination with amino acids; none of the combinations yielded urea at a higher rate than ammonium chloride alone.

During the first two weeks in November, the investigators turned to a new line of inquiry: the influence of glucose, fructose, lactate, and citrate, all substances involved as intermediates in carbohydrate metabolism. They had no specific hypotheses, but were exploring in this direction because a difference had been found in urea production in liver slices from well-fed and starved rats.

On November 15, Henseleit was continuing these experiments, but also ran a test with the amino acid, ornithine, and with a combination of ornithine and ammonium chloride. The combination produced urea at an unexpectedly high rate, and Krebs immediately turned his attention to the ornithine effect. The laboratory logs (and Krebs' later recollections, as well) do not provide conclusive information as to why the ornithine experiment, which represented a departure from the current activity, was run at that particular time. Krebs in his recollections insisted that he took ornithine just because it was available. But Holmes speculates that he chose ornithine because the metabolic fate of ornithine was an unsolved problem. It is possible to speculate further about the reasons for the experiment, but we will leave the question unanswered here.

1.2.2. Determination of Scope

In investigating the ornithine effect, Krebs employed "a standard biochemical strategy: if a given compound exerts some particular action, check whether derivatives of that compound have similar actions." None of the substances tested had effects similar to the ornithine effect, and Krebs became more and more convinced that the effect was quite specific to ornithine, although he had no clear hypothesis of a mechanism to account for it. This phase of the inquiry extended from the middle of November to the middle of January, 1932.

1.2.3. Discovery of Reaction Path

On January 14, Krebs and Henseleit used, for the first time, new apparatus that permitted accurate comparison of the amounts of ammonia consumed with the amounts of urea formed. Although some of the results of the first experiments were ambiguous, it was fairly clear by January 23 that the ammonia was the precursor of all of the nitrogen in the urea.

Now some function had to be found for the ornithine, and Krebs gradually arrived at the conclusion that it served as a catalyst. While this conclusion might seem obvious to *us*, it was much less obvious in 1932, when the study of catalytic reactions was relatively new.

A known reaction existed, the conversion of arginine to urea and ornithine, that could serve as the second stage of the cycle. Krebs had, in fact, studied this reaction in an experiment performed the previous October. At some point, it occurred to him that this reaction might enter into the picture. The fact that arginase is abundant in the livers of animals that excrete urea seemed significant. While Krebs was trying to conceive of a specific reaction path for the catalytic action of ornithine, he continued to direct Henseleit in experiments to elucidate further the ornithine effect, and also its interaction with arginine. During March, they also performed experiments to show specifically that the ornithine effect could be obtained with very small amounts of ornithine (in relation to the amounts of urea produced), and must therefore be catalytic. A very successful experiment of this kind was performed on April 13, in which 24.5 molecules of urea were formed for each molecule of ornithine that was present.

Gradually, Krebs inferred a specific reaction path consistent with all the known facts. On chemical grounds, it was evident that the conversion of ornithine to arginine could not proceed in a single step, and the theory was improved when Krebs found in the literature a 1930 paper reporting a substance, citrulline, that had the properties of a satisfactory intermediate between ornithine and arginine. Even before he obtained some citrulline, with which he could test this hypothesis, he felt sufficiently confident of his theory (*sans* the citrulline intermediate) to publish it. On April 25, five days before his paper appeared, he performed a test with citrulline, and by the middle of May, on the basis of further experiments, Krebs sent off a second paper describing the elaborated theory. The ornithine cycle as it was understood and depicted in 1932 is shown in Figure 1. Other researchers have since

further elaborated the steps in the cycle, and the ornithine cycle as we understand today is somewhat more complex. (See Lehninger 1982)

2. Description of KEKADA

In this section, we describe the KEKADA system, a computer program that simulates Krebs' discovery process.

2.1. Production System

The KEKADA system is implemented in the production system language OPS5 (Brownston, Farrell, Kant & Martin, 1985).

A production system consists of two main components: a set of condition-action rules or *productions*, and a dynamic *working memory*. The system operates in cycles. On every cycle, the conditions of each production are matched against the current state of the working memory. From the rules that match successfully, one is selected for application. When a production is applied, its actions alter the state of working memory, so that new productions may match the working memory on the next cycle. The cycles of matching and acting continue until no rules are matched by the working memory elements or a stop command is encountered.

2.2. Representation of Processes

The discovery heuristics of the KEKADA system are stated as OPS5 productions. Each rule contains a set of *conditions* describing the system's hypotheses or specifying patterns that may occur in the data. In addition, each rule contains a set of *actions*, which are responsible for formulating hypotheses, changing confidences in the hypotheses, suggesting new experiments, etc.

On each cycle, one of the matching rules is selected for action and the associated actions are carried out. When two or more rules match, the system prefers the rule that matches against elements that have been added to memory most recently; if there is more than one such rule, then it chooses the one that is most specific.

2.3. Representation of Data

Working memory elements are represented as attribute-value pairs. Among the important categories of working memory elements are *process*, *substance*, *experiment*, *supplementary fact*, and *hypothesis*.

Process. Process elements , which describe chemical reactions, have the following attributes: inputs, outputs, likely locus of reaction, name, and a flag indicating whether the description of the process may be incomplete. An *is-a* attribute names the class of processes to which the individual process belongs.

Substance. Substance gives information about a given substance (an amino acid or some other substance). As attributes, it has the name of the substance, its chemical formula, the classes to which it belongs, its cost, and its availability.

Experiment. The attributes of experiment elements are: inputs, conditions for carrying out, place for carrying out, initial quantities of inputs, flags indicating what is to be measured when the experiment is carried out.

Supplementary Fact. Supplementary facts, which give additional information about a process, have the name of the process, a locus, and a measure of confidence that the process takes place at this place. They also have attributes that name a condition and give a measure of the confidence that the process takes place under this condition.

Hypothesis. A hypothesis is a description of how a phenomenon or process that has been noted might have taken place. Associated with a hypothesis is a measure of confidence in its truth.

A hypothesis about a reaction is represented at one of the following four levels of abstraction: (1) the reaction is viewed in terms of the inputs and the outputs. (Examples:"in a reaction some amino acids may produce urea" or "ornithine and ammonia produce urea"), (2) its description is given in terms of compound groups. (Example: "NH_2COOH group in arginine comes from ornithine"), (3) its description is given in terms of simple groups. (Examples: "amino acids contribute their amino group to urea" or "ornithine may donate an amino group to urea", (4) its description is given at the atomic level (Example: "C in urea comes from carbon-dioxide").

These levels of abstraction are among the levels that have been in widespread use in chemistry

since the mid-nineteenth century.

2.4. Representation of Confidence Measures

Confidence in a hypothesis is represented by a 5-tuple:

1. Success: the number of experiments that have verified a universal hypothesis about a class or a hypothesis in general
2. Failure: the number of experiments that have falsified a hypothesis.
3. Failed-effort: the amount of effort spent to find positive instances.
4. Implied-success: a fact that is a positive indication, but inconclusive, that the hypothesis may be true.
5. Implied-failure: a fact that indicates, but not conclusively, that the hypothesis may be false.

These attributes seem to represent many of the ways in which people evaluate hypotheses, for they make such comments as: "There are many facts indicating the truth of this." "If after spending so much effort I still cannot prove this, probably it is false." "Three experiments have disproved this hypothesis."

We convert the values of the attributes into numbers by assuming that each fact increments the appropriate attribute by one unit. That is to say, if a fact indicates that a hypothesis is probably false the implied-failure slot is incremented by one. This rough scheme seems to work satisfactorily for a realm like scientific discovery where matters are, at best, highly conjectural.

2.5. Processes and Heuristics

The overall organization of KEKADA is based on the two-space model of learning proposed by Simon and Lea (1974) shown in figure 2. The system searches in an instance space and a rule space. The possible experiments and experimental outcomes define the instance space, which is searched by performing experiments. The hypotheses and other higher-level descriptions, coupled with the confidences assigned to these, define the rule space. On the basis of the current state of the rule space (what hypotheses are held, with what confidences), the system chooses an experiment to carry out. The outcome of the experiment modifies the hypotheses and confidences.

Operators to carry out the search in the instance space: The heuristic operators used to

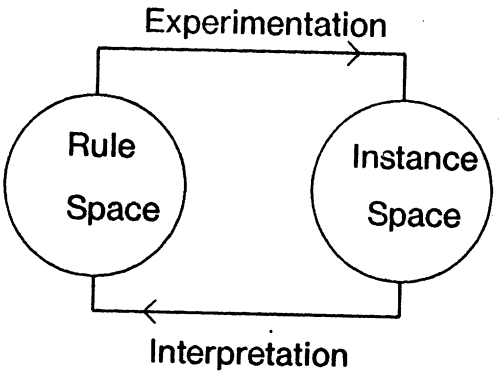


Figure 2. Two-space Model of Learning

search the instance space fall in two categories:

1. **Experiment-proposers**, which propose experiments based on existing hypotheses.

