

**NOTICE WARNING CONCERNING COPYRIGHT RESTRICTIONS:**  
The copyright law of the United States (title 17, U.S. Code) governs the making of photocopies or other reproductions of copyrighted material. Any copying of this document without permission of its author may be prohibited by law.

# **The Processes of Scientific Discovery: The Strategy of Experimentation**

**Deepak Kulkarni and  
Herbert A. Simon**

Departments of Computer Science  
and Psychology  
Carnegie-Mellon University

This research was sponsored by the Defense Advanced Research Projects Agency (DOD), ARPA Order No. 4976, monitored by the Air Force Avionics Laboratory Under Contract N00039-85-C-0134.

The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the US Government.

## Table of Contents

1. Introduction	1
2. The Ornithine Cycle	4
2.1. Importance of the discovery	4
2.2. Background of the discovery	4
2.3. Course of Krebs' Research	5
2.3.1. Overview	5
2.3.2. The Ornithine Effect	6
2.3.3. Determination of Scope	8
2.3.4. Discovery of Reaction Path	8
3. Description of KEKADA	9
3.1. Production system	9
3.2. Representation	10
3.2.1. Representation of Processes	10
3.2.2. Representation of Data	11
3.2.3. Representation of Confidence Measures	12
3.3. Processes	13
3.3.1. Strategy and Basic Heuristics	13
3.3.2. Problem-Proposers	15
3.3.3. Decision-Making	16
3.3.4. Experiment-proposers	17
3.3.5. Expectations	17
3.3.6. Experimentation	18
3.3.7. Hypothesis-generators	18
3.3.8. Hypothesis-modifiers	19
3.3.9. Hypothesis-removers	19
3.3.10. Confidence-modifiers	19
3.4. Subject-matter knowledge	20
4. Simulation of the Discovery of the Ornithine Cycle	21
4.1. Overview of the simulation	21
4.2. Simulating the Ornithine Effect Discovery	21
4.3. Simulating Determination of Scope	25
4.4. Simulation of Reaction Path Discovery	26
4.4.1. Scenario 1	28
4.4.2. Scenario 2	29
4.4.3. Scenario 3	31
4.4.4. Discovery of citrullin as an intermediate	31
5. Comparison of History with Simulation	31
6. Conclusions	32
7. Acknowledgements	33
I. Glossary	34

Strategy of Experimentation

### **List of Figures**

- Figure 2-1:** The Ornithine cycle
- Figure 3-1:** Two-space Model of Learning
- Figure 3-2:** Interaction of heuristics
- Figure 4-1:** Progress of KEKADA in the discovery
- Figure 4-2:** Ornithine as catalyst

# The Processes of Scientific Discovery: The Strategy of Experimentation

Deepak Kulkarni and Herbert A. Simon  
Carnegie-Mellon University

## 1. Introduction

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 [5], of Lenat on AM [10] and of Langley, Simon, Bradshaw, and Zytkow on BACON [2] 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 perforce 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 themselves came from was left 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 [9]. 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 of urea synthesis was a major event in biochemistry, and Holmes' reconstruction of the process from published papers, laboratory

## Strategy of Experimentation

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, 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 [1]. 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 practicable 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, accounts from diaries and laboratory notes, and accounts by observation.

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.

4. **Accounts from observations.** The discovery is chronicled by a historian (or the

## Strategy of Experimentation

scientist himself or herself) who observes the discoverer's activities over a long period of time, also obtaining additional information from discussions. Such accounts are very rare. June Goodfield's account of Anna Brito's research is an example [8].

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 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 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 the domain of recombinant RNA [6].

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 diaries 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. But of course even an account from logs may have major or minor gaps -- completeness cannot be guaranteed.

Accounts by observation may provide a much better idea of the progress of the work,

for the observer can ask questions about ideas and motives not recorded in the log. But this source is very rarely available.

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 simulates the process. 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.

## **2. The Ornithine Cycle**

### **2.1. Importance of the discovery**

We paraphrase here (with his kind permission) Holmes' [9] account of the discovery of the ornithine cycle. 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.

### **2.2. 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 (erroneous) 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.



## Strategy of Experimentation

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 was "impossible to prove experimentally which of the several theories of the reaction mechanism derived from the test tube processes was the one that occurred physiologically. "

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" [11]. This was the situation that prevailed, in 1931, when Krebs began his research on this topic.

### **2.3. Course of Krebs' Research**

#### **2.3.1. Overview**

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 he 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

## Strategy of Experimentation

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.

### 2.3.2. 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 of 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 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

## Strategy of Experimentation

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

## Strategy of Experimentation

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) provide no information as to why the ornithine experiment, which represented a departure from the current activity, was run at that particular time. It is possible to speculate about the reasons for the experiment, but we will leave the question unanswered here.

### 2.3.3. 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.

### 2.3.4. 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. Measuring the effects of concentrations on the rates of reactions was, of course, a standard technique. 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

## Strategy of Experimentation

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.

## 3. Description of KEKADA

### 3.1. Production system

The KEKADA system is implemented in the production system language OPS5 [4]. This version of OPS5 is, in turn, written in LISP.

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

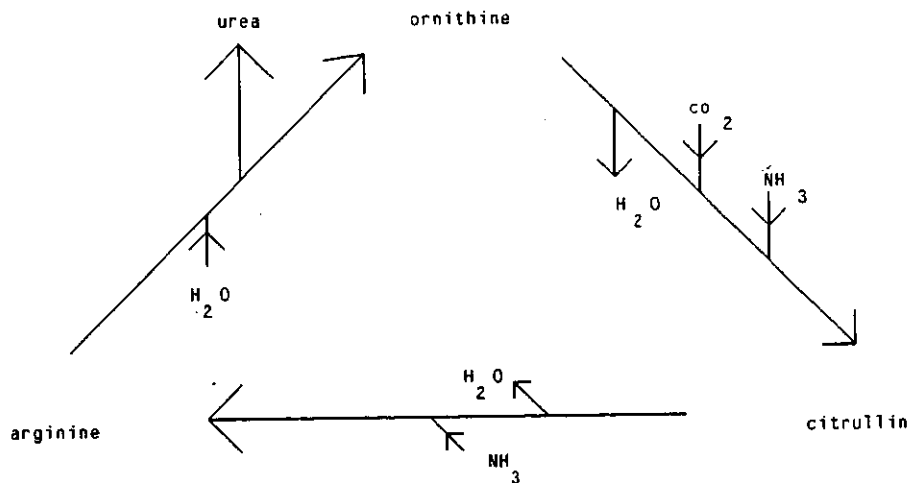


Figure 2-1: The Ornithine cycle

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.

### 3.2. Representation

#### 3.2.1. Representation of Processes

The discovery heuristics of the KEKADA system are stated as OPSS 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.

Conditions and actions can be written directly in OPSS, or, alternatively, as LISP functions. 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

## Strategy of Experimentation

than one such rule, then it chooses the one that is most specific.

### 3.2.2. 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 facts*, and *hypothesis*.

**Process.** Process elements 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, the class 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, and a status flag labeling it either as "done" or "not-done. "

**Supplementary Facts.** 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.

At present, a chemical substance is represented only by the compact form of the formula for the molecule: for example, the system knows that ornithine is  $C_5H_{12}N_2O_2$ , but does not know the exact bond or geometrical structure beyond this. The system also knows that ornithine is an amino-acid, specifically, an alpha-amino-acid, and that it has amino and carboxylic group. In the further development of the system, it may be necessary to represent exact chemical structures.

**Hypothesis.** Hypothesis represents a hypothesis or a possible description of how a phenomenon or process that has been noted takes place. Associated with a hypothesis is a

## Strategy of Experimentation

measure of confidence in the truth of the hypothesis.

The description for a hypothesis both states the hypothesis and provides an attribution of source. Thus, the possibility that "ornithine may be donating N to urea" may be represented by the elements: "description input1-donates element N", where input1 in the reaction is ornithine. The description slot may have such values as "catalyst", "donor", "indirect-metabolic-effect", and so on.

This representation for hypotheses works because productions can be written in terms of these descriptions. But there are problems with this representation, for the system cannot introduce new concepts beyond those already given to it.

### 3.2.3. Representation of Confidence Measures

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

1. Success: the number of experiments which have verified the hypothesis.
2. Failure: the number of experiments which have falsified the 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 us 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. "

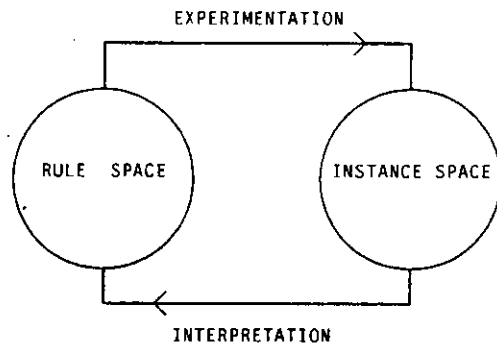
We convert the values of the attributes into numbers by assuming that each fact increments the appropriate attribute by one unit. That is to say, if a fact indicates that a hypothesis is probably false the implied-failure slot is incremented by one. This rough scheme seems to work satisfactorily for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



### 3.3. Processes

#### 3.3.1. Strategy and Basic Heuristics

The overall organization of KEKADA is based on the two-space model of learning proposed by Simon and Lea [12] shown in (Figure 3.1).