2. **Experimenters**, which carry out experiments.

Operators to carry out the search in the rule space: The heuristic operators used to search the rule space fall in the following categories:

1. **Hypothesis or strategy proposers**: When the system has decided to focus on a particular problem, these decide which hypothesis or hypotheses to focus on or which strategy to adopt for the work on the problem.

2. **Problem-generators**, which propose new problems or subproblems on which the system can focus attention.

3. **Problem-choosers**, which choose which is task the system should work on next.

4. **Expectation-setters**, which set expectations for the experiments to be carried out.

5. **Hypothesis-generators**, which generate new hypotheses about unknown mechanisms or phenomena.

6. **Hypothesis-modifiers**, which modify the hypotheses on the basis of new evidence.

7. **Confidence-modifiers**, which modify confidences about hypotheses on the basis of the interpretations of experiments.

Heuristics to make choices: In KEKADA, only certain alternatives are applicable at any stage. If more than one alternative is applicable, heuristics called **decision-makers**, are used to choose between the operators. **Decision-makers** determine, for example, which of the various problems proposed by problem-proposer heuristics will be worked on.

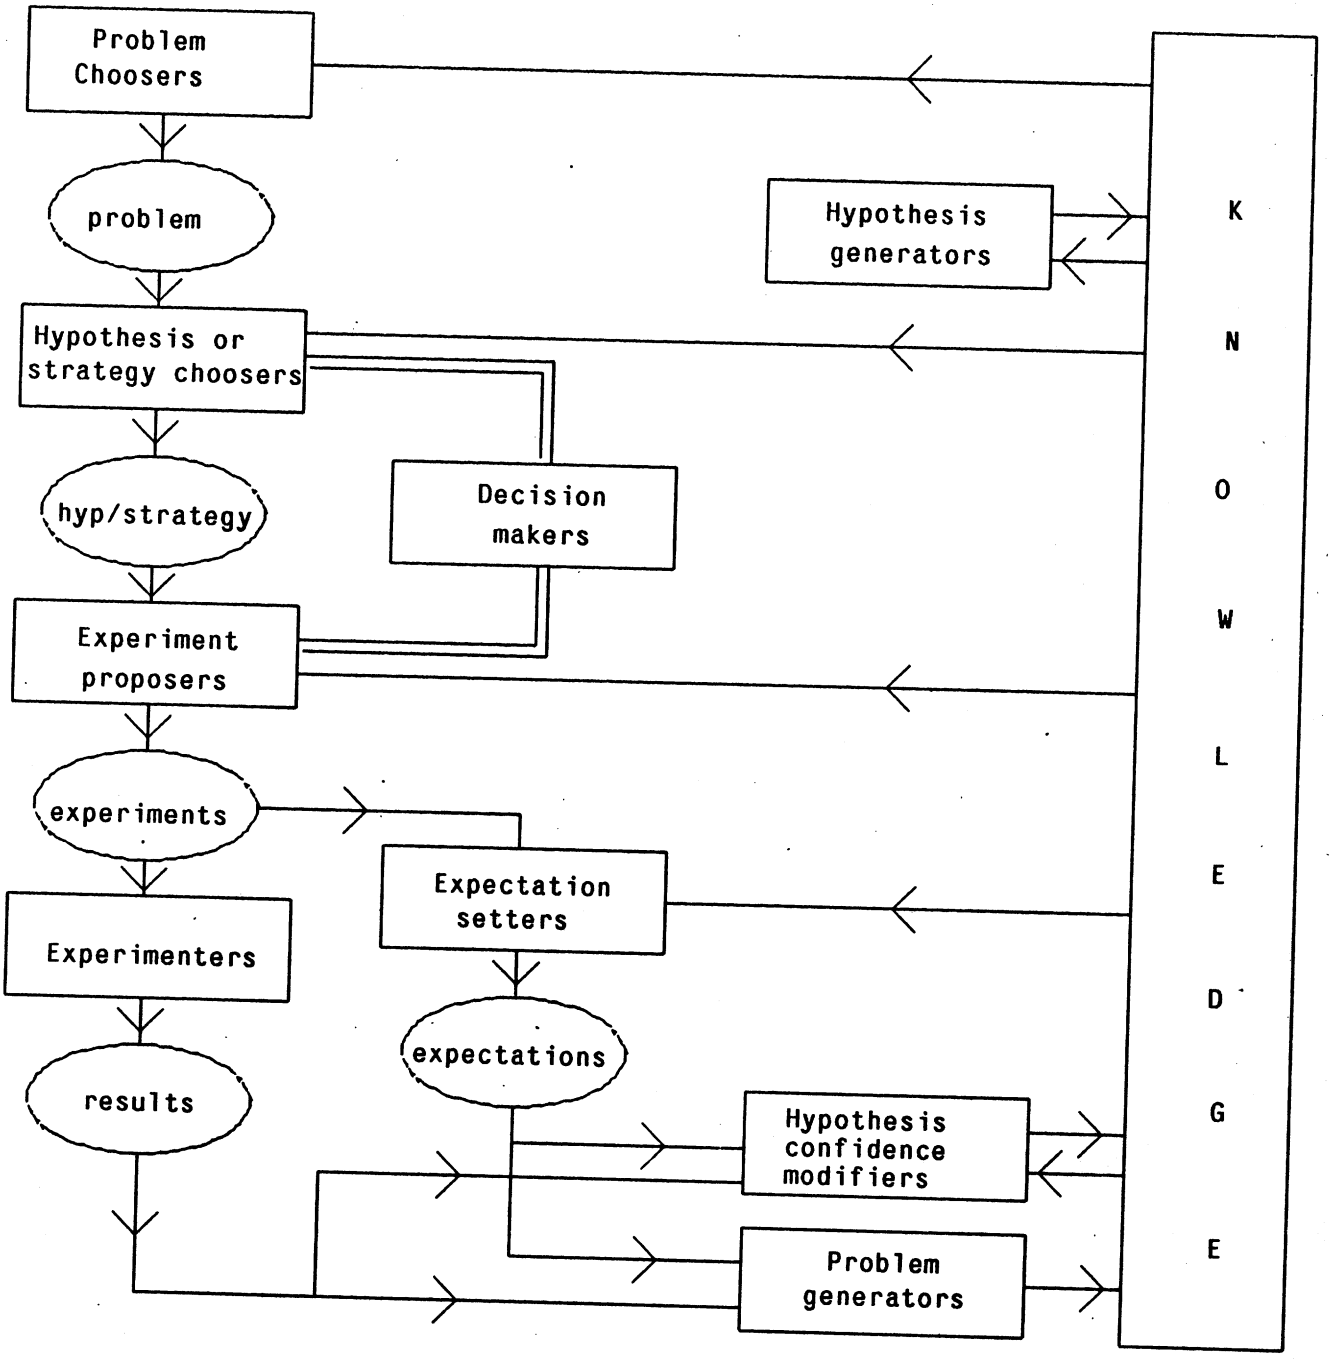


Figure 3. Interaction of heuristics

2.5.1. Interaction of Heuristics

We now describe in more detail how the heuristics in various categories interact as the system works on a problem. If the system has not decided on which task to work (or in situations where new tasks have been added to the agenda), *problem-choosers* will decide which problem the system should start working on. *Hypothesis-generators* create hypotheses when faced with a new problem.

Thus at any given stage a certain number of hypotheses with varying confidences are present in working memory.

When working on a given task, the *hypothesis* or *strategy proposers* will choose a strategy to work on. Then the *experiment-proposers* will propose the experiments to be carried out. Both of these type of heuristics may need the *decision-makers*. Then *expectation-setters* set expectations and *experimenters* carry out experiments. The results of the experimenters are interpreted by the *hypothesis modifiers* and the *confidence modifiers*. When applicable, *problem-generators* may add new problem to the agenda and preempt the system to focus on a different problem.

```
; Name of the rule
(p note-surprise

 ; LEFT HAND SIDE (Condition of the rule)
 ;if this rule is of type problem-generator
 (context ↑name problem-generator)
 ;if given experiment with inputs <i1>, <i2>, and <i3> is found to
 ; to have output <o> and the rate-of-output <r-o>
 (experiment ↑status just-done ↑input1 <i1> ↑input2 <i2> ↑input3 <i3>
     ↑expected-output <e-o>
     ↑expected-lower-bound <lb> ↑expected-upper-bound <ub>
     ↑output <o> ↑rate-of-output <r-o>)
 ;and if expectations set with the experiment are: output <e-o>
 ; upper-bound on the output-rate <ub>, and
 ; lower-bound on the output-rate <lb>
 ; and if the results of the experiment violate these expectations
 -(experiment ↑status just-done ↑input1 <i1> ↑input2 <i2> ↑input3 <i3>
     ↑expected-output <e-o>
     ↑expected-lower-bound <lb> ↑expected-upper-bound <ub>
     ↑output <e-o> ↑rate-of-output { >= <lb> <= <ub>})

 ;THEN
-->

 ;RIGHT HAND SIDE (Action taken if the condition is met)
 ;Note this as a surprise and add to the agenda, with associated
 ;information on actual and expected outputs.
 (bind <newid>)
 (make agenda ↑task-name <newid>)
 (make surprise ↑name <newid> ↑input1 <i1> ↑input2 <i2> ↑input3 <i3>
     ↑expected-output <e-o>
     ↑expected-lower-bound <lb> ↑expected-upper-bound <ub>
     ↑output <o> ↑rate-of-output <r-o>))
```

Figure 4:OPS5 code for the surprise-detector heuristic

Since the interaction of these mechanisms can produce surprise, a very important incident in the discovery process, we will discuss how the concept of 'surprise' is represented in the program, before proceeding to discuss the heuristics in more detail.