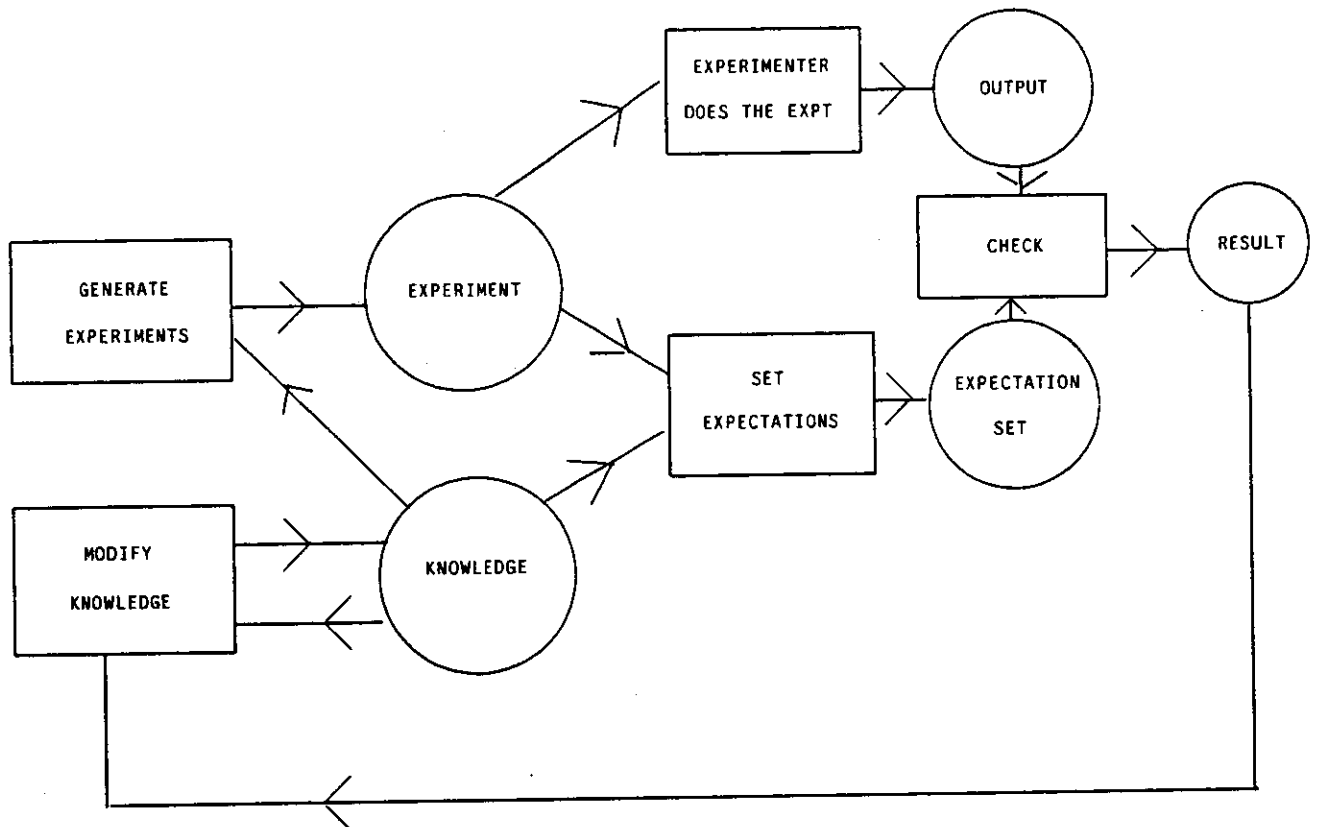
**Figure 3-1: Two-space Model of Learning**

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. The heuristics used to choose and guide these processes fall into several categories:

1. **Problem-proposers**, which propose new problems or subproblems on which the system can focus attention.
2. **Decision-makers**, which choose, for example, which of the various problems proposed by problem-proposer heuristics to work on.
3. **Experiment-proposers**, which propose experiments based on existing hypotheses.
4. **Expectation-setters**, which set expectations for the experiments to be carried out.
5. **Experimenters**, which carry out experiments.

## Strategy of Experimentation

6. **Hypothesis-generators**, which generate new hypotheses about unknown mechanisms or phenomena.
7. **Hypothesis-modifiers**, which modify the hypotheses on the basis of new evidence.
8. **Hypothesis-removers**, which remove the hypotheses when the confidence in them reaches a low enough value.
9. **Confidence-modifiers**, which modify confidences about hypotheses on the basis of the interpretations of experiments.



**Figure 3-2: Interaction of heuristics**

At any given moment (see Figure 3.2)<sup>1</sup>, a certain number of hypotheses with varying

<sup>1</sup>Figure 3.2 is a close analogue of General Rule Inducer schema proposed by Simon and Lea [12]

## Strategy of Experimentation

confidences are present in working memory. Experiment-proposers suggest experiments to be carried out on the basis of these hypotheses and the problem that is in the focus of attention. They may also suggest experiments that are likely to refine existing descriptions and to provide useful data. Expectation-setters set expectations about the outcomes of these experiments. Then the experimentation heuristics carry out these experiments (the results being provided interactively by the user of the system), and the results are interpreted by the confidence-modifiers. When applicable, the problem-proposers suggest new problems or subproblems to focus on. Decision-making heuristics decide which problem or subproblem to work on, and make other essential choices. When applicable, hypothesis-generators propose new hypotheses, hypothesis-modifiers alter existing hypotheses and hypothesis-removers remove the hypotheses with very low confidence from consideration. We will now discuss each of these classes of heuristics in more detail.

### 3.3.2. Problem-Proposers

The problem-proposers generate new problems and subproblems on which the system can focus attention. If at any given stage more than one problem or subproblem is proposed, the decision-making heuristics choose among them. The following heuristics are used by the problem-proposers:

1. If a new method is found to be useful for a given process, it may be worthwhile applying it to similar processes.
2. If the outcome of an experiment violates expectations for it, then note it as a *surprise* and make it the focus of attention.
3. If the goal is to study a reaction involving a class, then begin by studying the reaction for a particular member of the class. Thus, if the goal is to study how urea is formed from amino acids, start by studying urea synthesis from a particular amino acid.
4. If the focus is on a surprise, and there is a hypothesis that the reason for the surprising outcome may lie in a subprocess, propose to study that subprocess in detail.
5. If the focus is on a surprise, and there is a hypothesis that the surprising outcome may be common to a class of substances, propose a study to determine the scope of the phenomenon.

### 3.3.3. Decision-Making

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. For choice among biological processes, 2. For choice among substances, 3. For defining an initial ordering.

**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.

1. If the output of a process is not measurable, eliminate it.
2. If the typical rate of production of a process is significantly less than that of another process, prefer it.
3. If there is a process whose understanding is considered important by the discipline and another that is not, give preference to the important problem.
4. If there are no other criteria for choice between two processes, choose one of them randomly.

**For choice among substances:** The following set of rules is used for deciding which one of the given set of substances should be chosen for study.

1. If the cost of a substance to be tested is too high, eliminate it.
2. If a substance to be tested is not easily available, eliminate it.
3. If the cost of two substances is low and both are available, and they are being tested as being similar to a particular substance, then give preference to the substance that is closer to the given substance. (In the present implementation, a partial ordering is defined on various substances indicating their similarity to ornithine.)
4. If there is no other criterion for choice between two substances, choose one of them randomly.

**Defined priority:** Sometimes experience before the present research program was undertaken or the nature of the hypotheses define a partial order on the hypotheses that determines which hypothesis should be tested before which. For example, hypothesis that a given surprising reaction may be common to a class of substances is normally considered before other hypotheses. Work on this kind of a hypothesis is likely to be very productive. Correspondingly, in two cases the hypotheses in the system are considered in a pre-defined order. But we do not have the data on Krebs' previous experience. In the case where we have used a pre-defined order, it is possible that he actually used decision-making rules like

## Strategy of Experimentation

the ones proposed above.

In running this system for the urea example, in one case the choice was made by the user, because the biochemical heuristics Krebs used to make the choice are not clear to us. Interaction with the user allows the system to make the discovery of the ornithine cycle along different pathways. Introduction of additional biochemical heuristics may allow these choices to be made automatically instead of interactively.

### 3.3.4. Experiment-proposers

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

1. If there is a hypothesis that A produces C with B as an intermediate product, then carry out experiments on A and B, and compare rates of formation of C from A and B.
2. If 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.
3. If there is a hypothesis 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 in combination, measuring appropriate parameters to determine the quantity of C in relation to the quantities of A and B.
4. If there is a hypothesis that the reactant A in an experiment is a catalyst, then carry out the experiment over long periods but with very low concentration of A.
5. If there is a hypothesis that the reason for a surprising outcome may lie in an unknown substance, guess the substance to be a substance related with the process. i.e. ( substances that earlier experiments seem to have associated with the given process or the class of the process.) and carry out an experiment on it.
6. If the goal is to study a particular reaction in detail, carry out the reaction under various conditions.

### 3.3.5. Expectations

Expectations are set by three rules:

1. When the same experiment was carried out before at least once, 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.
2. When 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

## Strategy of Experimentation

knowledge of the investigator.

3. When 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.

### 3.3.6. Experimentation

In the current system, there are no experimentation heuristics. The outcomes of experiments are supplied interactively by the user.

### 3.3.7. Hypothesis-generators

The hypothesis-generators are ten in number:

1. 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.
2. 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.
3. 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).
4. If there is low confidence in a set of hypotheses created by the above rules, create hypotheses asserting which reactant donates which element to the output substance.
5. If there is low confidence in a set of hypotheses created by the above rules, find a related fact with high confidence (the heuristics for doing this have not yet been implemented) and hypothesize a relation using (5a) or (5b). (5a) If the related reaction and the given phenomenon both produce the same output, hypothesize a class containing both and predict that it will produce this output. (5b) If there is evidence for a hypothesis that the given reactant could be an intermediate, then create this hypothesis. (If the given reaction and the surprising phenomenon contain two common substances, then they may be related.)
6. 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.
7. If A is a catalyst, and there is a hypothesis that D is an intermediate product, and it is known that D produces the output and A, then conclude that A produces D, which in turn produces the output and A.
8. If the output from A and B is different than the sum of the outputs from A and B, then create hypothesis that "there is mixed action from A and B" otherwise

## Strategy of Experimentation

create the hypothesis that the effect is additive.

9. If a one-step stereochemical reaction from inputs to the outputs is not possible, then create the hypothesis<sup>2</sup> that an intermediate exists. Otherwise create a hypothesis that there is an one-step stereochemical reaction.
10. If the goal is to find an intermediate in a reaction and there is another reaction in which input B produces the same outputs as in this reaction or if B is produced by same inputs as in the case of the reaction, then consider the hypothesis that B is an intermediate in this reaction. [In work continuing on this problem, we are constructing a quantitative version of this heuristic. ]

### 3.3.8. Hypothesis-modifiers

The present version does not contain rules for the modification of hypotheses. A possible rule of this kind would be: If there is low confidence that a class C always exhibits certain phenomena, then specialize the class to C', and hypothesize that C' exhibits the phenomena. A complementary modifier would generalize a hypothesis to a larger class.

### 3.3.9. Hypothesis-removers

Hypotheses in the system are in 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.

1. If amount of effort spent on an existential hypothesis reaches a certain high value, make the hypothesis inactive.
2. If the number of experiments that falsify a given hypothesis reaches a certain high value (at present 1), make the hypothesis inactive.

Another rule that might be needed: If the amount of effort spent on solving a problem reaches certain high value , consider the problem not worthy of pursuing further.

### 3.3.10. Confidence-modifiers

Three rules modify confidences in the hypotheses that the system holds:

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

---

<sup>2</sup>KEKADA at present does not consider the possibility of an indirect metabolic effect, which Krebs considered. Later we plan to incorporate this in the system.

2. 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.

### 3.4. Subject-matter knowledge

When Krebs started his research in 1931, he already had a great deal of knowledge about biochemistry, urea synthesis and amino-acids. Correspondingly KEKADA has such knowledge when it begins. This knowledge takes two forms. Some of it is in the form of the domain-specific heuristics described above. Other knowledge is created by using 'make' statements before KEKADA is run. A 'make' statement creates a working memory element of specified type. These working memory elements constitute system's initial knowledge. Prior knowledge falls in 3 categories: knowledge about substances, knowledge about processes, and knowledge about previous experiments.