2.6. Surprise

The ability to react to surprise, and to attempt to explain the puzzling phenomenon, plays an important role in many discoveries. KEKADA has an ability to notice a phenomenon as 'surprising.' Before any experiment is carried out, expectations are formed by *expectation-setters* and are associated with the experiment. These expectations consist of expected output substances of the reaction, and expected lower and upper bounds on the quantities or the rates of their outputs. If the results of the experiments violate these bounds, this is noted as a *surprise*. We give in figure 4 a slightly simplified version of the OPS5 code (See, Brownston, Kant, Farrell, & Martin, 1985) which implements the PG1 heuristics: if the outcome of an experiment violates the expectations for it, then make the study of this puzzling phenomenon a task and add it to the agenda. The bold lines beginning with a semicolon (;) are comments about the OPS5 code.

Now we will the discuss the heuristics in the program in detail.

2.7. Problem-choosers

[PC0] Take into consideration all the tasks on the agenda.

[PC1] If no analytic methods exist to measure the outputs of a process or to carry out the process, eliminate it.

[PC2] If the task is not regarded as very important by the discipline, eliminate it.

[PC3] If a new method significantly increases the rate at which a task can be carried out and its accuracy, then prefer it over another method , other things being equal.

[PC4] If there are no other criteria applicable, then make a random choice.

[PC5] If you do not have the skill to study a task, eliminate it.

[PC6] Other things being equal, prefer the task that can be studied more accurately.

[PC7] Other things being equal, prefer the task which can be carried out fast.

[PC8] If a new task to study a puzzling phenomenon is being added to the agenda, prefer it over

all the other tasks, making it the focus of attention.

2.8. Problem-generators

[PG1] If the outcome of an experiment violates expectations for it, then make the study of this puzzling phenomenon a task and add it to the agenda.

2.9. Decision-makers

The decision-making process is represented by a set of rules. Different sets of rules are used for different types of decisions. There are three such sets:(1) Rules for choice among biological processes, (2) Rules for choice among substances, (3) Rules for defining an initial ordering.

Rules for choice among processes: The following set of rules is used for deciding which one of the given set of processes is to be chosen for study.

[DM1] If the output of a process is not measurable, eliminate it.

[DM2] If the typical rate of progress of a process is significantly more than that of another process, prefer it.

[DM3] If there are no other criteria for choice between two processes, choose one of them at random.

Rule for choice among hypotheses: [DM4] If confidence in one hypothesis is higher than in another hypothesis, with respect to any one of the slots, then prefer the former hypothesis.

Rules for choice among substances: The following rules are used to decide which one of the given set of substances should be chosen for study.

[DM5] If the cost of a substance to be tested is too high, eliminate it.

[DM6] If a substance to be tested is not easily available, eliminate it.

[DM7] If the cost of two substances is low and both are available, and they are being tested because they are similar to a particular substance, then give preference to the substance that is most similar to the given substance. (In the present implementation, a partial ordering is defined on various substances indicating their similarity to ornithine.)

[DM8] If there is no other criterion for choice between two substances, choose one of them at random.

Defined priority: [DM9] Sometimes the investigators' experience before his current research program was undertaken or the nature of the hypotheses defines a partial order on the hypotheses. For example, the hypothesis that a given surprising reaction may be common to a class of substances is normally considered before other hypotheses, for experience shows that work on this kind of a hypothesis is likely to be very productive. Correspondingly, the system has the following predefined order for hypotheses: (1) a causal explanation that substance S, which is previously known to have a stimulating effect on a process, may be necessary for the process, (2) divide and conquer, (3) a hypothesis about scope of a phenomenon, (4) any other hypotheses. But since we do not have exact data on Krebs' previous experience, in the cases where we have used a pre-defined order, it is possible that he actually used decision-making rules like other rules in the DM category.

[DM10] In running this system for the urea example, in a few cases where the biochemical heuristics Krebs used to make the choice are not clear to us, the choice was made by the user. Interaction with the user allows the system to make the discovery of the ornithine cycle along different pathways.

2.10. Experiment-proposers

These heuristics propose to carry out an experiment whose findings could change confidences in existing hypotheses or verify or falsify hypotheses.

[EP1] If the preferred strategy is to see if a surprising phenomenon is common to a class of substances, then use the decision-makers to choose a substance A in that class, and decide to study the phenomenon with A as a reactant.

[EP2] If you are studying a phenomenon with A as reactant, and there is a hypothesis that A produces C with B as an intermediate product, then carry out experiments on A and on B, and compare rates of formation of C from A and from B.

[EP3] If you are studying a phenomenon with A as reactant, and there is a hypothesis that A and B react to form C, carry out experiments on A and B in combination and on A and B separately.

[EP4] If the chosen hypothesis is that in the reaction under study A and B react together to form C, and that B is the source of one of the components of C, then carry out an experiment with A and B together, measuring appropriate parameters to determine the quantity of C in relation to the quantities

of A and B.

[EP5] If the chosen hypothesis is that the reactant A in an experiment is a catalyst, or if the chosen hypothesis is that A donates some element or group and no other possibility of A donating a group or element exists , then carry out the experiment over long periods but with very low concentration of A.

[EP6] If the chosen hypothesis is that the reason for a surprising outcome may lie in an unknown substance, guess the substance to one that is related to the process (i.e. a substance that earlier experiments seem to have associated with the given process or the same class of the process.) choose one of the substances using decision-makers, and carry out an experiment on it.

[EP7] If the goal is to study a particular reaction in detail, carry out the reaction under various conditions. (Draw on general knowledge about the process to design the experiment.)

[EP8] If the preferred hypothesis is to study the relation of a related fact to a surprising phenomenon, and the related reaction and the given phenomenon both produce the same output, create two new hypotheses and add them to the hypothesis set: (a) Hypothesize a class and predict that it will produce this output. (b) If there is evidence for a hypothesis that the given reactant could be an intermediate, then create this hypothesis. (Note that this rule operates as a hypotheisis generator or modifier.) Finally study one of the newly identified hypotheses.

2.11. Expectation-setters

[ES1] If the same experiment was carried out before, the expected value is the mean of the previous outcome quantities, while the lower bound is the lowest quantity observed previously minus a tolerance factor. The upper bound is the largest quantity observed previously plus a tolerance factor.

[ES2] If no experiments with the given inputs have been carried out before, and no experiments with similar inputs (e.g., experiments with different amino acids), then the expectation is a predetermined value assumed to reflect the prior knowledge of the investigator.

[ES3] If experiments are carried out on members of a class, the expectation for the class (that is, for all members of the class) is modified to reflect the outcome. Expectations for a class are used as expectations for members of the class not previously tested.

[ES4] When a new experiment has been carried out, update the summary information elements.

2.12. Experimenters

In the current system, there are no experimentation heuristics.

[E1] The outcomes of experiments are supplied interactively by the user.

2.13. Hypothesis-generators

[HG1] If a surprising outcome occurs involving A as one of the reactants, then hypothesize that there is a class of substances containing A (or its derivatives) that will produce the same outcome.

[HG2] If there is a surprisingly low output of substance A under some experimental conditions but not others, and if it is possible that another substance S is present in the latter conditions but not the former, hypothesize that the absence of S is causing the low output.

[HG3] If a reaction has subprocesses and the outcome of the reaction is surprising, hypothesize that the surprising result depends on one of the subprocesses (divide and conquer strategy).

[HG4] If a reaction produces some output, create hypotheses asserting which reactant donates which group to the output substance and that a reactant may be a catalyst.

[HG5] If a one-step stereochemical transformation from inputs to outputs of a reaction is not possible, then create the hypothesis that an intermediate exists. Otherwise create a hypothesis that there is a one-step stereochemical reaction.

[HG6] If the goal is to study a puzzling phenomenon and If the given reaction and the surprising phenomenon contain two common substances, then create a hypothesis that they may be related.

[HG7] If the output from A and from B is different from the sum of the outputs from A and B, then create hypothesis that there is mixed action from A and B otherwise create the hypothesis that the effect is additive.

[HG8] Properties of a class are true for a member.