1. Knowledge about various 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 particular substance to be used in the experiments. Besides KEKADA knows the partial order relation stating which of two substances is closer to ornithine. In future version, KEKADA will compute the partial order relation using simple heuristics comparing the structure of the molecules.
2. KEKADA also has knowledge about certain chemical processes. This includes the inputs of the process, the outputs, the class to which the process belongs and some supplementary facts. When the exact place or condition under which the process takes place is not known, supplementary facts may give the various possible places or conditions where the process might be taking place if any are known. Also associated with the supplementary fact is the confidence that the process does take place at the place. The knowledge also includes various possibilities previously considered likely regarding where the process takes place.
3. Before Krebs undertook the research program which led to the ornithine cycle discovery, he certainly had read about the experiments others had carried out on the urea synthesis. When the first experiments were carried out, it is assumed that his expectations about the outcomes were set by the previous experiments. e.g. when he carried out his first experiment on alanine, expectations could have only been based either on (1) the previous experiments on alanine itself or some amino-acids or (2) some theory (e.g. amino-acids get completely oxidized in the formation of urea) which was supported by these experiments. Therefore the summary of these previous experiments is made available to KEKADA. KEKADA uses this knowledge only to set the expectations in the initial experiments.



## **4. Simulation of the Discovery of the Ornithine Cycle**

### **4.1. Overview of the simulation**

We turn now to an examination of KEKADA's simulation of Krebs' path of discovery of the urea cycle, described earlier. As in that earlier section, we divide our account into three phases: discovery of the ornithine effect, the determination of scope, and the discovery of the reaction path. Major stages in these phases are depicted in the diagram on the next page. We will mark with an asterisk (\*) steps that are taken by decision of the user, operating with the system in interactive mode. One goal of future research with KEKADA is to find means gradually to remove these asterisks by providing richer knowledge and heuristics that enable the system to make the decisions independently.

### **4.2. Simulating the Ornithine Effect Discovery**

The first task of KEKADA is to select a research problem. On the basis of its prior knowledge, it fixes on the tissue slice method and decides to apply it to a different biological process (i.e., one other than cellular respiration). It examines processes similar to respiration and evaluates these alternatives.

KEKADA takes into consideration various processes 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. With such detailed knowledge, general heuristics for