Hypothesis modifiers:

[HM1] If A and B react to produce C, and B does not act without A, and the amount of product is large relative to the amount of A, then conclude that A is a catalyst.

[HM2] If the preferred strategy is to verify the existence of an intermediate in a reaction, then

carry out the following three steps: (1) Consider substances structurally intermediate between the inputs and outputs as possible candidates; (2) evaluate the plausibility of each candidate's being intermediate in the reaction; (3) choose the substance (if any) which has been evaluated most likely to be an intermediate in the reaction.

[HM3] (This actually is a set of heuristics.) Given a reaction in an incomplete and unbalanced form, use balance heuristics listed below to attempt to balance it.

Rules applicable at levels of abstraction corresponding to simple and compound groups:

[B1] If the coefficient of a substance in the reaction is known, then convert the groups contained in the substance into FLOATING GROUPS. (E.g., if ammonia is known to have one amino group and the coefficient of ammonia is 2, then produce two floating amino groups on the appropriate side.)

[B2] If no other rule is applicable, change the level of abstraction.

[B3] Cancel equal groups on the right and left hand sides

[B4] If a substance on one side has a group A, and there are no floating groups A on the same side, and there are a certain number of floating groups A on the other side of the reaction, then determine the coefficient of the substance by a simple match.

[B5] If there are floating groups of A on one side, and there is no reactant having A on the other side whose coefficient is not known, and one of the other substances present has group A, then guess this substance as the possible reactant of the reaction.

Rules applicable at atomic level of abstraction:

[B6] If the coefficient of a substance in the reaction is known, then convert the atoms of the substance into FLOATING ATOMS. (E.g., it is known that ammonia is NH_3 and that the co-efficient of ammonia is 2, then produce 6 floating atoms of H and 2 of N.)

[B7] If no other rule is applicable and the reaction is not balanced, then conclude that the reaction cannot be balanced.

[B8] Cancel identical atoms on the right and left hand sides

[B9] If the substance on one side has an atom A, and there are no floating atoms A on the same side, and there are a certain number of floating atoms A on the other side of the reaction, then

determine the coefficient of the substance by simple match.

[B10] If there are floating atoms of A on one side, and there is no reactant having A on the other side whose coefficient is not known, and one of the substance present has atom A, then guess this substance as the possible reactant of the reaction and attempt to balance this reaction.

[B11] If you can account for both the sides at the atomic level then the reaction is balanced.

Hypotheses in the system are in one or the other of two states: active or inactive. When KEKADA has very low confidence in an hypothesis; it removes that hypothesis from consideration and makes it inactive. The following heuristics are used by the hypothesis-removers.

[HM4] If the amount of effort spent on an existential hypothesis reaches a specified high value, make the hypothesis inactive.

[HM5] If the number of experiments that falsify a given hypothesis reaches a specified high value , make the hypothesis inactive.

[HM6] If by experiment it is found that the source of a group or element G is substance A, then eliminate hypotheses that any other substance donates group G, and create a clue that A donates G (i.e., increase the success-slot of the confidence in the hypothesis by 1).

2.14. Confidence-modifiers

The following rules modify confidences in the hypotheses that the system holds:

[CF1] If there is a hypothesis that A produces C with B as an intermediate, and if experiments show that the production from B is slower than from A, then increase the implied-failure of the hypothesis by 1 ; else increase the implied-success by 1.

[CF2] If there is a hypothesis that A and B react together to produce C, and A and B together do not produce more output than A or B individually, then increase the implied-failure by 1; else increase the implied-success by 1.

[CF3] The failed effort slot in the confidence slot stores the amount of effort spent on a hypothesis or a problem.

[CF4] If there is a hypothesis that a reaction will take place under certain conditions and there is a positive result from the experiment under the conditions, then the success slot is increased by 1.

[CF5] If there is a hypothesis that a certain reaction will take place under certain conditions and

there is a negative result from the experiment under the conditions, then the failure slot is increased by 1.

2.15. Hypothesis or Strategy Choosers

[HSC1] If no hypothesis is chosen for consideration, then evaluate the alternatives and choose one of them according to decision-making rules.

[HSC2] If the chosen strategy is to study a subprocess in detail, then choose one of the subprocesses to study using the decision-makers.

2.16. Subject-matter Knowledge

Any scientist has a certain amount of background knowledge when he begins his research. While he is doing research, he may acquire additional knowledge through literature surveys or through discussions with colleages. Scientists with different background knowledge may follow different courses of research. Correspondingly, KEKADA needs background knowledge before it is run and can acquire additional knowledge while it is running. Differences in its background knowledge may cause it to work on different problems or follow different courses of action on any particular problem.

When provided with knowledge corresponding to that which Krebs had, KEKADA follows a path of discovery similar to that actually followed by Krebs. We discuss this knowledge in further detail in the paragraphs below.

2.16.1. Background knowledge

The background knowledge takes two forms. Some of it is contained in domain-specific heuristics embedded in KEKADA, that are described in previous subsections. Other knowledge is created by using 'make' statements before KEKADA is run. 'Make' statements create initial working memory elements of various kinds. These working memory elements constitute the system's initial knowledge. Prior knowledge falls in 3 categories: knowledge about substances, knowledge about processes, and knowledge about previous experiments.

1. Knowledge about substances including the amino-acids , glucose, etc includes their chemical formulae, cost, availability and the class to which they belong. KEKADA also knows the typical low, medium and high quantity of a substance to be used in the

experiments. Besides KEKADA knows the partial order relation stating which of two substances is more similar to a given substance.

2. KEKADA also has knowledge about chemical reactions. This includes the inputs, the outputs, the class to which the reaction belongs and some supplementary facts. When the exact place or condition under which the process takes place is not known, supplementary facts may give various possible places or conditions where the process might be taking place. Also associated with each supplementary fact is the confidence that the process does take place at this place. The knowledge also includes various possibilities previously considered likely regarding where the process takes place.

3. Before Krebs undertook the research program that led to the ornithine cycle discovery, he had read about the experiments others had carried out on urea synthesis. It is assumed that his initial expectations about the outcomes were set either by the previous experiments or by some previously known theory. Therefore the summary of these previous experiments is made available to KEKADA. KEKADA uses this knowledge only to set the expectations for the initial experiments.

2.16.2. Acquiring knowledge through literature and from colleagues

Apart from the results of his own experiments, Krebs' research was also influenced by such factors as the availability of a new instrument and the research results published by other scientists. Correspondingly OPS5 allows the creation of new working memory elements at intermediate stages in the progress of KEKADA to allow such factors to enter.

3. Simulation of the Discovery of the Ornithine Cycle

We present here the log of a particular run of KEKADA described in terms of the numbered heuristics we have described. An asterisk (*) denotes repeated application of a set of heuristics. Seq*i* names the sequence of firings of heuristics that is enclosed in the following pair of dashed lines.

Heuristics	Results
PC0	Considers various alternative tasks on the agenda. Considers as possible candidates urea synthesis and synthesis of some fats, proteins, and fatty-acid degradation, etc.
PC1-7*	Chooses urea synthesis from among the various alternatives and creates a goal to study urea synthesis using the tissue slice method.
HSC1	Considers alternative hypotheses on urea synthesis, viz., amino-acids may produce urea, pyrimidines may do so, cynates may be precursors to urea, etc.
DM4*	Considers it likely that amino-acids may produce urea.

EP1	Considers various amino-acids as alternatives.
DM5-8*	Chooses alanine.
HG8	Assigns to alanine the properties of the class, amino-acid.
EP2-3	Decides for an experiment on alanine and on ammonia. Decides for an experiment on both combined together.
ES1-3*	Sets expectations for these experiments.
E1, ES4, CF1-2*	Asks user for the results of experiments, modifies confidences.
PG1, PC8	Notes the result of the experiment on alanine as surprising, and makes it focus of attention, creates the following hypotheses:
HG5, B1-11*	Studies alanine to urea reaction, *decides that intermediate exists.*
HG2	*Some essential substance is missing from the tissue slice preparation.*
HG3	*The reason for surprise may be one of the sub-reactions.*
HG1*	*The phenomenon may be common to some or all elements of a class.*

[seq0]

[Begin seq0]

HSC1	Evaluates the alternatives.
DM4, 9*	Decides to consider the hypothesis that an absense of a substance may be causing the surprise.
EP6	Guesses the substances which may be present-various substances involved in carbohydrate mechanism.
DM5*	Chooses glucose.
ES3	Sets expectations for the experiment.
E1, ES4	Asks user for output for an experiment on alanine and glucose.
CF3	Modifies failed-effort slot in hypothesis.

[End seq0]

[Repeats seq0 for various substances.]

HM4 Makes inactive the existential hypothesis that there may be a substance missing.

HSC1 Evaluates the alternatives.

DM4, 9* Decides to consider the hypothesis that the cause of the process may be in one of the subprocesses.

HSC2, DM1 Decides to study the subprocess of urea synthesis from ammonia.

EP7, ES1, E1, Es4, CF4-5*

Carries out experiments on urea formation on ammonia under various conditions of PH, aerobicity and in various organs, study quantitative relations.

[seq1]

[Begin seq1]

HSC1 Evaluates the alternatives.

DM4* Decides to consider the third hypothesis: that surprise may be limited to a class.

EP1 Decides to list possible amino-acids for consideration.

Dm5-8* Chooses cysteine.

HG8 Assigns properties of the class to cysteine.

EP2-3 Decides for an experiment on cysteine and on ammonia. Decides for an experiment on both combined together.

ES1-3, E1, ES4 CF1-2*

Sets expectations for these experiments. Asks user for the results of the experiment. Modifies the confidences in hypotheses.

[End seq1]

[Repeats seq1 on other amino acids, last one being ornithine]

PG1, PC8 Notices the ornithine effect and makes it the focus of attention. Creates following hypotheses.

HG7 New clue is created for *mixed action of both the inputs.*

HG4* *Hypotheses about who donates what to the reaction.*

HG5, B1-11* *Intermediate exists.*

HG4* *Possibility that ornithine or ammonia is catalyst.*

HG1* *Possibility that the phenomenon may be common to a class of substances.*

HG6* *Possibility of relation to similar reactions.*

[seq2]

[Begin seq2]

HSC1 Evaluates the alternatives.

DM4-9* Decides to study the scope of the phenomenon. Considers that the phenomenon may be common to amino-acids.

EP1 Considers various amino-acids.

DM5-8* Decides on an amino-acid as the choice.

HG8 Assigns properties of the class to that amino acid.

EP2-3 Decides for an experiment on the amino-acid leucine and on ammonia, separately and combined.

ES1-3, E1, ES4, CF1-3*

Sets expectations for these experiments Asks user for the results of experiments. Changes the implied-failure in hypotheses about how urea is formed reduce the failed-effort slot in the hypothesis asserting that the phenomenon may be common to a class.

[End seq2]

[Repeats [seq2] for various amino-acids]

HM4 Removes the description that some amino acids might produce urea.

[seq3]

[Begin seq3]

HSC1 Evaluates the alternatives.

DM4-9* Decides to study the hypothesis that the scope to the surprise may be common to some or all amines.

EP1 Considers various amines.

DM5-8* Decides on putrescine. Decides for an experiment on putrescine and ammonia.

HG8 Assigns the properties of its class to putrescine.

ES3, E1, ES4, CF3

Sets expectations for these experiments Asks user for the results of experiments. Reduces the failed-effort slot in the hypothesis asserting that the phenomenon may be common to a class.

[End seq3]

[Repeats [seq3] for various amines.]

HM4 Removes description that some amines might produce urea.

[Repeats [seq3] for various carboxylic acids.]

HM4 Removes description that some carboxylic-acids might produce urea.

HSC1 Evaluates the various alternatives.

DM10 User decides to study the hypothesis that source of NH_2 group in urea is ammonia.

EP4, ES1, E1 Carries out the experiment after setting expectations.

HM6 Concludes that the source of amino group is NH_3.

HSC1 Evaluates the various alternatives.

DM10 User chooses to study the related reaction: arginine reaction.

EP8, DM10 Two possible hypotheses are created: *arginine may be intermediate,* or *there may be a class of substances exhibiting reaction similar to arginine reaction.* Considers the second hypothesis.

EP1 Considers substances in guanidino class.

DM5* Chooses guanidine as substance for reaction.

EP1 Decides for the reaction on guanidine and ammonia.

HG8 Assigns properties of the class to guanidine.

ES3, E1, ES4, CF3

Carries out the experiment. Reduces the confidence in the existential hypothesis.

HSC1-DM10 Chooses the possibility that ornithine is catalyst.

EP5 Decides for an experiment to verify catalysis.

E1 Carries out experiments to check catalysis.

HM1 Concludes that ornithine acts as a catalyst.

B1-11* Balances the catalysis reaction.

HG5 *Creates hypothesis that there exists intermediate in the reaction.*

HM2, B1-11* *Creates candidates for intermediate.* Balances the reactions. Counts the number of inputs. Evaluates the intermediates. Chooses arginine.

HG5 *Creates a hypothesis that there exists intermediate in the reaction.*

(User, when asked to carry out a survey, creates elements corresponding to citrulline and other substances.)

HM2, B1-11* *Considers candidate substances* which are structurally intermediate between the inputs and the outputs of the ornithine to arginine reaction. Balances the reactions. Counts the number of inputs. Evaluates the plausibility of the candidate substances and chooses citrulline from them.

3.1. Overview of the Simulation

As we mentioned in the previous section, differences in background knowledge would lead KEKADA to follow a different research pathway. In the present section we will interprete the log we have displayed, which describes the behavior of KEKADA when placed in a situation similar to Krebs. In a few cases the choice between the alternatives was made by the user, because the heuristics Krebs used are not clear to us. Interaction with the user (which is indicated by (INT)) allows the system to make the discovery of the ornithine cycle along different pathways. It is possible to conjecture the reasons that might have lead Krebs to make the choices exactly the way he did, but given the uncertainty here, we decided to rely on user interaction to resolve the issue instead.

As in the earlier description of the actual history in Section 1 above , we divide our account into three phases: discovery of the ornithine effect, the determination of scope, and the discovery of the