## Strategy of Experimentation

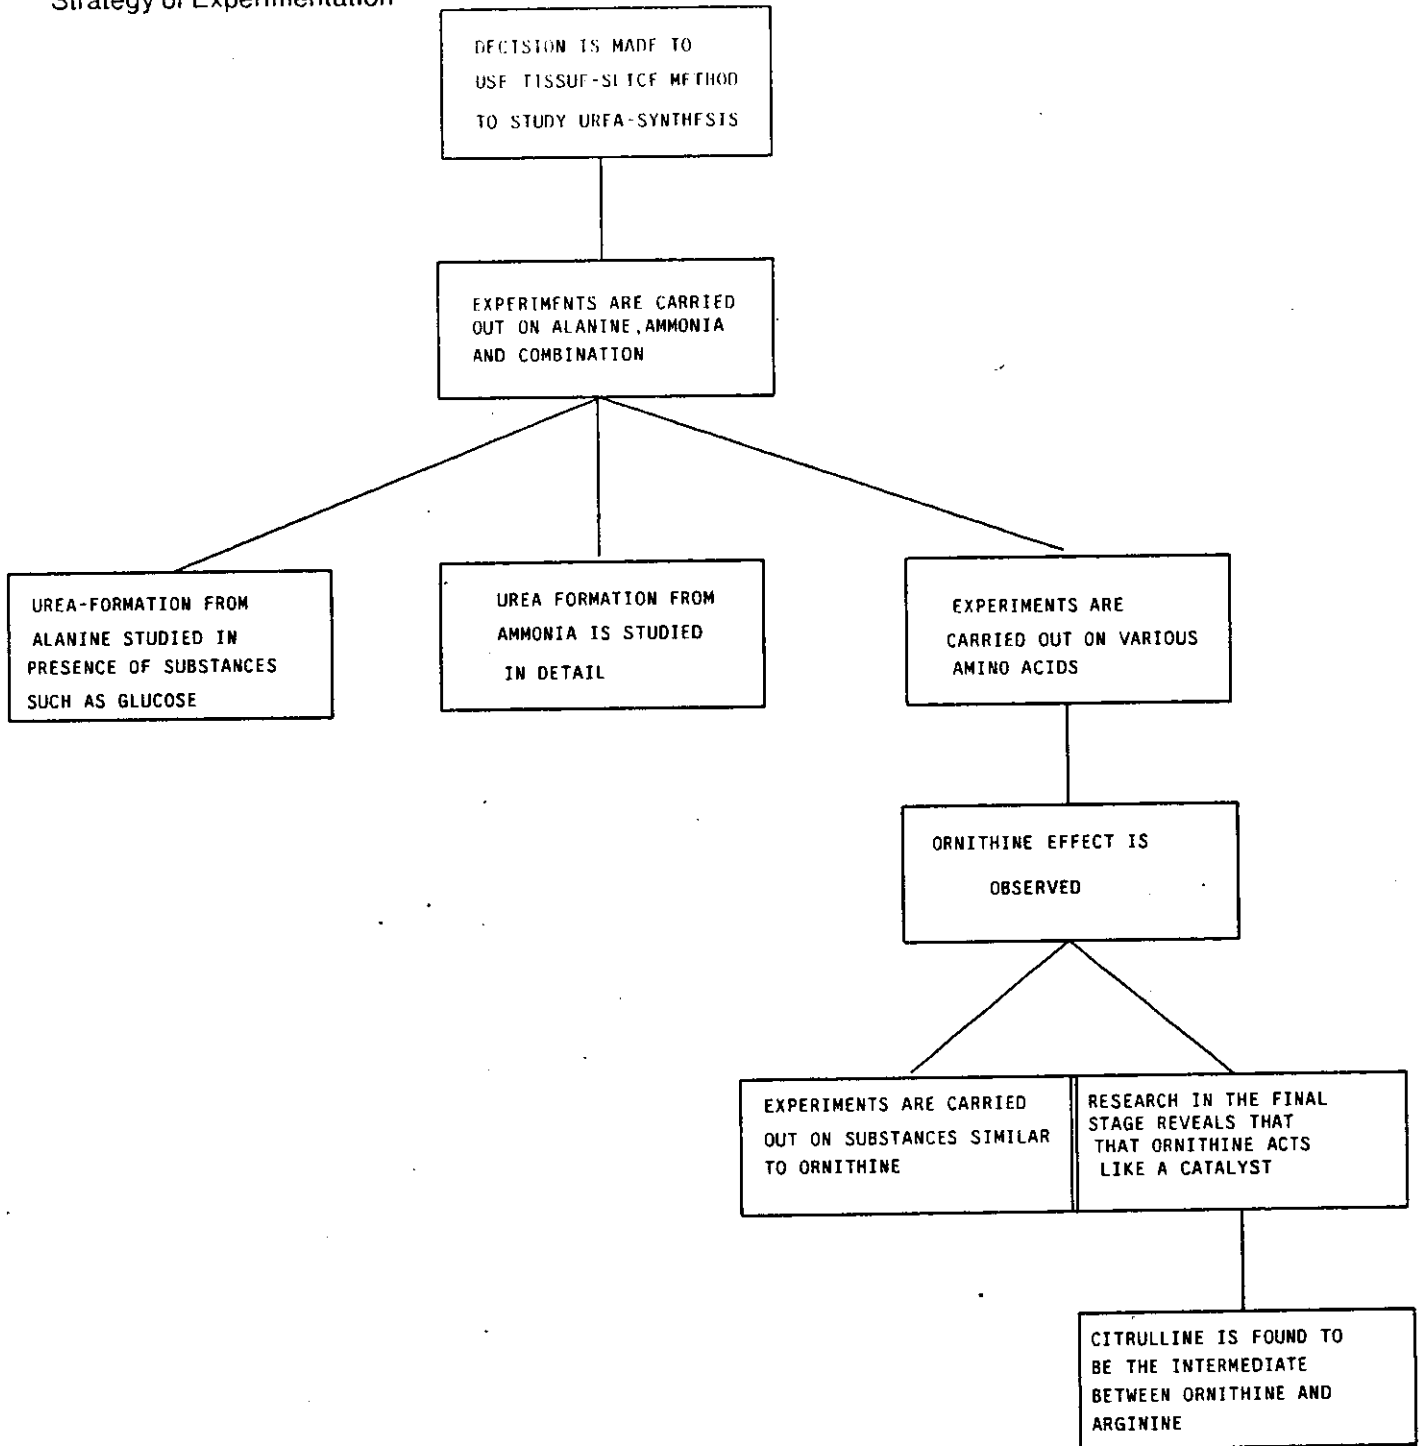


Figure 4-1: Progress of KEKADA in the discovery

## Strategy of Experimentation

problem selection might be strong enough to pinpoint the choice. Much more study will be needed before we will be able to estimate the sizes and shapes of these search spaces.

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 two possible mechanisms, which are stored in working memory:

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-making 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 quantity of the output, urea. The results of the experiment are provided by interaction with the user(\*). The user is asked for the output substance, the rate of production of the output, and the quantity of output produced. (It is assumed implicitly that one output will be measured.) In this experiment, the expected output was computed on the assumption of the complete oxidation of alanine.

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

## Strategy of Experimentation

amino acids. Certain possible explanations or hypotheses for this surprising result are now created by the hypothesis-generator and hypothesis-modifier heuristics. In the presence of appropriate facts of biochemistry, these rules produce corresponding hypotheses or modify hypotheses. Four 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 amino acids that will behave like alanine.
4. There may be a class of amino acids, 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

## Strategy of Experimentation

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.

On the basis of the third and fourth hypotheses that have been generated, KEKADA now repeats the original experiments with different amino acids. None of these experiments produce much urea from the amino acids, and the confidences in the various hypotheses are changed accordingly.

After carrying out these experiments, the expectation of output of urea from an amino acid is reduced, as is the expectation of increase in the production of urea from ammonia in the presence of amino acid. The next amino acid tested is ornithine. We have provided no special reason for inserting it into the list at this particular point, and, as we have seen, there is great uncertainty as to why Krebs tested it at just the moment he did. 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.

### 4.3. Simulating Determination of Scope

The ornithine effect now becomes the focus of attention. It is a common biochemical strategy that: if a surprising phenomenon is observed, try to see if its derivatives and substances similar to it also exhibit the same phenomenon. Thus 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, and (4) some derivatives of ornithine.

## Strategy of Experimentation

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 might be specific.

### 4.4. Simulation of Reaction Path Discovery

After experiments started indicating 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. Holmes says in his paper that "The clues Krebs had about the action of ornithine did not fit into the functional patterns associated with any of the previously identified classes of catalysts, " and implies that Krebs did not immediately consider catalysis as a serious possibility. However, on the other side of the balance, he began to pay a great deal of attention to the role of arginase (which will prove to be a secondary catalyst in the reaction), and he began a line of thinking that presently led him to the ornithine-arginine-ornithine cycle. It is not clear to us that this cycle was wholly unlike other catalyzed reactions known at the time, and we believe that further investigation into the biochemistry of the early 1930's is needed to clarify this step in the investigation.

As described below, KEKADA rediscovers the ornithine cycle in three different scenarios. In the first one, the possibility of catalysis is not given serious consideration before data start pointing to such a possibility. In the second one, the possibility of catalysis is entertained right from the beginning , but is not considered likely in preference to certain other possibilities. After spending some time on the other possibilities , the possibility of catalysis is considered a likely candidate and the experiments verify this fact. In the third scenario, the possibility of catalysis is considered very likely from the beginning.

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. "

## Strategy of Experimentation

These possibilities include:

1. Ornithine may be donating carbon to urea.
2. Ornithine may be donating nitrogen.
3. Ornithine may be donating both carbon and nitrogen.
4. Ornithine 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. 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 play any role as sole source of the urea in the absence of another substance. Secondly (\*) 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. The following are the related facts:

1. Arginase 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.
2. Lactase produces an increase in urea production from amino acid. "Increase in urea production" is common to this fact and the surprising ornithine effect.
3. Amino acids may be deaminated. Amino acid and ammonia are common to both 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. Carrying out experiments that would help in refining the description of the surprising phenomenon.
3. 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 not aimed at direct verification but at modifying confidences in the hypotheses.

The choice(\*) among these alternatives is made by interaction with the user.

Depending on which alternative the user chooses we enter different scenarios.

#### 4.4.1. Scenario 1

In this scenario the user, for some reason, feels that the catalyst possibility is not likely at all. First, the decision(\*) is made to determine the source of the nitrogen in urea. Experiments establish that this is the ammonia. This rules out the possibility that ornithine could be donating nitrogen.

Next, it is decided(\*) 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 created. The first is that ornithine and arginine belong to a class of substances that together have 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.

Next, one other related phenomena, the lactate effect, is studied. This study does not result in any possibility that these phenomena are related to the ornithine effect.

Next, the system decides (\*) to carry out an experiment to determine the source of the carbon in urea. This experiment is postponed to this stage because it needs to be carried out with very low concentrations of ornithine, and the equipment needed for the experiment only becomes available at this time. In this experiment, 25 molecules of urea are formed for every molecule of ornithine used. This proves conclusively that ornithine does not donate C or N to urea. As mentioned above there is already the clue in the memory that there is a mixed action of ornithine and ammonia. It is concluded that the ornithine is not consumed in the reaction, but is a catalyst. Later it is concluded that arginine is an intermediate in the catalysis reaction.



## Strategy of Experimentation

### 4.4.2. Scenario 2

In this scenario the user does consider the possibility of catalysis, but only after he considers a few other alternatives. First, the decision(\*) is made to determine the source of the nitrogen in urea. Experiments establish that this is the ammonia. This rules out the possibility that ornithine could be donating nitrogen.

Next, it is decided(\*) 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 created. The first is that ornithine and arginine belong to a class of substances that together have 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.

Next, one other related phenomena, the lactate effect, is studied. This study does not indicate any possibility that these phenomena are related to the ornithine effect.

Next, the system decides (\*) to carry out an experiment to find out whether ornithine is a catalyst. This experiment is postponed to this stage because it needs to be carried out with very low concentrations of ornithine, and the equipment needed for the experiment only becomes available at this time. 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.

Further, one of the possibilities that has previously been stored in memory is that arginine may be an intermediate. Since arginine produces ornithine and urea, it is concluded positively that arginine is an intermediate, and that the reaction path is the one shown in Figure 4.2.




Figure 4-2: Ornithine as catalyst

## Strategy of Experimentation

### 4.4.3. Scenario 3

The user from the beginning thinks that ornithine is likely to be a catalyst. KEKADA decides to carry out a verification experiment at once. Note that this does differ from the historical account. Later KEKADA pursues the goal of determining the unknown intermediate substance and concludes that arginine is an intermediate substance. This scenario is not very likely, because it does not explain why Krebs carried out some of the experiments in the last phase of the research nor the state of puzzlement that has been reported. Krebs thought for two weeks about the possibility of the relation of the arginase reaction to the ornithine effect and discussed it with his colleagues. If he was already looking for an intermediate that would produce urea and ornithine at this stage, it is very hard to give reasons why he shouldn't think arginine was the intermediate.

### 4.4.4. Discovery of citrullin 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 candidates as intermediates. Finally it concludes (\*) citrullin is the intermediate. The reaction pathway it knows at this stage is shown in the figure 2.1.

## 5. Comparison of History with Simulation

Some major points of comparison between the historical account of Krebs' discovery and our simulation require discussion. First of all, the system notes the low production of urea from alanine as a cause for surprise. It is not clear that, in historical fact, the outcome was regarded as very surprising, since Krebs continued to test a sequence of amino acids, and to perform other experiments that he hoped would throw light on the general situation. The surprise evoked by the system at this point perhaps makes the sequence appear more logical and determinate than it was. But this is a point that is not of much significance in the discovery, and the system here uses very general heuristics that would be appropriate in many contexts, and its course of action is reasonable, even if it may not be that followed by Krebs.