Strategy of Experimentation

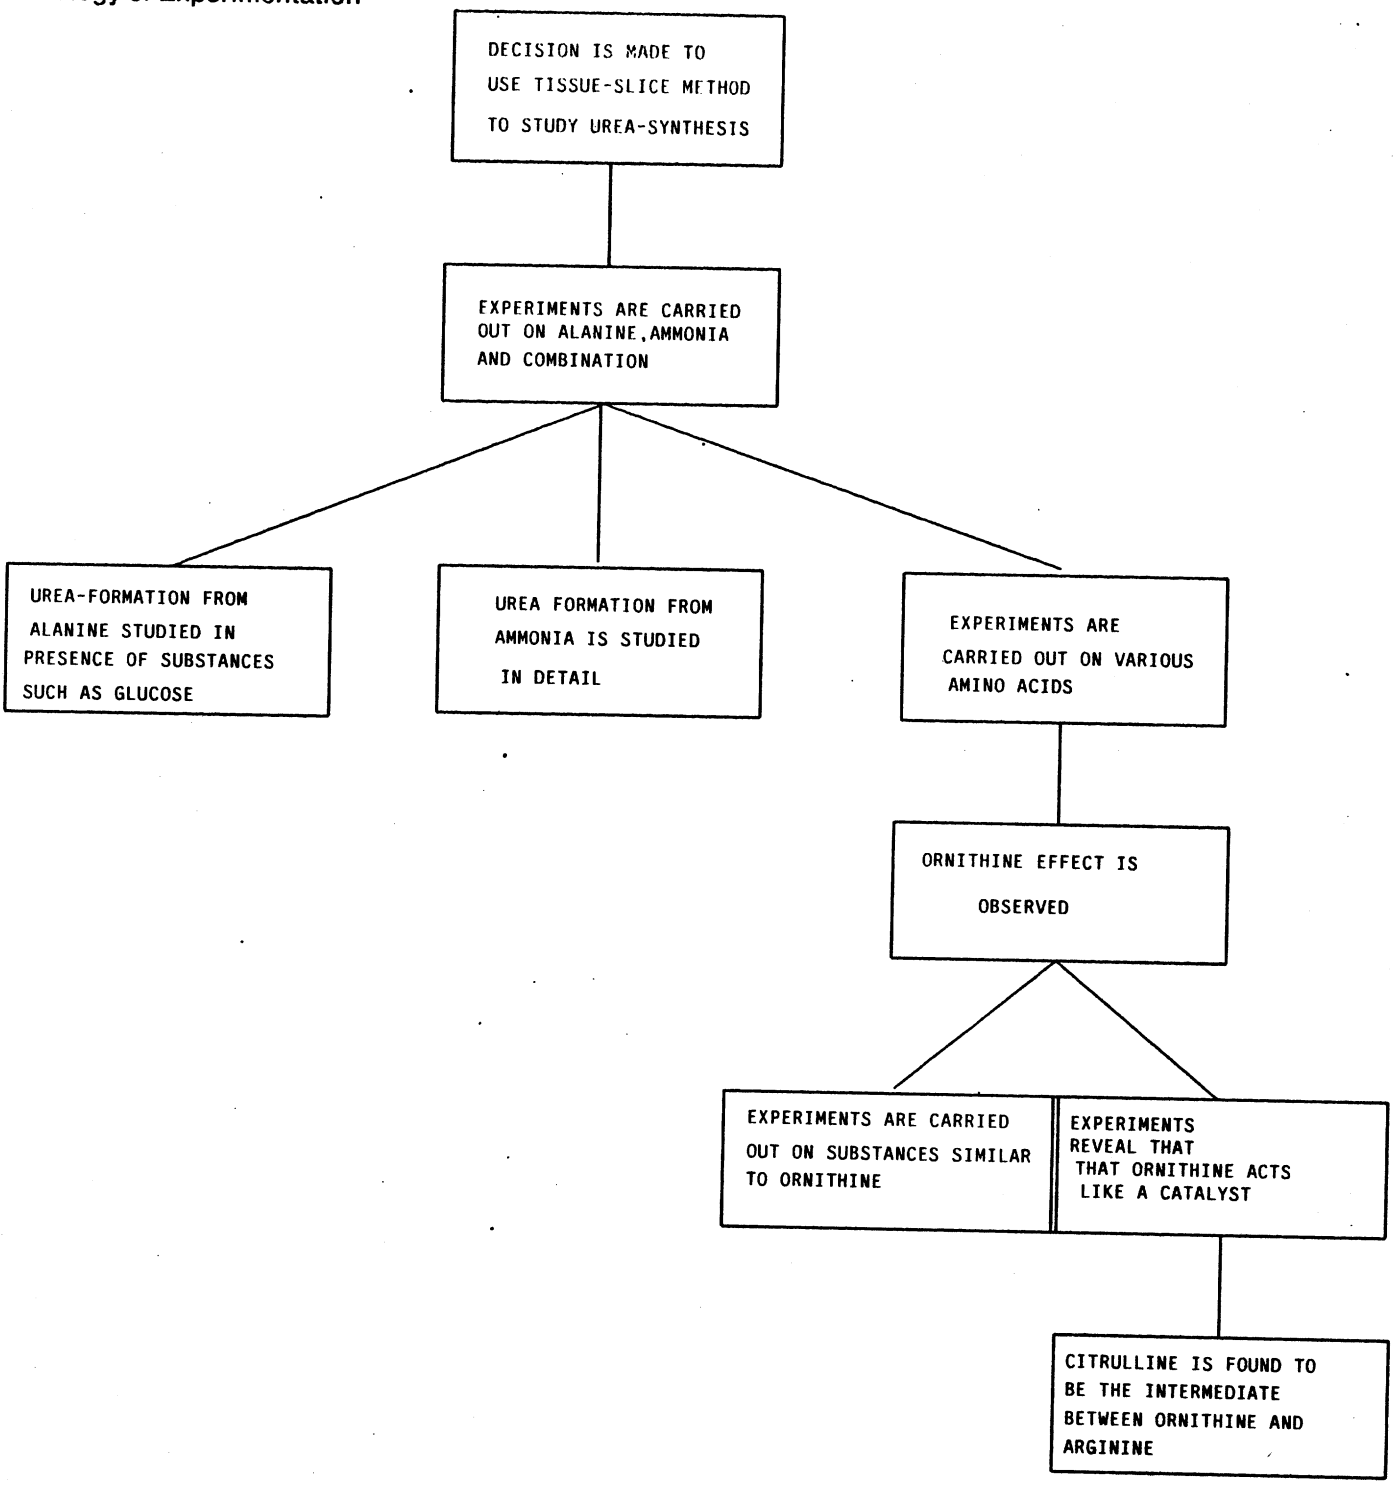


Figure 5. Progress of KEKADA in the discovery

reaction path. Major stages in these phases are depicted in Figure 5.

3.2. Simulating the Ornithine Effect Discovery

The first task of KEKADA is to select a research problem. It considers the various problems on its research agenda including urea synthesis and protein synthesis. Urea synthesis is a good choice for various reasons. Analytic methods are available for the measurement of urea. The rate of production of urea is quite high. It is also an unsolved problem regarded by the discipline as important.

Of course, these heuristics, interacting with the differing bodies of biochemical knowledge and skills possessed by different investigators might easily lead to the selection of different problems. In fact, few of Krebs' contemporaries were then studying the urea synthesis problem, and Krebs' specific choices were undoubtedly strongly influenced by his long exposure to the tissue slice method, and the comparative advantage that his skill with this method gave him in its use. Without a detailed knowledge of initial conditions -- in particular, of what the scientist knew and could do -- only hindsight could tell us what research problem he would choose.

Having selected its research problem, KEKADA now has the goal of finding the unknown mechanism by which urea is formed in living tissue. Prior knowledge in biochemistry proposes the following possible mechanisms, among others: (1) Amino-acids may be precursors of the urea. (2) pyrimidines may be the precursors of the urea.

The system considers the first alternative as more likely. It knows two possible ways in which this might happen.

1. Amino acids might donate their amino groups to form urea, with ammonia as an intermediate product in the process.

2. Amino acid and ammonia might react together to form urea.

A predetermined level of confidence has been assigned to each possibility. The inference is drawn that if ammonia is an intermediate, then urea will be formed more rapidly directly from ammonia than from an amino acid. The system decides to carry out an experiment with liver tissue on an amino acid, another on ammonia and a third on a combination of both. Differences in the outcomes of these three experiments should provide some evidence for choosing between the two hypotheses. Alanine

is selected (from a list of amino acids chosen by decision-maker heuristics) as the first amino acid to be tested.

Before the experiment is carried out, expectations are formed and associated with the experiment. These expectations consist of expected values, expected lower bounds, and expected upper bounds on the rates of production of the expected output, urea. The results of the experiment are provided by interaction with the user (INT), who is asked for the output substance, the rate of production of the output, and the quantity of output produced.

The first experiment on tissue slice with alanine produces very little urea, less than the lower-bound of the expectation. This result is noticed as a surprise, and whenever surprise occurs its cause becomes the focus of attention.

Now the system tries to discover why alanine, an amino acid, does not produce much urea in the tissue slice contrary to biochemical beliefs that amino acids are the sources of the nitrogen for urea, and that there should be no essential differences, on this point, among amino acids. Certain possible explanations or hypotheses for this surprising result are now created by the hypothesis-generator and modifier heuristics. In the presence of appropriate facts of biochemistry, these rules produce corresponding hypotheses or modify hypotheses. Three possible explanations are generated at this point:

1. Since alanine on liver tissue slice does not produce urea, and since it is assumed that alanine in the living organism does produce urea, there must be some essential substance, present in the organism, that is missing from the tissue slice preparation.

2. Using the heuristic that if there is a defect in a process made up of subprocesses the defect may be in one of the subprocesses, the inference is drawn that the defect may be in the subprocess that converts alanine into ammonia, or the subprocess that converts ammonia into urea.

3. There may be a class of substances, other than alanine, that produce urea.

The various experiments that the system now carries out are driven by these hypotheses, together with the two hypotheses about the urea synthesis mechanism introduced earlier. At the beginning, the system has no bias about these hypotheses -- confidence neither in their truth or their falsity. As the system carries out various experiments, the confidences in the hypotheses are modified according to the experimental results.

In response to the possibility that there is some other substance in whose presence alanine

produces urea, the system tries to identify this substance. Substances related to the surprising fact are considered likely candidates, especially substances that earlier experiments appear to have associated with urea synthesis. Here KEKADA adds such substances as glucose and fructose and reruns the experiments, without any change in outcome. These results do not falsify the assumption that there exists a substance in whose presence alanine would produce urea, but they do reduce confidence in the assumption. Each failed guess about the substance increases the failed-effort value by one, and when that value reaches a specified level, confidence in the hypothesis is low enough to remove it from further consideration.

The second -- divide-and-conquer -- hypothesis leads KEKADA to study the formation of urea from ammonia, and to repeat experiments to confirm previous knowledge about the reaction. The system confirms that aerobic conditions are required and that the pH must lie in a certain range. Experiments are also carried out to verify that only liver tissue is able to carry out the reaction. The experiments confirm previously established effects but do not reveal any reason for the surprising phenomenon.

The possibility next considered is that there may be a particular class of amino-acids that produce urea. On the basis of the third hypothesis that has been generated, KEKADA now repeats the original experiments with different amino acids. The first experiments do not produce much urea from the amino acids, and the confidences in the various hypotheses are changed accordingly. The expectation of output of urea from an amino acid is reduced, as is the expectation of an increase in the production of urea from ammonia in the presence of amino acid.

The next amino acid tested is ornithine. Krebs had claimed that he chose ornithine just because it was available. As we indicated in Section 1, Krebs' claim is disputable and Holmes has speculated that Krebs chose ornithine because the metabolic fate of ornithine was an unsolved problem. At present KEKADA chooses ornithine just because it is available, but it is possible to make KEKADA to follow the other scenario by keeping 'metabolic fate of ornithine' as a sufficiently interesting problem on the agenda. The experiment shows that ornithine produces little urea; ammonia alone produces urea at about the expected rate; but ornithine and ammonia together produce urea at about double that rate, which is much above the expectations. This result is noticed

as a surprise.

3.3. Simulating Determination of Scope

The ornithine effect now becomes the focus of attention. It is a common chemical strategy, if a surprising phenomenon is observed, to see if its derivatives and substances similar to it also exhibit the same phenomenon. The idea is that it is more productive first to determine the scope of the phenomenon and then to think about the specific mechanism of the reaction.