## Strategy of Experimentation

Second, Krebs, unlike KEKADA, shows some uncertainty about the validity of the tissue-slice method, and is reassured by the fact that healthy rats produce much more urea than ill ones (hence tissue slices less than healthy rats). The system requires no such reassurance. For a brief time Krebs also entertained previously known possibilities that cyanate may be intermediate and that pyrimidines could be donors. At present, KEKADA does not entertain such possibilities. But if KEKADA is supplied with these possibilities before it is run, (corresponding to the fact that Krebs knew of such possibilities before he started working on the problems) KEKADA would carry out some experiments entertaining these possibilities without needing any additional heuristics.

The last stage in Krebs' discovery is hardest to understand in detail -- especially since the log of experiments does not by itself give a detailed trace of the theory-building effort. If Krebs thought the ornithine path was catalytic after determining the scope of the urea reaction, why did he take so much time to conclude that arginine could be the intermediate and to arrive at the cyclic path? He was aware very early of the presence of large quantities of arginase in active liver tissue, and of the reaction by which arginine produces ornithine and urea. Since catalysis was not a new idea, why was the discovery of the cyclic mechanism regarded as so significant, and why were other scientists -- for example, some scientists working later on the citric cycle -- unable to conceive of such cycles? Further inquiry into the conceptual world of chemists at the time of Krebs' discovery is needed to answer these questions.

## 6. Conclusions

The goal of this research is to understand as well as possible the heuristics Krebs employed in his discovery of the ornithine cycle. Certainly our present version of KEKADA will need some revisions to bring it into closer conformity with the actual discovery procedures that Krebs used. A principal reason for interest in these heuristics is to understand to what extent they constitute a rather general set of heuristics for scientific discovery, to what extent they are specific to biochemical research, and to what extent they are idiosyncratic to Krebs' approach to scientific problems.

## Strategy of Experimentation

By constructing such simulations for a number of historical instances of important scientific discoveries we may hope to gain deeper insight into the discovery process and the relation of its heuristics to the heuristics of other types of problem solving and everyday thinking.

Holmes, in his account of the discovery of the ornithine cycle, remarks: "One might have expected, as I once did, that the cycle might have occurred to Krebs at some moment during his investigation, in the manner of gestalt switch." But looking at his historical reconstruction and our simulation program, the discovery appears as a gradual development, the steps connected in a logical order. This leads one to suspect that other discoveries that appear to result from a "flash of insight, " (i.e. an unmotivated advance) appear so only because we have not been able to examine the path of discovery in sufficient detail to expose the reasons that motivated each step.

## 7. Acknowledgements

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 report. Among others whom we would like to thank for comments on the manuscript or for information about the chemistry of our problem are Prof. John Modell, Prof. Ho, Prof. Hicken, David Steier, Yumi Iwasaki, Craig Knoblock, N. Ganesh, U. Shenoy and R. Ramanujan.

## I. Glossary

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

**Guanidino:** The Guanidino group is characterized by  $(\text{NH}_2\text{-C}(\text{NH})\text{-NH-})$ . Arginine and creatine are examples of guanidino-bases.

**Perfusion method:** In 1920s, the perfusion method was one of the methods used to study experimentally the metabolic activities occurring in that 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 kept under conditions very close to normal physiological conditions.

## References

- [1] Anzai, Y. and Simon, H. A.  
The theory of learning by doing.  
*Psychological Review* 86:124-140, 1979.
- [2] Langley, P., Zytkow, J., Simon, H. A., and Bradshaw, G. L.  
The search for regularity: Four aspects of scientific discovery.  
In R. S. Michalski, J. G. Carbonell, and T. M. Mitchell (editors), *Machine Learning, Volume 2*. Tioga Press, Palo Alto, CA, 1984.
- [3] Baldwin, E. .  
*Dynamic aspects of biochemistry*.  
The Macmillan Company, 1947.
- [4] Brownston, L., Farrell, R., Kant, E., and Martin, N.  
*Programming expert systems in OPS5: an introduction to rule-based programming*.  
Addison-Wesley, 1985.
- [5] Buchanan, B., Feigenbaum, E., and Lederberg, J.,  
*A heuristic programming study of theory formation in science*.  
Technical Report AIM-145, Stanford University Artificial Intelligence Project, June, 1971.
- [6] Friedland, P. and Iwasaki, Y.  
The Concept and Implementation of Skeletal Plans.  
*Journal of Automated Reasoning* 1:161-208, 1985.
- [7] Friedland, P. and Kedes, L.H.  
Discovering the secrets of DNA.  
*CACM* 28(11), 1985.
- [8] Goodfield, G. J.  
*An imagined world: a story of scientific discovery*.  
Harper & Row, 1981.
- [9] Holmes, F. L.  
Hans Krebs and the discovery of the ornithine cycle.  
In *Proceedings of the Symposium on Aspects of the History of Biochemistry*. FASEB, April, 1979.
- [10] Lenat, D.  
*AM: An Artificial Intelligence Approach to Discovery in Mathematics as Heuristic Search*.  
Technical Report AIM-286, Stanford Artificial Intelligence Laboratory, July, 1976.
- [11] Loffler, W.  
Zur Kenntnis der Leberfunktion unter experimentell pathologischen Bedingungen.  
*Biochem. Z.* :112-164, 1920.
- [12] Simon, H. A. and Lea, G.  
Problem solving and rule induction: A unified view.  
In L. W. Gregg (editor), *Knowledge and Cognition*. Lawrence Erlbaum Associates, Hillsdale, N. J., 1974.

## Strategy of Experimentation

- [13] Simon, H. A.  
*Models of discovery.*  
D.Reidel Publishing Company, 1977.