The hypothesis generated at this point is that the ornithine effect may be common to a class of substances similar, in one way or another, to ornithine. Using the system's general heuristics, four possibilities are generated for substances that may exhibit the ornithine effect: (1) certain carboxylic acids, (2) certain amino acids, and (3) certain alpha-amines.

Using the same heuristics as before, a whole series of experiments is carried out with such substances, none of which, except control experiments with ammonia, produce much urea. These outcomes produce low confidences in all of the above possibilities and indicate that the ornithine effect may be specific.

3.4. Simulation of Reaction Path Discovery

After the experiments began to indicate that the ornithine effect was specific, Krebs must have entertained some hypotheses regarding what the ornithine effect meant. Catalysis is one such possibility. Here, the historical account by Holmes leaves some questions unanswered. It is not clear how seriously Krebs considered the possibility of catalysis right from the beginning and at what stage he started considering it seriously . Given the uncertainty about how seriously he considered various alternatives at this stage, we decided to allow the user to make a choice between various hypotheses at this stage. This allows KEKADA to make the discovery in various different scenarios. Presently we will be describing one such scenario.

At this stage, just after the phase of determining scope is over, KEKADA has failed to identify a class of substances all of which would exhibit the ornithine effect. Without such guidance, the number of possible reaction paths is large and the system is able to generate only very, incomplete process descriptions that are viewed only as vague possibilities. These hypotheses are created at a

higher level of abstraction, where all the details need not be specified. The possibilities include:

1. Ornithine may be donating a carbonyl group to urea.

2. Ornithine may be donating an amino group.

3. Ornithine may be acting as a catalyst.

4. Ammonia may be donating an amino group.

5. Ammonia may be acting as a catalyst.

When dealing with an unknown phenomenon, KEKADA converts various facts disclosed by the experiments and by other work in the literature into clues. (By a clue we mean a hypothesis that has a high enough confidence to be considered true.) . Here two clues are known at the outset. First, since ornithine and ammonia produce much more urea than either produces by itself, it is noted that "there is mixed action of both inputs." From this it may be inferred that one of the inputs may not be a sole source of the urea in the absence of another substance. Second, it is noted from chemical structure that ornithine cannot produce urea by direct reaction. This creates the clue that an intermediate substance exists.

Besides generating these hypotheses, the system notes certain facts as related to the surprising event. One of the related facts is:

1. Arginine produces urea and ornithine. This fact, known from the literature, is considered relevant because two substances, urea and ornithine, are common between this reaction and the surprising phenomenon.

At this stage, the system considers the following alternative actions:

1. Studying one of the related facts to generate new hypotheses that would, in turn, suggest new experiments.

2. Performing experiments as directed by the hypotheses. Since the hypotheses under consideration do not all constitute concrete and complete descriptions of processes, these experiments are aimed at modifying confidences in the hypotheses and refining them.

The choice(INT) among these alternatives is made by interaction with the user. In this scenario the user, for some reason, feels that the catalyst possibility is not likely at all. First, the decision(INT) is made to determine the source of the amino group in urea. Experiments establish that this is the ammonia. This rules out the possibility that ornithine could be donating an amino group.

Next, it is decided(INT) to study if the fact that arginine produces urea and ornithine is related to

the surprising phenomenon, and, if so, in what way.

First, a number of hypotheses about the relation are generated from the clues, the surprise, and other knowledge. Two possibilities are considered. The first is that arginine belongs to a class of substances that has the ability to produce urea. The second possibility is that arginine is an intermediate. Confidence in the first possibility was reduced by experiments on various guanidino compounds that produced no urea. For reasons that are not clear to us, Krebs did not consider the second possibility very seriously at this point, and we did not permit KEKADA to explore it very much. KEKADA carries out an experiment to compare the rate of production of urea from ornithine and from arginine.

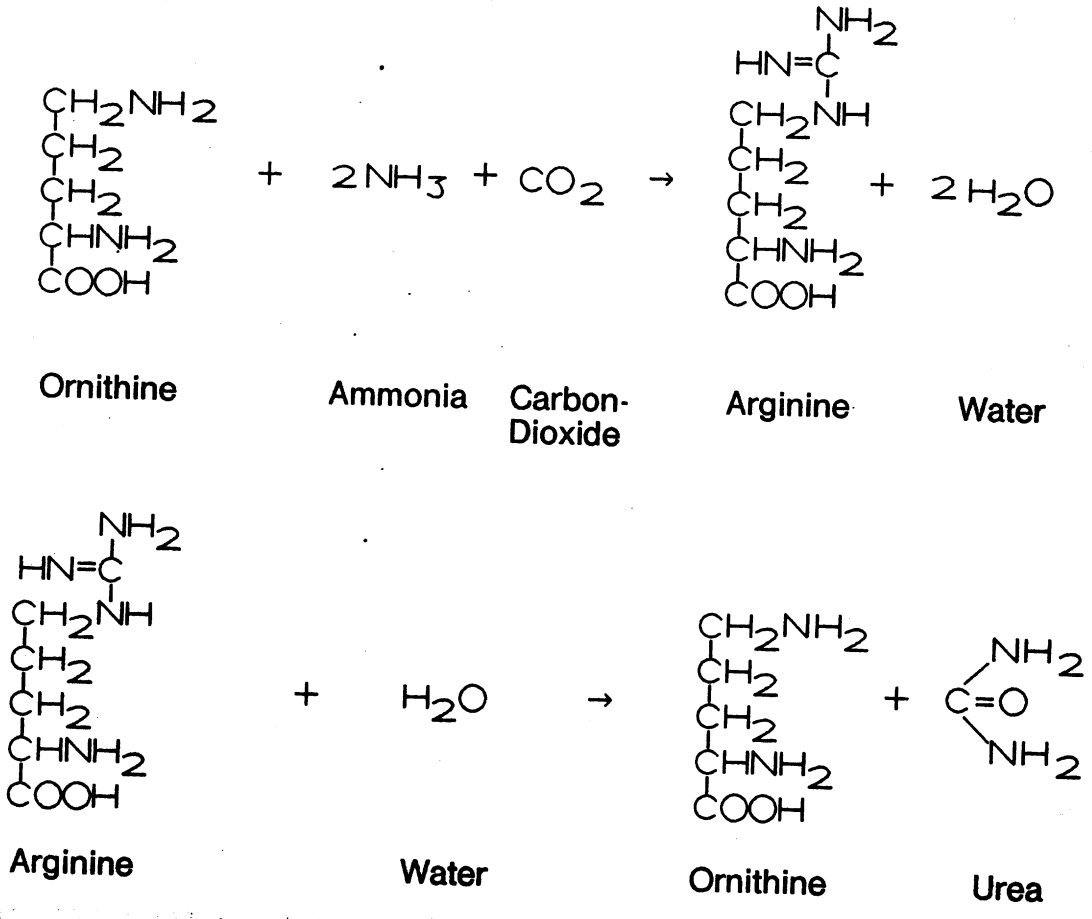


Figure 6. Ornithine as catalyst

Next, the system decides (INT) to carry out an experiment to find out whether ornithine is a

catalyst. In this experiment, 25 molecules of urea are formed for every molecule of ornithine used. This proves conclusively that the ornithine is not consumed in the reaction, but is a catalyst. Later it is concluded that arginine is an intermediate in the catalytic reaction.

3.4.1. Discovery of Citrulline as an Intermediate

On chemical grounds, KEKADA concludes that the conversion of ornithine to arginine could not proceed in a single step and decides to pursue the goal of finding the intermediate. It then creates possible candidate substances which are structurally intermediate between the inputs and outputs of the reaction producing arginine from ornithine. For each candidate substance, it evaluates the plausibility of its serving as the intermediate substance. Citrulline is the clear choice preferred by reaction-balancing heuristics. Besides, the system has the knowledge of Ackermann's work in which he showed that citrulline can be produced by biological action from arginine. Therefore it concludes citrulline is an intermediate substance in the reaction that produces arginine from ornithine. The reaction pathway it knows at this stage is shown in Figure 1.

4. Generality of the Simulation Program

In the introduction, we argued that Holmes reconstruction of Krebs' discovery of ornithine cycle is reliable data on which to build a theory of discovery. Now if we compare the course of work of Krebs with that of KEKADA, we find that there are only minor differences, which can be explained by focus of attention shifts[2] and small differences in the initial knowledge with which KEKADA and Krebs started. Apart from these differences, KEKADA follows the same strategy of experimentation as Krebs and its motivations for carrying out various experiments are the same as the motivations of Krebs, whenever these are indicated by evidence in the diaries and retrospective interviews. As KEKADA accounts for the data on Krebs' research, it constitutes a theory of Krebs' style of experimentation. Next we must ask that how general this theory is.

(1) KEKADA contains many general heuristics that are applicable in a large number of situations. Figure 7 shows that KEKADA has 31 domain-independent and 33 domain-specific heuristics. The

[2] A slightly more elaborate hypothesis evaluation system could explain a few differences in the order in which KEKADA and Krebs carry out their experiments.

CATEGORY OF HEURISTICS	DOMAIN-INDEPENDENT	NO	DOMAIN-SPECIFIC	NO
PROBLEM CHOOSERS	PC0-8	9		
PROBLEM GENERATORS	PG1	1		
DECISION-MAKERS	DM1-4	4	DM5-10	6
EXPERIMENT-PROPOSERS	EP1,EP6,EP7	3	EP2-5 EP8	5
EXPECTATION-SETTERS	ES1-4	4		
HYPOTHESIS-GENERATORS	HG1,3,8	3	HG2,4,5,6,7	5
HYPOTHESIS-MODIFIERS	HM4-5	2	HM1-3,B1-11,HM6	15
CONFIDENCE-MODIFIERS	CF3,CF4,CF5	3	CF1,2	2
HYPOTHESIS/STRATEGY CHOOSERS	HSC1 HSC2	2		
BACKGROUND KNOWLEDGE			DOMAIN-SPECIFIC	
TOTAL		31		33

Figure 7. General heuristics in KEKADA

domain-independent heuristics are some that scientists in various disciplines continue to use in making discoveries. Of domain-specific heuristics, DM5 to DM8 are actually applications to chemistry of more general domain-independent heuristics. Of the other domain-specific heuristics, , for all except B*, DM9 and EP3 we have historical evidence (Baldwin, 1979; Fruton, 1972; Luck, 1932; Holmes, 1986) that they were in common use in the study of metabolic reactions in biochemistry in early 20the century, before 1931 and for some years. Thus they constituted accepted domain-specific strategies which a newcomer like Krebs was likely to know after a brief introduction of the field. The B* heuristics are also quite general in their applicability , for they can be used to balance not only the reactions in this discovery, but many other reactions as well.

(2) As is shown in in the log in the section 3, most of KEKADA's heuristics are used a number of times in the particular scenario given. EP8, HG2, HG7, and HM1 are the only domain-specific

heuristics that are fired only once, but their potential utility in other research situations is clear.

(3) Some of KEKADA'S heuristics were also used in different forms by AM a mathematical discovery system, in the course of a wide variety of discoveries (Davis & Lenat, 1980).

(4) Thanks to Holmes (1986) , we now have data on a second major discovery of Hans Krebs, that of glutamine synthesis. A hand-simulation indicates that, the path Krebs followed there is wholly consistent with the current theory. We will report in more detail on the KEKADA simulation of the research on glutamine synthesis in another paper.

These considerations show that although KEKADA was hand-crafted to fit our knowledge of the procedures Krebs used in his discovery of the urea cycle, the structure and the heuristics it embodies constitute a model of discovery of wider applicability.

5. Conclusions

The immediate goal of the research reported here was to model as concretely as possible the heuristics Han Krebs employed in his discovery of the urea cycle. This was viewed, in turn, as a first step toward characterizing the heuristics used by scientists for planning and guiding their experimental work.

A number of very fundamental questions can be addressed if we are able to obtain a clear picture of the heuristics guiding particular discoveries, especially if that picture is sharp enough to permit us actually to simulate the discovery process. How specific are the guiding heruistics to the precise domain of the research problem? Conversely, which of the heuristics are applicable to other problems in the same discipline or even in other, distant, scientific disciplines. To what extent are the strategies of experimentation idiosyncratic to a particular scientist, arising out of his special knowledge, skills, and interests? To what extent are they based specifically on the current state of the art in the research problem domain? To what extent do they represent general strategies of problem solving search?

Our examination and simulation of the history of Krebs' discovery show that answers to these kinds of questions can be found. For example, we were able to show that nearly half of the heuristics Krebs used were quite general, being relevant not only beyond the urea synthesis problem, but beyond chemistry to a wide range of research situations. On the other side, we found that Krebs'

choices of problem and technique were much determined by the special opportunities provided by his training in Otto Warburg's laboratory. The tissue culture method, acquired there, was his "secret weapon," his source of comparative advantage.

The relative generality of KEKADA, and the ease with which it can be provided with knowledge and heuristics specific to a particular research domain allow us to view the control structure of KEKADA and its domain-independent heuristics as a model of scientific experimentation that should apply over a broad domain. We have already found that it can give a good account of Hans Krebs' research on glutamine synthesis, and we are currently applying it to other research problems as well.

Computer programs like BACON provided sets of processes that were shown to be sufficient for inducing numerous scientific laws from data. The present research carries our understanding of scientific discovery several steps further, by providing a detailed account of the successive steps in the discovery process, as well as showing how it reaches its final product.

The elucidation of the step-by-step progress of Krebs toward the discovery of the urea cycle shows the discovery being produced by a whole sequence of tentative decisions and their consequent findings, and not by a single "flash of insight." i.e., an unmotivated leap. It would appear that whenever we are able to build our models of the discovery process on detailed data, like that provided by Holmes in this instance, scientific discovery becomes a gradual process guided by problem solving heuristics similar to those used in other intelligent human endeavors. This conclusion will have to be tested, of course, with the data for many more instances of discovery before we can assess the generality of the model of experimental research provided by KEKADA. We are now undertaking a number of such addition tests.

References

Anzai,Y.and Simon,H.A. (1979). The theory of learning by doing. *Psychological Review, 86,* 124-140.

Baldwin,E. (1947). *Dynamic Aspects of Biochemistry.* New York:The Macmillan Company.

Brownston,L.,Farrell,R.,Kant,E.,and Martin,N. (1985) *Programming expert systems in OPS5: an introduction to rule-based programming.* Reading,MA:Addison-Wesley.

Buchanan, B.G. and Feigenbaum, E.A. (1978) Dendral and Meta-Dendral: Their application Dimension. *Artificial Intelligence, 11,*5-24

Davis, R. and Lenat, D. (1980). *Knowledge-based systems in Artificial Intelligence.* New York:McGraw Hill.

Friedland,P. (1979.) Knowledge based exeperiment design in molecular genetics. *Proceedings of the Sixth International Joint Conference on Artificial Intelligence,* 285-287.

Fruton, J.S. (1972). *Molecules and life: historical essays on the interplay of chemistry and biology.* New York:Wiley-Interscience.

Holmes, F.L. (1980) Hans Krebs and the discovery of the ornithine cycle. *Federation Proceedings, 39*(2), 216-225

Holmes,F.L. (1986) Personal Communication.

Langley,P.,Simon, H.A., Bradshaw,G.L., and Zytkow, J.M. (1987) *Scientific Discovery: Computational explorations of the creative processes.* Cambridge, MA :MIT Press

Lehninger,A.L. (1982). *Principles of Biochemistry.* New York: Worth Publishers.

Loffler, W. (1920). Zur Kenntnis der Leberfunktion unter experimentell pathologischen Bedingungen. [On the Knowledge of the Function of Liver under experimental pathological conditions] *Biochemische Zeitung,* 112:164.

Luck, J.M. (1932). *Annual review of biochemistry.* Stanford,CA:Stanford University Press.

Shrager, J. and Klahr, D. (1986). Instructionless learning about a complex device:the paradigm and observations. *Int.J.Man-Machine Studies,* 25,153-189.

Simon, H. A. and Lea, G. (1974). Problem solving and rule induction: A unified view. In L.Gregg (Ed.),*Knowledge and Cognition,* L.Erlbaum Associates, Hillsdale, N.J.

Simon, H.A. (1977). *Models of discovery.* Boston:D.Reidel Publishing Company.

I. Glossary

Alanine: $CH_3CH(NH_2)COOH$, is the simplest of the optically active amino-acids.

Ammonia: NH_3

Arginase: Arginase is the enzyme that catalyses the hydrolysis reaction in which arginine produces ornithine and urea.

Arginine: See figure 6 for the chemical formula.

Cysteine: This amino acid has chemical formula $CH_2(SH)CH(NH_2)COOH$

Cadaverine: $H_2N(CH_2)_5NH_2$

Guanidino: The Guanidino group is characterized by (NH_2--C(NH)--NH--). Arginine and creatine are examples of guanidino-bases.

Ornithine: See figure 6 for the chemical formula.

Perfusion method: In the 1920s, perfusion was one of the methods used to study experimentally the metabolic activities occurring in an organ. In the perfusion method, the organ under study is artificially provided with an independent circulation, driven by a mechanical pump, of blood of an individual of the same species or of certain physiological salines. The organ is thereby maintained under conditions very close to normal physiological conditions.

Lysine: This is the next higher homologue of ornithine. The chemical formula is $H_2N(CH_2)_4CH(NH_2)COOH$.

Tissue-slice method: In this method the experiment is carried out with thin tissue slices. Provided certain conditions are fulfilled, these slices will survive for some hours, apparently in a manner that closely approximates the physiological. Slices are easy to prepare and manipulate. The size of the avarage cell is such that the proportion of damaged cells to undamaged is very small, and the debris of the damaged cells can be removed by washing.